Marshall here, welcome back to The Realignment.
Today's episode is focused on a topic that has dominated the headlines of late, the rise of artificial intelligence.
It's been a lot of interest, hype, demerism, apocalyptic thought, basically any reaction you could possibly imagine to the leaps and bounds companies like OpenAI and their products like Chat and GPT-4 have made over the past year.
The big lesson I've learned from my coverage of the metaverse and cryptocurrency topics during the post-COVID tech boom and bust is that it's best to zoom out without getting too focused on specific claims and predictions when it comes to emerging technologies.
Instead, one should seek a broad outline and then fill in the frame over time which allows you to check and then just revise your thoughts as you go along.
My guest today then is Paul Chare, author of 4 Battlegrounds, Power in the Age of Artificial Intelligence.
Paul is a former Army Ranger and is currently Vice President at the Center for a New American Security.
We look at the AI topic through the lens of how Washington and the national security community are responding, especially as other countries and regions, both allied and rival, seek to make their own advances.
Hope you all enjoy the episode.
We'll be back with another Martial & Saga Q&A episode for Supercast subscribers this Friday so be sure to subscribe at realignment.supercast.com or click the link at the top of the show notes.
Huge thank you to Link and Never for supporting the show.
See you all next time.
Paul Chare, welcome to the realignment.
Thanks.
Thanks for having me.
The thing I'm obsessed with as an interviewer talks with a lot of people who publish books is when people publish books because the timeline is usually like a year and a half to two years before that perfectly coincide with shifts in the narrative.
So your book perfectly came out just as like the Silicon Valley space was pivoting from like Web 3 hype train to AI hype train, which makes you the right person to speak with because you didn't just jot this out last week to coincide, you were thinking about this deeply one, two, three, four, five, six years before we even got here.
So I just want to say, A, congrats to you, but B, kind of speak to that the hype direction.
How do you be substantive in a moment like that?
We'd love to hear that the personal side of that.
Yeah, thanks.
I mean, certainly, you know, as I say, it's better to be lucky than good.
The timing has been great, particularly with the chat GPT moment and all of the excitement we've seen around generative AI and of course, the gold rush we're seeing now into investments in generative AI.
I've been working on this project for about five years, started working on this when there

was all of this interest and excitement coming out of AI research labs about deep learning.
And there was really a not at all interest in the national security community, which was really far behind this.
But I, you know, like many people working on AI, there was this feeling of like there's something really interesting happening here, something really exciting.
I don't know where AI is going, but it's going to be big.
I think we're seeing that play out and it's, it's really exciting.
It's great.
I'm glad that the book is coming out at a time when hopefully people are going to be really interested in the topic.
Yeah.
Here's what would be helpful then.
Let's just do a kind of place setting, throw a lot of terms, chat GPT, machine learning, obviously this is all under the bucket of AI.
Just like position us like if I had a terminology level, because these are referring to different phenomenon.
There's different tech, different aspects have more national security relevance than other ones.
Just set the table that way.
Sure.
So I think it's most helpful to think of the term artificial intelligence as a field of study like chemistry or civil engineering, where the field of AI is about trying to design machines that are intelligent.
Now we'll set aside maybe for the moment like what is intelligence, turns out that's really difficult and AI researchers do not agree at all.
But in practice, what AI researchers have done is basically pick really hard tasks that we think require intelligence and then try to build a machine that can do that and then see where that gets you.
And within this field of AI, there's a whole variety of different techniques and approaches that have come in and out of fashion over the years.
We are now a decade into the deep learning revolution, and that's been an explosion in AI since 2012, and that in particular been one type of AI called machine learning.
So older model AI systems were rule based systems that were based on a set of rules that came from human experts.
Great example of this is a commercial airline autopilot.
It's got a set of rules for what the airplanes should do in any given circumstance based on the sensors coming in, the information, the aircraft, and it's rules that come from conversations with computer scientists and pilots and aerospace engineers.
And those rule based systems, those expert systems that draw on human expert knowledge are really good.
We use them all the time.
We use them in things like tax preparation software.
We just don't think of them as AI.

So there's this phenomenon where AI is like the weird stuff that doesn't quite work yet.
And then as soon as something works, we just call it software.
And machine learning works totally different way, where instead of being based on a set of rules, the algorithms are trained on data.
And you could take data sets, in some cases, massive data sets and train these algorithms on data.
And they actually learn from the data to identify patterns the same way that we do.
And that's a really powerful technique.
It's good for things like text generation, like generative models like chat GPT, or generating images like some of the art generators that we've seen explode onto the scene in the last year, as well as classifiers that can identify objects or faces based on images, identifying subtle patterns.
There's a whole slew of different applications of this approach.
And in particular, a type of machine learning called deep learning that uses deep neural networks, which is a connection as paradigm that's loosely modeled on how brains work.
And it's a really powerful technique, the data requirements are massive.
So models like chat GPT use hundreds of gigabytes of data.
But it's a very powerful approach.
And it's one where continue to see investment in progress in AI.
So the way you explain that, I'm curious about your opening point that for a long time, the national security community wasn't interested in this space, in this direction, I think is this hearing that articulation, you would just think, OK, obviously there are going to be some deep implications to this.
Kind of explained the, like, I hate this cliche, the gap between Silicon Valley and Washington DC there.
There are many, there are many.
And this issue particularly, there's just an information gap where people in DC often are, I find, several years behind the cutting edge of what's going out in the research space.
There's maybe, you know, like a five year gap in just understanding what's happening.
And so I found several years ago, I'd go talk to scientists and engineers who were working on frontier AI systems at labs like DeepMind and OpenAI, and you could see the AI progress if you were paying attention.
It was remarkable with things like AlphaGo beating the top humans ago.
Huge breakthrough in AI, and in particular in machine learning.
But people in the national security world really weren't paying attention yet.
That's changed.
And one of the big things that's really been a catalyst in the national security world has been China.
China a few years ago declared its intention to be the global leader in AI by 2030.
And that got people in Washington to sit up and pay attention more than any other research breakthrough because even if they're not exactly sure what AI is, they definitely know they don't want China to be ahead.
You know, it's interesting.

What do you think about at a conceptual, almost historical level, the idea of countries, not just China, but the U.S. countries in Europe deciding they are going to, at almost governmental level, own a technology?
I'm not sure that's how the first industrial revolution happens, but that's kind of like a bottom-up organic phenomenon.
Like, how do you think about this idea of countries in the 2020s trying to own a space like that?
Well, I think the industrial revolution is a great paradigm for what we're seeing in AI.
In the sense that it's a very diffuse technology, it has many, many uses.
And so while the industrial revolution led this process of industrialization across society, that we saw accelerated economic productivity and military power and its huge transformational effects on society, AI is likely to do the same.
And so AI is not a technology where a government can bottle it up.
It's like the opposite of stealth technology that was invented at secret defense labs.
That can be controlled and has been slow to proliferate.
AI is coming out of the commercial sector.
The proliferation is very, very rapid.
We're seeing, in some cases, just a few months from a cutting-edge research breakthrough, generative image models that do art generation, a great example of this, to then things go open source.
And there's an open source model, like in the case of art generation, stable diffusion, that anyone can get access to and download and use and modify.
So no one's going to be able to own that technology, but it is possible to be in a leadership position in shaping how the technology evolves, in developing it, and then applying its useful applications and shaping the governance for this technology.
Sort of what are the laws and norms and ethical frameworks for how it's used?
So I think the next question would basically come down to, obviously, China, the Chinese Communist Party has had a lot of headlines relating to really contentious relationships between the government and tech companies.
From a pure, let you kind of be more independent in the 2010s to there being a aggressive crackdown.
A lot of the AI work that we're talking about or that users or listeners have experienced has really been private company-led in the United States.
Do you think the contentious relationship between tech and government in China is going to, let's say, complicate the ability to own the space as much, or is this just an issue that's been quote-unquote resolved?
Well, I certainly don't think it's been resolved.
And I think it remains an open question whether the Chinese Communist Party is going to be able to reassert its control over the tech sector, that the party sees has gotten a little too big for the britches of the last couple of years, and they're reining them in and whether the party can do that without killing the golden goose here.
China has seen tremendous advances in terms of economic growth over the last several decades.

It has been very effective in climbing up the value chain of technologies where now they have Chinese tech companies that are global leaders in technology and in AI, companies like Baidu, Alibaba, Tencent, SenseTime, iFlytech.
But the government's been very active in wheeling some of these companies in.
The government has sort of strangely been very proactive in AI regulation.
So while there are no laws inside China that can constrain the Chinese Communist Party, the way there are laws in other countries about constraining what the government might do here in the US or in Europe, for example, the government in China has been very proactive in putting in place laws that are going to constrain what Chinese tech companies can do.
So historically, countries have had to face a choice, often called the dictator's dilemma, about whether they can embrace economic growth and openness and opportunity.
But with that has historically for other countries, the liberalization of economic opportunity has led over time to more political liberalization or countries can constrain freedom, constrain the free flow of ideas and basically head down the path of North Korea.
China has been able to thread that needle and they've been able to see unprecedented economic growth while the Chinese Communist Party keeping an ironclad grip on power.
But I think it remains an open question whether they'll be able to continue to do that or this crackdown will stifle Chinese innovation and growth going forward.
And this is where the Soviet example is so fascinating because obviously the Soviets missed the boat on computing.
They missed the boat when it comes to the internet, but all of the various issues with the Soviet economic system did not prevent them from leading early space program.
The various aspects of it weren't a constraint there.
So to what degree do you think a closed, top-down economic system has either an advantage or disadvantage when it comes to the AI technology specifically?
There's elements of both.
And this is what I think makes analyzing the political economy say the US versus China difficult, but also important to get down into the details.
So if you look, for example, at surveillance technology, China has half of the world's one billion surveillance cameras.
The Chinese Communist Party has been very aggressive in investing in a massive surveillance network, and not only is that helping them surveil and monitor their citizens in ways that are very troubling and dystopian from the standpoint of human freedom, but from the standpoint of lifting up the Chinese tech industry, it's been beneficial in accelerating Chinese progress through government funding, through data that then Chinese AI companies are collecting on Chinese citizens, and because Chinese companies now have the opportunity to actually deploy these systems out in the real world.
So for facial recognition, for example, China is going to be ahead of US companies, whereas in the US, there's been a grassroots movement against facial recognition at the state and local level.
Many states and cities banning the use of facial recognition by law enforcement.
And a lot of the big tech companies in the US, Amazon, Microsoft, IBM, they said they're

not going to sell facial recognition to law enforcement, either for the time being until there's better regulation or the case of IBM, they're just getting out of the business entirely.
And so that's one place where you can see China's model is going to accelerate progress in that very discrete domain of facial recognition technology.
Those Chinese companies are going to have a leg up.
It's not clear, though, that that is something that's applicable overall.
And in many cases, we see that the Chinese government's very active industrial policies, they're pretty wasteful and they're pretty inefficient.
They can get results, but there often comes at massive expenses.
Now the US government, of course, has been pretty hands off in terms of US private industry, but the government here in the United States is also starting to weigh back into industrial policy.
The US government now is going to spend $52 billion on subsidies for the semiconductor industry.
And that's pretty much an open experiment on how that's going to go and whether we're actually going to see the results that the US government wants to see.
You know, this is so fascinating.
I want to spend a few more seconds here.
So when you're talking about these different technologies, I went to the Soviet example with the space program because, you know, if we're going to find a technology category, it seems to be closest to be one where it's more of a decently militarized.
It's very top down by its very nature.
There's a lot of need to like pour right cash and resources into it.
That was one of the Soviet economic system was just built for versus something like the Internet, something like that.
So something I'm curious about is how do you see that dynamic playing out across all the different technologies we throw out?
I'm sure you're always seeing all these different things like, do you think there's, is there equivalent beyond just surveillance cameras then for a technology where you're like, oh, yeah, that seems to be one, which like either like the Chinese or the American or let's bring in the Europeans.
Maybe there's a technology that the Europeans are actually pretty well set up to dominate or at least develop strongly in.
Well, I think in a lot of applications, it's going to be a brutal level playing field.
One area where the U.S. and allies have a huge advantage is in the hardware that's used for applications in the semiconductors and the most advanced chips that are needed for these really capable systems like chat GPT.
So the semiconductor industry is very globalized, but it goes through a couple key choke points that give some countries outsized power over who gets access to the most advanced chips in the world.
The United States is one of those, but other countries include Japan, South Korea, Taiwan and the Netherlands.
And those countries have outsized influence over the physical hardware that AI runs on

because to train a large language model like chat GPT on hundreds of gigabytes of text, you need massive amounts of computing power.
And a lot of these large models run on thousands of specialized chips, graphics processing units or GPUs running for weeks at a time churning through all of this data.
Well, 90% of the most advanced chips in the world are made in Taiwan.
So Taiwan is the Saudi Arabia of computing hardware.
They have this really outsized geopolitical significance because of their role in making the most advanced chips.
But the technology that's used to make the chips themselves comes from other countries.
And it's the US, Japan and the Netherlands that control 90% of the global market for the technology that's used to manufacture chips.
And in fact, they've recently used this to hold China back in hardware.
So in October, the Biden administration put out new expert controls on the equipment that's needed to manufacture chips going to China.
So holding back China's domestic semiconductor industry, and more recently, we've seen the Netherlands, Japan have come on board to join the US in these restrictions.
And that's going to really hold China back.
You know, I'd love to go back to something you said earlier.
You were talking about in the category of surveillance technology, how the status quo is one where the Chinese are able to make more progress.
I want to unpack the word progress in this category because if we could take a step back, it seems that in most categories of technology, 99.9% of them progress sounds great.
Progress with my iPhone getting cheaper yet faster.
My Mac is better than ever.
Electric vehicles, et cetera, et cetera, et cetera, like consumers and just citizens in general like progress.
In the case of, let's say something like nuclear weapons, like more megatons for less, you know, dollars percent, I think that makes people feel uncomfortable.
And I think the same thing is true when it comes to surveillance technology.
So can you just interrogate the word progress and how it's clearly something that people are uncomfortable with, with AI, especially because in the popular imagination, when we think of the furthest extent of AI, we're definitely thinking of Skynet and Terminator.
And any progress that moves us towards that is obviously not good progress.
So just talk about progress in this context.
It's so different in other areas of tech.
Yeah, well, I think one of the things this highlights is that using technology, inventing technology is inherently a political act.
There's some implicit values that you have when you're designing and using technology.
What are you making it easier to do?
What are you making it harder to do?
So in this case, China is ahead of the U.S., if you want to say that, I'm building a massive dystopian surveillance state to monitor and control its citizens' behaviors.
That's not a race we want to be in.

So that's not one where we're like, oh, we'd like to have more surveillance cameras than them.
We'd like to have tighter control over the population.
And some of the things that China is doing are remarkable in how sweeping they are.
So it's not just surveillance cameras at train stations and airports and hotels and bus stations and public areas and sidewalks and public squares.
It is all of those things.
And the cameras are quite ubiquitous, they're quite obvious when you walk around major Chinese cities.
There's light pools with cameras all over them because, of course, the Chinese Communist Party wants you to know that they're watching you.
But it's also cameras being used for things like catching people jaywalking or using too much toilet paper and public restrooms.
So it's this really extensive paradigm of control over its citizens' behavior.
Because if you can control how much toilet people are using, then it's going to be easier to crack down on political dissidents, right?
That very intensive social and political control go together.
What I do think is important to think about leading on, would be great to have democracies leading on, is setting the norms and laws for how technologies are used.
Not to necessarily have better surveillance, but have a better model for how surveillance technology should be used.
So even more troubling than what China is doing is we're seeing the export of Chinese technology and social software, if you will, the laws and norms behind how China is using this technology.
80 countries around the world have now purchased Chinese policing and surveillance technology.
And they're using it domestically within their own countries.
And so having democracies come together to come up with an alternative model and say, here's a way to use some of these technologies to do legitimate things, like protect your border points and your entry, make sure that someone coming into your country is who they say they are.
There might be legitimate uses for police to crack down on crime.
But to do someone a way that protects individuals' privacy and human freedom is important.
And that's where we want to focus on leading.
Okay.
Here's a question, though.
What degree in this context can the U.S. actually start the, or actually to what degree can the U.S. actually set norms in the sense that if, let's say you're a, I don't want to engage in stereotyping.
Let's just say you're in a regime that is very friendly with the Chinese because unlike the United States, they're just not offering requirements or there's not the same degree of a conditionality, why would other countries that see the opportunity when it comes to top-down control, why would they say, okay, the U.S. is doing it, let's follow their model?
If they're not already like a Western-friendly democracy, it's basically what I'm getting

at.
Yeah.
So let's give us some examples about how some of these norms can diffuse.
I knew Bradford was talking about Europe engaging in a race to the top on regulatory standards.
And Europe's GDPR, their data protection regulation, is a great example of this.
Europe got out in front of the rest of the world, passed the GDPR, the Data Regulatory Regime.
Now, whether or not you like GDPR, it has been very influential in shaping what other countries are doing, including in China.
So when I visited China and I spoke with legal scholars, they were debating the GDPR and what it might mean for a consumer data privacy regulation inside China.
Now that's really focused on consumer privacy, so privacy of Chinese citizens from tech companies, not from the government.
That's not on the table for discussion inside China.
But GDPR gave Europe a first mover advantage when now everyone else around the world, whatever regulatory standard they come up with, they have to first think about the GDPR and how does that kind of affect them.
And a lot of companies then, they're going to choose to default, comply with GDPR, not just in Europe, but maybe elsewhere.
Technical standards are another place where this becomes a main issue, where engineers get together through these very technocratic organizations like the ITU and ISO to come together and talk about technical standards for all sorts of things, for wireless connections, for surveillance technologies to make sure the technology is going to be working together, they can pass data, they can communicate effectively.
And some of those are going to have impacts, not just on whether the technology works better, but how it works and whether it works in ways that might make it easier to protect people's privacy or make it easier to spy on people.
And that's a place where we've seen China being increasingly active in global standard setting.
And it's really important that we protect these standard setting bodies.
We make sure that they are technocratic first and foremost.
And the standards that we're adopting globally are ones that are going to protect privacy and freedom.
So, let's understand that then, because I think when most people think of things they wouldn't want AI to do, they wouldn't want AI to kill you.
Also, I think there are a bunch of listeners who now also don't want AI monitoring their public toilet paper usage.
Let's put those two to the side.
What does the actual gamut of, let's say, objectives actually look like when it comes to legal restrictions?
I don't even want to say restrictions, because I think that has a moral weight to it.
So, what does the legal gamut look like?
Well, we're seeing a whole host of new challenges rise up as we're starting to build these AI

systems.
So, we see already debates in the US about law enforcement using facial recognition.
We see debates about generative AI art.
Who owns that art?
That's going to be the case for both images and for text.
It might be, maybe not too far down the road, the case for inventions.
Who owns some new invention that an AI system comes up with?
It cannot be cooperated by the user.
What if you use the AI to do that?
What if you work with the AI?
How much do you have to modify it to have ownership over it?
It's also an issue when we come to think about the data that these AI systems are trained on.
This is a big issue for these generative art models.
They're trained right now on copyrighted data.
And so, the people that are training these models say that they believe that using copyrighted data to put in the training database is fair use under US copyright law.
But that's an open question.
And others are suing these companies saying that's not, that in order to use copyrighted data, you need to get permission.
Either way, we need to find some resolution to this issue and having clarity will help then unlock innovation.
So, everybody knows, if I wanted to be able to use this data, this is what I have to do to get there.
And those are just some of the things that we're running into now.
So I think it's worth keeping in mind that when we see powerful new technologies, we don't want to strangle those technologies, but we also want to use them thoughtfully.
And one of the big myths that can be really destructive here in the United States is this idea that digital technologies shouldn't be regulated.
Well, the only reason why we have clean air, water, safe food to eat, safe highways, safe air travel is because of government regulation of all of those areas.
In fact, in a number of states, even regulate tape recorders, we need two-party consent before recording.
So there's going to be lots of areas where AI is going to have to be regulated, too.
Yeah, I want to, we spend some time dunking on DC being behind the curve.
But I think you said something earlier that a lot of people in the tech world would just object to, but I think it's just true, which is that, like, your point that technology is ultimately like a political act, like how it's used, how it's governed, those different aspects.
Could you just speak to the, like, Silicon Valley side of the audience about how, at the end of the day, we're just not capable of having a conversation around AI that doesn't inevitably get political?
Because I was just saying, like, I've had plenty of VCs on this podcast at, like, very,

like, marquee firms who have said things like, seriously, DC, leave us alone, like, let us do our thing.
And the DC side of me is like, well, at the end of the day, if there's a political implication in what you're doing, it's not really even a choice, it's just actually a reality of the space there.
And so can you just talk about this dynamic?
I think it's important to keep in mind that if we're going to have an economy that's productive, we're going to have a society that's protecting social welfare and individuals being, we're going to have to regulate these powerful technologies in some ways.
And the people that are inventing them in Silicon Valley or elsewhere, and the people that are regulating them in Washington or in state capitals or in cities are going to have to work together.
So I understand that there's lots of frustration for people coming out of the tech sector with oftentimes just a lack of understanding and awareness in Washington.
I get it.
I live here in Washington.
There is a tech literacy problem.
But the best way to overcome that is to have greater dialogue and understanding.
Anything people in Washington are receptive to, okay, what are the approaches that we'd like to see?
But the answer can't be just let it rip.
That's not going to be good for society.
And I'm actually encouraged by things like debate surrounding chat GPT, that people are now able to see AI technology in at least one specific application and begin to envision all of the challenges that this is going to bring.
I suspect that issues like, are students going to use it to write essays is going to be the least of our problems when we think about generative AI.
I'm much more concerned about things like mass producing misinformation, massive amounts of text spam flooding the web, and how are we going to deal with that?
And we're going to have to find ways to have those on the tech side and those on the policy side.
Or as soon as people will say, you know, a suits and hoodies working together to solve these problems.
Yeah, that's a very, it was just shocked back into 2014 hearing articulation of it.
Okay.
So now we're in the other half of the show, I want to do a few quick bits and then get into your previous work and the actual four key elements that you're bringing to the table with the book here.
So firstly, this episode has not been doom and gloomy, but it's definitely like looking at the downsides of things.
Can you just tell me like why I should be excited about AI?
If we're comparing it to like electricity and we're comparing it to previous industrial revolutions, like the electricity answer is, hey, you could now stay up way later.

That's awesome.
Like what's just like a tangible way that you're without getting into too far like unhelpful science fiction imagery, like why should people be excited?
Also like putting aside, they've seen like they've seen like the, you know, tax generation, they've seen like the AIR, like what's like a monumental way you think someone should be excited about this?
I think AI overall is likely to accelerate economic productivity and scientific progress in ways that will have a whole sorts of valuable effects on society, whether that's in medicine, finance, transportation, or other areas, that there's all sorts of things where if we can build out data sets for what we want a machine to perform, and if we have a clear metric for what better performance looks like, we can train machines to do that.
We can offload tasks that humans are doing to machines now, which humans focusing on higher level tasks that maybe machines can't do.
And it's going to accelerate productivity and economic growth.
It's going to lead to a better society, ideally in many ways, a safer one overall, self-driving cars alone, we can get there to have safe autonomous cars that are better than people would save tens of thousands of lives on the in the US every single year.
And so there's tremendous advantages from AI.
Yeah, the book, our conversations about some of the scary stuff, but I wouldn't want that to take away from this idea that I think overall, AI has a lot of benefits in society.
And actually, I'm quite optimistic about that.
Yeah, I want to reverse the way you typically ask this question, then, instead of asking what keeps you up at night in this category, what doesn't keep you up at night in this category of conversation?
Oh, what is a false concern, I guess?
I think a lot of the concern that comes up in science fiction is this very anthropomorphic vision of AI, these AIs that like sort of one day becomes self aware and turn on us.
And I think it's actually not a very helpful way to think about AI.
The way this is going to be a little gloom and doom, I'm sorry.
The way that I think about it is probably worse, which is to say that a lot of the systems we're building, they actually don't think like people, which is not a good thing for us because it means that our mental model for how other intelligent systems interact, what aligns, you and I to have a conversation or listeners to listen to us and envision who we are and what we're thinking about.
We have sort of this mental model for another person, and we often put it on other objects.
We put it on our pets, people name their Roombas, and then we stick it on AI systems.
So you can interact with something like chat GPT and you're talking to it, and it can fool you into thinking that it thinks like a person and it doesn't.
What's going on under the hood is super weird and alien, and so I think a lot of the concerns that we see in science fiction are probably not what we're going to see, but we're going to have different challenges, and we're going to have to work hard to shed some of this anthropomorphic framing that we often have when we look at AI.
That's so interesting.

So even I'm thinking of like, you know, Isaac Asimov, is it the four laws?
How many laws are there in Isaac Asimov?
Three laws.
Three laws.
Three laws.
So don't kill people, and I can't remember, you know, the other two, even though it's simple, I should remember the other two.
But the kind of the point is, if I'm thinking about the way we articulate how you would protect yourself against AI, I guess you are thinking, we're worried, it's going to start asking itself, who am I?
What am I?
Why do I serve this, you know, meat bag human?
But even to your point, like, that's actually very, that's very, that's basically say, what if the human brain were in a machine?
How would it think about things?
And your point is basically that, like, it's not that that wouldn't necessarily be an issue or a concern, it's just that if we're actually going to game this out and think about it at a longer time scale, whatever disastrous result could come from this wouldn't probably be one that we could conceive of in the mid fifties or nineteen sixties in terms of, like, actually that concern there.
Exactly.
Because these things, again, under the hood, these sort of really weird alien forms of intelligence.
So we're not, I think it's, there's this often this implicit assumption, both in the public and I think among some researchers that we're moving up this staircase in intelligence, and at some point in time we'll get to, like, human intelligence.
We'll build an AI that is human intelligence.
I think what we're actually finding is it's not the case that human intelligence is one specific spot in a vast space of possible intelligences with many different dimensions to it and we're sort of exploring the space and we're building lots of different AI systems that think very differently than people.
So if you look at something like chat GPT, for example, it can engage in chatting with people and it sort of comes across like a person.
Now you can modify these systems, you can give them what's called reinforcement learning with human feedback.
You can basically give it feedback from people about this is a good answer, a bad answer to tweak some of these chatbots to make them more honest or more helpful to people.
And they can, you know, start giving answers back that are maybe what you'd like to get from an AI.
So you say, okay, you're an AI and it says, okay, what can I do to help you want to assist you?
But then people have found ways to jailbreak these systems to get around this and they can get the AI to say like, oh, I'm an evil AI, I'm going to commit world domination and I'm going to hack all the computers and murder everybody.

And that in and of itself also, like this is great, okay, but it's also like misleading because that's not actually what the AIs do.
All it's doing is generating new text.
So the reason why it's doing that isn't because the AI decided it wants to take over the world.
It's that you've created a block of text that's telling a story and AI just continuing the story.
And if it's a good story, this text generator will repeat back things that sound good.
If it's a scary story about how some AI takes over the world, it'll keep playing that role that you've effectively nudged it into playing.
And so it's really important to try to like wrap your head around like, what's going on with these systems so that we can use them safely?
Because right now the chatbot is just a chatbot.
But eventually systems will be doing more powerful things and we do want to make sure that they're doing what we would like them to do.
So in this last section, I want to, before we just get to going through the four key elements, people have a good takeaway.
I want to kind of get to the story so far from the way that you really introduced the book.
So you're talking about the robotics revolution of the 2000s.
You've wrote a book previously about autonomous weapons and you tell a story about being in Iraq during the surge, talking about how there was just this exact like natural security defense application of robotic technology.
I'd love for you to just kind of like introduce that robotics revolution, the surge story, and then illustrate how robotics was a perfect illustration of that.
And then just the follow-up is I'm wondering if there's like an AI equivalent defense related version of the story that we could see happening like right now.
There's a lot of questions, there's just three of you, so feel free to pause and re-ask.
I hate shotgunning questions, but yeah, take it wherever you want to go.
Yeah.
There was this single moment, at least in my life, where I realized that robots were going to be transformative to warfare, and that for me at least became this kind of stepping stone to look more broadly at AI and how it's transforming global power.
So I was in Iraq during the surge, so in 2007 to 2008 timeframe, I was a soldier in the army and we were driving down the road, we came across a roadside bomb, it improvised explosive device or ID, as one did at the time, they were everywhere.
Now we saw it first, which is the preferred method of finding them rather than just running into it.
So we pulled over and we called up the bomb disposal crew.
So they come out, they were very busy, so it took a while, they came out and then this huge armored vehicle, and I'm expecting to see a person come out in a giant bomb suit.
Heart Locker, like very specific imagery of that.
Exactly, right?
That people may have seen in the Heart Locker, these huge suits that the bomb disposal team

uses.
So I'm expecting to see one of those.
I'd just been in country like a month, so I was like, oh, this has been really interesting.
And instead, this little robot comes out.
And the light bulb went on for me, I was like, oh, yeah, send the robot to defuse the bomb.
That makes way more sense than having some poor person stick in their head over the bomb and snip on the wires.
And the US military invested heavily in thousands of air and ground robots in Iraq and Afghanistan, for very special purpose missions, like aerial reconnaissance, drones, or bomb disposal, really small ground robots.
But the more I thought about it, the more I started thinking about, there's all sorts of things in warfare that are dangerous, what would be great to put that robot out in front.
And so when I left the army, I went to the Pentagon and I worked there as a civilian.
This was one of the areas I worked on, how can the US military invest in robotics to make the US military more effective to save US service members' lives?
And then as the deep learning revolution began to kick off, it was clear to me that the story was much bigger than just robots, that the things we were seeing happen in artificial intelligence were really deeply profound.
We're going to have huge effects on not just global power, but also our lives in the coming decades.
So you would say, though, this given the nature of, let's say, like the most active conflict people could think of, like the war in Ukraine, there is not quite an example of, there hasn't been an AI equivalent of your experience with the IED and the robot, basically, correct?
There is, well, it depends on who's equivalent, right?
So the US Defense Department's first big project to use technology coming out of the deep learning revolution was Project Maven.
That was a project that was used to use AI image classifiers, which by now is a pretty mature technology.
After taking a whole bunch of images, labeling them, so people go through and label all the images, this is a truck, this is a building, this is a person, and then train an AI model based on these labeled images, and AI model can then learn to identify things because not all trucks look the same, but there is a general characteristic of truckness that people can learn.
And it shows that you can train an AI on that if you show it enough trucks over time.
So the DOD invested in this AI image classifiers through Project Maven and used it to process drone video feeds.
And this is really important for US military because they're basically collecting more imagery off of drones than they can possibly process with people.
So AI is really helpful for that.
Well, that project's five years old now.
And so the use of AI is already something that's underway inside the US military.
And we know that the United States is helping Ukraine with intelligence.
How much AI pulls a role in that is a little bit unclear.

It's probably the case that it's somewhat involved.
How significant it is, I think we don't really know because that's going on behind closed doors.
Man, I think a lot of folks in the audience probably have heard of Project Maven in the sense that it was also very controversial over employee protests about Google, at Google about working with Pentagon and US military.
So given the fact that we're talking about AI and how there are obviously all these deeply civilian-centric applications, we're having this conversation in the context of deep geopolitical precarity.
Do you have any insight or perspective on how that discussion has evolved as the US has transitioned much more into a, let's say, great power competition framework?
Because from my perspective, just being of similar age to a lot of the protesting employees, a lot of the 2017 incident was very tied to both issues at the US-Mexico border, but also a war on terror framework where you're going to see it in the context of the drone war and you're going to see it in the context of what's happening in Yemen or northern Pakistan.
Is there a difference in reaction in Silicon Valley when it's actually, this is happening here, it's happening in China?
Get on board, yes or no.
I'm just curious if you have any thoughts on that.
Well, I don't know whether the shift towards a concern about China has won over any of the people that were protesting the US military's use of AI with Project Maven.
And we saw protests from employees at Google and Microsoft and Amazon, employees wrote open letters protesting those companies working with the military.
What has been the case is that other than this one discrete incident of Google discontinuing their work in Project Maven in this very public fashion, while US leaders in the national security world really panicked that they were going to be shut out of this game-changing technology because, of course, over in China, Chinese tech employees are not going to write a letter protesting the government.
That's going to land them in jail.
And so while US national security leaders really were concerned there would be locked out of this game-changing technology, you see that hasn't been the case.
There's been a whole slew of defense-oriented AI startups that are competing in the AI space.
They're working with the defense community.
And in fact, all of the major tech companies, including Google, are working with the US military.
And so the reality is that I don't know whether the individuals have been won over by some of their shift, but at the corporate level, there's ample companies that want to work with the US government.
The biggest challenges that I see are on things like contracting and acquisitions.
It's not the ethical issues.
It's the slowness of government contracting that makes it difficult for companies to work with the US government.
So last question to take us out.
The book is super, super meaty, so I don't feel at all guilty of offering people a bit

of a free preview.
And you even say this in the book jacket.
So it's really right there.
You can get this on Amazon.
There are four key elements to how people should conceive of this.
I'd love for you just to walk through all four of those real quicks.
So we've got data, computing power, talent, and then institutions.
And that will take us out of the episode.
Yeah.
So if you think about AIs like another industrial revolution, we saw during the industrial revolution that nations rose and fell on the global stage based on rapidly the industrialized and even the key metrics of power change.
So coal and steel production became a key input for national power.
Coal became a geo-strategic resource that countries willing to fight war is over.
And I argue that data, computing hardware or compute, human talent, and institutions, the organizations that are needed to take all of these raw inputs and translate them into useful AI applications are going to be the key battlegrounds of AI competition.
And in areas like data and institutions, it's a fairly level playing field between the US and competitors like China.
But the United States has huge advantages in hardware, as we've talked about because of the role of the US and other allied countries in the semiconductor supply chain, but also in talent because the best AI scientists in the world, including from China, want to come to the United States and the US is a magnet for global talent.
And if the US can harness these advantages in hardware and talent, the US can remain a leader in artificial intelligence.
Well, I think that's an excellent place to leave it.
Thank you so much for joining me on the realignment.
Thank you.
Thanks for having me.
Hope you enjoyed this episode.
If you learned something like this sort of mission or want to access our subscriber exclusive Q&A, Lotus episodes and more, go to realignment.supercast.com and subscribe to our $5 a month, $50 a year, or $500 for a lifetime membership rates.
See you all next time.