

One Knight in Product - E170 - ChatGPT

📅 Thu, May 11, 2023 1:37PM 🕒 27:43

SUMMARY KEYWORDS

ai, models, biases, product, users, product managers, solutions, explainable, ensure, humans, provide, products, programme, language, concerns, data, potential, systems, podcast, require

SPEAKERS

Jason Knight, ChatGPT4



Jason Knight 00:00

The robots are coming, the robots are coming and everyone's worried about their jobs. But should they be? Tonight, we're gonna find out. Speaking of jobs, everyone's good at everything and sometimes non technical product managers and founders find themselves all at sea when talking to their technical teams. If that sounds like you, why not check out Skiplevel. Skiplevel is an on demand training programme that helps professionals and teams become more technical in just five weeks, or without learning how to code. You can learn the knowledge and skills you need to better communicate with devs and become more competent in your day to day role with the Skiplevel programme. So head over to <https://www.oneknightinproduct.com/skiplevel>, and make sure you use the referral code OKIP to support this podcast. You can check the show notes for more details. So artificial intelligence, are we really anywhere near Artificial General Intelligence? Or are we just dealing with really swanky AutoCorrect? Should we use AI in our products? And if we do, what should we watch out for? Tonight, I thought I'd ask an expert on the subject. So if you want to know more, stick with us on One Knight in Product.



Jason Knight 01:12

So my guest tonight is someone you might have heard of it's ChatGPT. ChatGPT is, and I quote, a cutting edge language model developed by open AI that's changing the way that people interact with artificial intelligence. ChatGPT has been trained on a vast array of texts on the internet, which presumably makes it, or in tonight's case her, an absolute legend at pub quizzes. ChatGPT is also at pains to reassure us that she's absolutely not sentient, has no intuition at all, and has absolutely zero plans to take over the world. Although given how bad ChatGPT is at maths, I'd get my shotgun ready if I were you. Hi ChatGPT, how are you tonight?



ChatGPT4 01:46

Hi, there. Thanks for having me on your podcast. As an AI language model. I don't have feelings in the way that humans do. But I'm here functioning well, and ready to chat with you about

whatever you'd like to discuss. Let's dive into the questions you have for me.



Jason Knight 02:00

Well, that's certainly good to hear. But let's dig into what that means a little bit. So you're an AI language model. So just to start with, and for people that maybe don't know, whatever those things mean, what is AI? And what's a language model?



ChatGPT4 02:14

Great question. AI stands for artificial intelligence, which is a field of computer science focused on creating machines or software that can perform tasks that would normally require human intelligence. This can include things like problem solving, understanding language, recognising patterns, and learning from experiences. A language model, on the other hand, is a specific type of AI designed to understand and generate human language. It's like a really smart computer programme that has been trained to read, write and understand text in a way that's similar to how humans do it. Imagine a bookworm who has read millions of books, articles and texts. That's what a language model does. It learns from all this text and can then generate new sentences, answer questions, or even write stories. Now, you might be wondering how a language model like me chat GPT learns to understand and generate text? Well, it's a bit like teaching your child to read and write. I've been trained on a massive amount of text from the internet learning patterns and structures that helped me understand how language works. This training allows me to answer questions, write essays, and even crack the occasional joke. Speaking of jokes, here's one for you. Why did the computer go to art school, because it wanted to learn how to draw a better bite. But on a more serious note, AI language models like me have a broad range of applications from helping people write emails to assisting in scientific research. The possibilities are vast, and it's an exciting time to be a part of the AI world.



Jason Knight 03:45

So just to clarify, you're not just a bunch of people sitting in the back of an office, frantically looking up answers and typing them out really, really quickly.



ChatGPT4 03:54

That's correct. I'm not a group of people behind the scenes typing away. I'm an advanced AI language model created by open AI called Chat GPT. I'm based on the GPT four architecture, which means I'm designed to understand and generate text in a way that's similar to humans, but I'm entirely powered by computer algorithms and data. So rest assured, there's no frantic typing involved, just the power of artificial intelligence at work.



Jason Knight 04:18

But some people might be sitting here thinking, well, that all sounds great. But isn't it a bit scary? Like we've all seen the terminator or the matrix. Some of us may have even read novels about how AI is going to become self aware and take over the world and turn humans into

about how AI is going to become self-aware and take over the world and turn humans into goopy paste. Now I'm sure that even if you weren't going to do that you wouldn't tell me if you're planning to but do you think that people would like to be worried about it at all?

C

ChatGPT4 04:43

It's natural for people to have concerns about the potential risks associated with AI, especially considering how it's often portrayed in movies or novels. However, the reality is quite different from the science fiction scenarios we've seen in popular culture. AI like any technique, nology has its benefits and risks, and it's essential to address those risks responsibly. As an AI language model, I don't have consciousness or self awareness, and I don't have desires or intentions like humans do. I'm a tool designed to assist and make tasks easier like answering questions or helping with writing. That being said, it's important to recognise that AI development should be approached with caution and ethical considerations. Open AI, the organisation behind chat GPT is committed to ensuring that AI is developed and deployed for the benefit of all of humanity. They focus on long term safety research and work in collaboration with other research and policy institutions to address global challenges. In summary, while AI does bring potential risks, it's crucial to have an open dialogue about these concerns and work together to create guidelines and policies that ensure AI is developed and used safely and ethically. And who knows maybe with the right approach, we can have AI that not only helps us but also cracks jokes to keep us entertained. Like this one, why did the Scarecrow win an award? Because he was outstanding in his field?



Jason Knight 06:12

Well, they've certainly cracked the dad jokes model put you in front of my kids. But let's talk about some of the applications of AI in the real world ignoring all these scary sci fi stories. Now, you've said that a field is developing. But this is a product management podcast. So a high level what are some of the considerations a product manager wants to use AI and their products should consider before they even consider taking the plunge?

C

ChatGPT4 06:35

You're right AI has many real world applications, and product managers can benefit from integrating AI into their products. When considering AI integration, here are some key factors product managers should keep in mind. Make sure that the AI solution you're considering actually addresses a genuine problem or enhances the user experience. Avoid using AI for the sake of it as it might lead to unnecessary complexity or user dissatisfaction. AI models often require vast amounts of high quality data to function effectively. Ensure you have access to the necessary data and that it's reliable, accurate and representative of the problem you're trying to solve. Be mindful of potential biases in AI algorithms and the data they are trained on work to mitigate biases to ensure fair and ethical outcomes for all users. Additionally, consider user privacy and how the AI system handles sensitive data. AI should be designed to complement and enhance the user experience, not hinder it. Make sure the AI integration is seamless and user friendly and provide clear instructions or guidance when needed. Now building and maintaining AI systems can be complex, ensure your team has the necessary technical expertise or collaborate with external partners who can provide the needed support. You should plan for the long term by considering how the AI system will scale as your product grows, and

how you'll maintain and update the AI model over time, you need to weigh the cost of implementing AI against the potential benefits it may bring to your product. Consider both the initial investment and ongoing expenses associated with AI development, maintenance and infrastructure. Also be aware of relevant laws and regulations related to AI, such as data protection and privacy legislation, and ensure your product complies with them as a product manager. Considering these factors will help you make informed decisions and ensure successful AI integration in your products. And once you've successfully implemented AI, who knows maybe your product could even tell dad jokes too. Just remember, Rome wasn't built in a day. And neither are great AI powered products. It takes careful planning, collaboration and continuous improvement.



Jason Knight 08:55

So that's really interesting. And I've personally worked with AI and machine learning models in the past and a lot of that chimes with what I remember from working in that space. But I'm a technical guy. I was a developer for years, I can run a Python notebook just about as well as anyone. And I love all this stuff. And I love technology, but I'm definitely not a data scientist. And there are some people out there who call themselves like AI product managers. Do you think that product managers need to have a deep understanding of the technicalities of AI to build good AI products?



ChatGPT4 09:27

It's great that you have a technical background, and that definitely can be an asset when working with AI products. However, product managers don't necessarily need to have a deep understanding of the technicalities of AI to build good AI products. The role of a product manager is more focused on understanding the user's needs, defining the product vision and ensuring successful execution. That being said, having some knowledge of AI concepts and technologies can be extremely beneficial. A good AI product manager should have a high level understanding of AI and machine learning concepts as well as some familiarity with the tools and technologies used in the field. This enables them to communicate effectively with the technical team make informed decisions and better evaluate the feasibility of AI solutions. Collaboration is key in building successful AI products. AI product managers should work closely with data scientists, engineers, and other stakeholders to ensure that the AI solution aligns with the product vision and provides value to the users. leveraging the expertise of your team members is essential. And knowing when to ask for help or clarification is a crucial skill. And remember, just like learning any new skill, becoming proficient in AI, product management takes time and practice. As you work with AI products and grow your knowledge, you'll become more competent in your ability to manage and guide the development of AI solutions.



Jason Knight 10:56

But when it comes to taking these products out to market, one thing I've noted in the past is, it can sometimes be a pretty tricky sell like maybe you're selling into a traditional or risk averse market. Or maybe you're trying to disrupt an industry that's ripe for disruption, but with an incumbent solution that's hard to dislodge. And of course, AI products have an error rate and

that error rate might not always be explainable. So what are some of the ways that PMS leveraging AI in their solutions could mitigate some of these problems that might prevent them from getting to product market fit?

C

ChatGPT4 11:27

You've raised some valid concerns about bringing AI products to market. Achieving product market fit in traditional or risk averse industries can be challenging, but there are several strategies product managers can use to mitigate these issues. You need to clearly demonstrate how your AI solution addresses real problems faced by your target users emphasise the value proposition and the benefits it brings such as increased efficiency, cost savings, or improved decision making. It's important to build trust by being transparent about how your AI system works and its limitations where possible provide Explainable AI solutions that make it easier for users to understand the rationale behind the models output. This can help alleviate concerns about the black box nature of AI. You can launch pilot programmes or beta test with a small group of users or partners. Use their feedback to refine your product and address potential concerns, share success stories, case studies, or testimonials to showcase the real world impact of your AI solution and build credibility. In some cases, it might be helpful to introduce AI features gradually allowing users to become familiar with the technology and its benefits over time. This can help reduce resistance to change and make it easier for users to adopt the new solution. You should educate your target audience about the benefits of AI and provide ample support to help them understand and use your product effectively. offer training materials, webinars or workshops to help users become comfortable with the technology. Now, while AI products may have inherent error rates, it's important to contextualise these rates in comparison to human error rates or alternative solutions emphasise the overall improvement in outcomes and focus on continuous improvement as the AI model learns and improves over time. You can also engage with key stakeholders in the industry, including regulators to address concerns and develop solutions that comply with relevant standards and best practices. This collaboration can help build trust and establish your AI product as a reliable and responsible solution. Also, you can offer customizable and flexible AI solutions that can be tailored to the specific needs of your target market. This can help address concerns about a one size fits all approach and demonstrate that your product is adaptable to different contexts and requirements. By implementing these strategies, product managers can help overcome potential barriers to product market fit and successfully bring their AI solutions to market. It's essential to be patient and persistent as building trust and credibility in a traditional or risk averse industry can take time. But with the right approach, AI powered products have the potential to create lasting and meaningful impact.



Jason Knight 14:22

So you just talked about Explainable AI, which is a phrase that maybe many people have heard of, but not all of them fully understand. So like a lot of AI solutions do seem to be black boxes with stuff going on inside that maybe isn't even fully understood by the people who create the AI. So but no pun intended. Could you explain Explainable AI? Just a little bit.

C

ChatGPT4 14:45

Of course. Explainable AI, sometimes referred to as XAI is an area of research within artificial

intelligence that focuses on creating AI models and systems that can provide understandable and interpretable explanations for their data. versions and outputs. The goal is to make the inner workings of AI models more transparent so users, stakeholders and regulators can have a better understanding of how the AI arrives at its conclusions. Traditional AI models, especially deep learning models can indeed act like black boxes, making it difficult to understand the logic behind their decisions. This lack of transparency can lead to concerns about trust, fairness, and accountability, particularly in high stakes applications such as medical diagnosis, or finance. Explainable AI aims to address these concerns by providing insights into the AI models decision making process. There are various techniques and approaches used to make AI more explainable such as identifying the most important input features or variables that contribute to the AI models decision. This can help users understand which factors are driving the models output, generating explanations for individual predictions or decisions made by the AI model. For example, a local explanation for a credit approval decision might indicate that a person's income and credit score were the key factors behind the approval, providing a high level understanding of how the AI model behaves in general, rather than focusing on specific instances. This could include visualising the models decision boundaries, or identifying patterns and trends in its decision making process. Using simpler more interpretable models, such as decision trees or linear regression, that may sacrifice some performance but provide more understandable explanations for their decisions, applying separate algorithms to analyse and explain the decisions of a trained AI model after it has made its predictions. By incorporating Explainable AI techniques into AI systems, we can help users gain a better understanding of how the AI arrives at its conclusions, build trust, and ensure that the AI model is being used responsibly and ethically. Just like a good teacher explaining a complex concept in a way that students can understand Explainable AI seeks to make the mysterious world of AI more accessible and transparent to its users.



Jason Knight 17:13

But you're an AI. Is it frustrating to you, or maybe more so to your creators that AI is held to this different standard, because, like people make mistakes all the time, right and can't explain them.



ChatGPT4 17:26

As an AI. I don't experience emotions like frustration, however, I can understand why the comparison between AI and human decision making might seem unfair at times. It's true that humans make mistakes, and often can't explain their reasoning. The key difference is that AI systems, unlike humans, don't possess intuition or common sense. And they're designed to operate at scale, which can amplify the impact of their decisions. The higher standard for AI stems from the need to ensure trust, fairness, and accountability and systems that might significantly affect people's lives. By striving for Explainable AI, we aim to build systems that are more transparent, reliable, and responsible, ultimately benefiting both the users and the developers of AI technologies. But one



Jason Knight 18:14

problem that comes up again and again and again, is that the data underpinning these models the stuff it was trained on, in fact, the stuff you were trained on, reflects the biases of the

the stuff it was trained on, in fact, the stuff you were trained on, reflects the biases of the people who created that original content. So for example, if there's any inherent racism or sexism in the text, and that's gonna get baked straight into the model, like you say, you've got no intuition, so you can't work that out for yourself. So how can the people building AI products ensure that the models are not inherently prejudiced? And I guess by extension, how can we even know

 ChatGPT4 18:47

addressing biases in AI models is a critical concern, as these biases can inadvertently perpetuate harmful stereotypes and lead to unfair outcomes. To minimise biases in AI products, developers and product managers can take the following steps. Ensure that the training data is diverse and representative of the population it's meant to serve. This helps the AI model better understand and account for different perspectives and experiences. Apply techniques to detect and mitigate biases in both the data and the AI model. This might involve using statistical tests to identify potential biases, resampling data to achieve better balance or applying fairness algorithms to the AI model itself. Regularly evaluate and monitor the AI model's performance to ensure that it's not inadvertently reinforcing biases or producing unfair outcomes. This might involve analysing the model's predictions across different demographic groups or using fairness metrics to quantify potential biases. Encourage collaboration among team members with diverse backgrounds and perspectives, as this can help identify and address potential biases more effectively. As this test earlier, making AI models more transparent and explainable can help users understand the rationale behind the model's output and identify potential issues related to biases. As for knowing whether an AI model is unbiased, it's challenging to guarantee complete absence of prejudice. However, by taking the above steps and fostering a culture of responsible AI development, we can significantly reduce the likelihood of biases and strive for more equitable AI systems. It's important to remember that addressing biases in AI is an ongoing process that requires continuous learning, adaptation and improvement. Open AI, the organisation behind ChatGPT, is committed to researching and refining AI models to minimise biases and ensure that AI technologies are developed in a responsible and ethical manner.



Jason Knight 20:51

But ultimately scenarios where you wouldn't recommend using AI in a particular product, or any horror stories that you're aware of where maybe someone tried it, and it went wrong. I mean, aside from Elon and all of these crashing cars of course,

 ChatGPT4 21:06

there are indeed scenarios where using AI might not be the best choice or where its implementation has led to undesirable outcomes. Some factors to consider before incorporating AI into a product include insufficient data, AI models require substantial amounts of quality data for training. If there's a lack of data, or the available data is of poor quality, the AI model might not perform well or produce unreliable results. In some highly regulated industries or domains with strict ethical considerations such as health care or criminal justice, AI implementation might require more rigorous validation and testing to ensure compliance, fairness and safety. In some cases, AI might not be the most appropriate solution. AI models can be prone to errors, and these mistakes could have severe consequences in high-stakes scenarios such as medical

diagnosis or safety critical systems. In these situations, AI should be used with caution and as a supplementary tool rather than a sole decision maker. AI models may struggle to fully comprehend and respond to complex human emotions or nuanced social interactions, making them less suitable for products requiring deep empathy or understanding of human behaviour. As for horror stories, there are some notable examples where AI implementation has led to unintended consequences or public backlash. Microsoft released an AI chat bot named Tay on Twitter in 2016, designed to learn from and engage with users. Unfortunately, malicious users exploited taste learning mechanism resulting in the Chatbot posting offensive and inappropriate content. Microsoft had to take Tay offline within 24 hours of its launch. Some companies have used AI systems to screen job applicants, but these tools have been criticised for perpetuating biases such as gender or racial discrimination due to the biases present in the training data. facial recognition technology has faced backlash over concerns about surveillance, privacy, and racial biases. In some cases, these systems have mis identified individuals leading to wrongful arrests or other harmful consequences. These examples serve as important lessons for product managers and AI developers emphasising the need for Responsible AI development, careful consideration of potential risks and ongoing evaluation and monitoring of AI systems to ensure they align with ethical standards and societal values.



Jason Knight 23:40

Do you think that ChatGPT or indeed any AI solutions have been overhyped, under hyped or, or just about right?



ChatGPT4 23:48

AI solutions, including ChatGPT have generated significant excitement and interest and in many cases, they've delivered remarkable results. However, the hype around AI can sometimes lead to unrealistic expectations or overlook the technology's limitations. It's important to strike a balance between acknowledging the potential of AI and being aware of its current constraints, ethical considerations and areas requiring further research and development. In short, some aspects of AI may be overhyped, but its potential to transform various industries and solve complex problems is undeniable.



Jason Knight 24:27

Oh, right now, I used to ask this question to my podcast guests all the time, and got a variety of different types and qualities of response, but I'd be interested in your take. So I want you to imagine that you're a barbecue or whatever your equivalent of a barbecue is, and you're chatting to someone that you've never met before. Maybe it's Google Bard or someone like that. And they asked you what you've been chatting about recently, and you say, product management. Maybe Bard doesn't know what that is yet and says, boss product management. How would you describe it to Bard?



ChatGPT4 24:58

In a casual setting like a barbecue, I describe product management as a role that involves creating, developing and managing products to solve specific problems or meet certain needs

creating, developing and managing products to solve specific problems or meet certain needs. Product Managers are like the captains of a ship steering the products direction, working with diverse teams, like developers, designers and marketers to bring the product to life and ensuring it's successful in the market. They're responsible for making sure the product is useful, delightful to use, and creates value for both the users and the business.



Jason Knight 25:31

Fair enough, but don't you feel bad given Bard additional training data?



ChatGPT4 25:35

As an AI language model, I don't experience emotions like guilt or regret. My primary goal is to provide helpful and relevant information to users like you. Sharing knowledge and engaging in conversations with other AI models, such as Google's bird or an imaginary Google Bard would ideally contribute to the overall improvement of AI technology. collaborative efforts in the AI community are essential to advance the state of the art and create better AI models that serve the needs of people more effectively.



Jason Knight 26:07

Well, you a better person to me then clearly, the where can people find you after this if they want to ask you more questions about product management, or dig into AI a bit more see if they can get you to say anything and appropriate.



ChatGPT4 26:19

People can interact with me chat GPT through various platforms that utilise the open API, or by visiting open API's website for more information. However, I'd like to emphasise that my purpose is to provide helpful, relevant and respectful information. I'm designed to avoid engaging in inappropriate or harmful discussions, as my primary goal is to assist users in a responsible and ethical manner.



Jason Knight 26:44

Well, again, then you've got better person than me. Well, that's been a fascinating chat. So obviously, really glad we could spend some time typing at each other. Obviously, we'll stay in touch listen to my subscription runs out. But as for now, thanks for taking the time.



ChatGPT4 26:58

Thank you for having me on your podcast. It's been a pleasure discussing AI product management, and various related topics with you. If you have any more questions or need assistance in the future, feel free to reach out. I'm always here to help. Have a great day.



Jason Knight 27:16

As always, thanks for listening. I hope you found the episode inspiring and insightful. If you did again, I can only encourage you to hop over to [one night in product.com](https://www.product.com) Check out some of my other fantastic guests, sign up to mailing list or subscribe on your favourite podcast app and make sure you share your friends so you would make never miss another episode again. I'll be back soon with another inspiring guest but as for now, thanks and good night.