

[Transcript] AI Hustle: News on Open AI, ChatGPT, Midjourney, NVIDIA, Anthropic, Open Source LLMs / OpenAI's DALL·E-3: Major AI Image Generation Upgrade Revealed

Welcome to the OpenAI podcast, the podcast that opens up the world of AI in a quick and concise manner.

Tune in daily to hear the latest news and breakthroughs in the rapidly evolving world of artificial intelligence.

If you've been following the podcast for a while, you'll know that over the last six months I've been working on a stealth AI startup.

Of the hundreds of projects I've covered, this is the one that I believe has the greatest potential, so today I'm excited to announce AIBOX.

AIBOX is a no-code AI app building platform paired with the App Store for AI that lets you monetize your AI tools.

The platform lets you build apps by linking together AI models like chatGPT, mid-journey, and 11Labs.

Eventually, we'll integrate with software like Gmail, Trello, and Salesforce so you can use AI to automate every function in your organization.

To get notified when we launch and be one of the first to build on the platform, you can join the wait list at AIBOX.AI, the link is in the show notes.

We are currently raising a seed round of funding.

If you're an investor that is focused on disruptive tech, I'd love to tell you more about the platform.

You can reach out to me at jaden at AIBOX.AI, I'll leave that email in the show notes.

So if you head on over to the Dolly 3 landing page, which you're going to hear is, you know, OpenAI says, quote, Dolly 3 understands significantly more nuanced and detail than our previous systems, allowing you to easily translate your ideas into exceptionally accurate images.

So the big thing I think to know here is that Dolly 3 is actually going to get embedded into chatGPT.

So this is not something that is like in the past, if you use Dolly 2, you went over to your website, play around with it.

Now you're going to be able to literally, if you have the premium account of chatGPT, you're going to be able to access and use Dolly 3 right on chatGPT.

Now this is not launched right now.

They said that this is going to be coming to enterprise customers and chatGPT plus customers in October.

And it's also apparently via the API and in labs later this fall.

So it's going to be starting to get integrated into a ton of different products, which this is actually really cool.

For those that don't know, I'm currently building an AI platform that allows you to link together a ton of different models and do some really impressive stuff.

One of the big struggles we've had is that the major image generator is mid journey.

By far the best one.

They do not have an API out yet, which is how you integrate it with software, which is super annoying.

So you got to make all these workarounds.

[Transcript] AI Hustle: News on Open AI, ChatGPT, Midjourney, NVIDIA, Anthropic, Open Source LLMs / OpenAI's DALL·E-3: Major AI Image Generation Upgrade Revealed

It's a ton of work.

And this is really big.

The fact that this is going to be, you know, later this fall, the API for this is going to be available and this is going to be amazing.

So I'm super excited about this.

This is going to spur a ton of innovation.

I know a lot of other developers are and people working on startups are very, very excited about this.

So here is what they're saying about it right now.

They say, modern text to image systems have a tendency to ignore words or descriptions forcing users to learn prompt engineering.

Dolly three represents a leap forward in our ability to generate images that are exactly adhere to the text you provide.

This is really interesting, right?

If you've tried to generate anything on, you know, mid journey, you have to like throw random keywords in there to get it to do stuff.

It's getting better to the point where it's going to be like eventually talking to chat GPT say you just tell it what to do in plain English and it's going to do it.

You don't have to change together, you know, keywords and other things.

So I think this is very, very interesting.

They show an example of a picture where you literally where they literally said the sidewalk bustling with pedestrians enjoying the nightlife a bustling city street under the shine of a full moon at the corner stall a young woman in a fiery in fiery red hair dressed in a signature velvet cloak is haggling with a grumpy old vendor.

The grumpy vendor a tall sophisticated man is wearing a sharp suit sports a noteworthy mustache in his animated convert conversing on his steampunk telephone.

Okay.

All of that right there.

They did an image that is literally exactly that.

It's just like it is exactly what the words described.

So if you've ever tried to use mid journey or anything else, you know, like the challenge where you're like, you know, like seven, you know, millimeter lens, camera, far back.

You know, just like you throw like all these different like, like UX UI, like keywords on there that aren't really part of a story.

So now you literally tell the story, you type the exact words and it can generate this thing be really interesting for things like books.

It's going to be able to like take the page of a book and say create an image of, you know, what's happening on this page.

Like so crazy, this is going to be amazing.

So one thing that they have noted is that even with the same prompt, Dolly three gives a lot better images over Dolly two.

So they have the same prompt, which is an expressive oil painting of a basketball player dunking depict it as an explosion of a nebula.

[Transcript] AI Hustle: News on Open AI, ChatGPT, Midjourney, NVIDIA, Anthropic, Open Source LLMs / OpenAI's DALL·E-3: Major AI Image Generation Upgrade Revealed

And the, you know, Dolly two gives us pictures kind of, I don't know, splotchy, not very, not very clear.

They have it on Dolly three.

And it's like, it's a really impressive picture.

It's exactly what you'd kind of imagine or think of when you hear that prompt.

So Dolly three is built straight into chat should be to like I mentioned, and this kind of lets you use chat should be as a brainstorming partner to help you refine your prompts. Right.

Because even though it does exactly what you say in your text, sometimes you may want to add or remove things.

This is kind of nice where you literally are in chat should be an environment you're familiar with while you're developing these images.

And you can also ask chat should be what you want to see.

And just, you know, throw it to a paragraph, throw it into simple sentences is going to generate it.

So when prompted with an idea, chat should be T is going to automatically generate Taylor detailed prompts for Dolly three, which is really cool, right?

So chat should be is actually going to be able to create prompts.

I've actually been using this for a while.

And I wasn't sure if this was something that they were tying in, but it looks like it's getting a lot more tied in where chat should be taken generate prompts that really bring your idea to life.

If you like a particular image, but it's not like exactly what you're going for, you can also ask chat should be to make small tweaks with just a few words.

And it's, you know, you're constantly like reiterating and kind of creating that the perfect image for yourself, which is really, really cool.

Like I said, this is going to be available in early October for paying customers.

And I think what's really interesting, they said, quote, as with Dolly to the images you create with Dolly three are yours to use, and you do not need our permission to reprint, sell or merchandise them.

So this is really interesting.

They're essentially saying you have the copyright to these images that you're generating, which is very and literally like commercial rights are selling, you can they're telling you you can resell them.

I don't know if that means that they have, you know, done some unique things with their image data set where they're essentially using images, they have the rights to I'm not sure like where the line is on that, but it says that they've taken steps to limit Dolly three's ability to generate, you know, of course, violent adult or hateful content.

That's makes sense.

They've also essentially said that it has, you know, mitigations to decline requests to ask for a public figure by name.

So, you know, essentially, I think if you're saying like a picture of Donald Trump, you know, flying an airplane or a picture of, you know, Biden throwing a grenade or something

[Transcript] AI Hustle: News on Open AI, ChatGPT, Midjourney, NVIDIA, Anthropic, Open Source LLMs / OpenAI's DALL·E-3: Major AI Image Generation Upgrade Revealed

crazy, right?

It's not going to do that, which is interesting.

I'm sure that's controversial.

People are going to want jailbroken versions that they can generate public figures doing things, even if it's satire, I'm sure, or parody and stuff.

So it's going to be interesting.

I think mid-journey was in a bit of controversy earlier this year when they essentially, if you tried to ask for a picture of Xi Jinping, the president of China, it wouldn't allow you to do it, but you could do other people, and I think that's from like some Chinese pressure or other things.

So, you know, there's interesting things, but it looks like they're doing the same thing here, and at least Open AI seems to have a blanket policy where it's no public figures versus just, you know, a couple random Xi Jinpings that are complaining.

So in any case, they said that they improved safety performance and risk areas like generation of public figures and harmful biases related to people in partnership with, they have a bunch of big red teaming thing and a bunch of domain experts that have kind of helped do this.

They said we're, they said they've done this to help inform our risk assessments and mitigation efforts in areas like propaganda and misinformation, interesting.

I'm sure there's a lot of people thought of different opinions on that, but in any case, it says we're also researching the best ways to help people identify when an image was created with AI.

We're experimenting with a provenance classifier, a new internal tool that can help us identify whether or not an image was generated by Dolly 3, and we hope to use this tool to better understand the ways generated images might be used.

So they'll share more on that soon.

Very interesting.

They're coming up with a way to actually detect which images are generated by AI.

This is something that Google has been working on as well, so they're not the first people on this, but they, depending, maybe some people to crack it on a high level.

One thing that they have mentioned is really giving creative control.

They said Dolly 3 is designed to decline requests that ask for an image in the style of a living artist.

So I guess you could still say do X, Y, and Z in the style of Vincent Van Gogh if they're dead, but if it's a living artist, then it won't generate, it's essentially designed to not generate images in that same style.

So living artists today are people, a lot of them have been complaining like, hey, these image generators are just knocking off my work, so we'll be able to do that.

Another interesting thing that they've also said is that creators can now also opt their images out from training of our future image generation models.

Now, what's, you know, the keyword here is future image generation models.

So the model they have today is using everyone else's, but I went over to it and I saw they literally have a form where you put your name, your email, if you're the rights owner or

[Transcript] AI Hustle: News on Open AI, ChatGPT, Midjourney, NVIDIA, Anthropic, Open Source LLMs / OpenAI's DALL·E-3: Major AI Image Generation Upgrade Revealed

if you're requesting it.

On behalf of someone else, you can have a general description of your image and you can actually upload images that you would like to be excluded from future training models.

So you're an artist, you have, you know, let's say a library of like 150 really famous images you've painted, maybe you've sold them, you can upload that library and it will not include those in its generation.

Now this is definitely a step behind something like what Adobe is doing where you, you know, they have a platform that's actually opted in, they have the rights to all the content on the platform and they train models off of that.

Definitely a step behind that as this is still generating from people's stuff.

Very, very interesting.

The images look quite impressive, way more like something you'd see on mid-journey.

I'm looking at a picture of like a stormy ocean out the window.

There's a coffee cup on the, on the window sill that has a wave splashing of coffee.

Really realistic, looks really, really impressive.

So there's some really cool stuff I'm seeing indoor renderings.

So for like architects and stuff outside of buildings, insides of buildings, interior design, ocean stuff, hyper realistic, all sorts of different types of art.

There's one that's like a digital illustration of a beach scene crafted from yarn and it literally looks like yarn, looks like a piece of art a person could have made.

All sorts of really, really impressive art is on this.

So I think this is definitely going to be the next step.

I'll tell you what I think is the biggest deal about this.

Definitely this is impressive.

Definitely they're doing a bunch of new things around safety and like, you know, helping artists.

The biggest thing in my opinion is the fact that this thing is going to have an API.

And if, you know, mid-journey does not come out with an API before then, this is going to be what gets integrated into all software and sees the most adoption.

So very, very exciting.

And we're excited for early October to start testing out Dolly 3 and giving you an in-depth analysis once it is live.

If you are looking for an innovative and creative community of people using chatGPT, you need to join our chatGPT creators community.

I'll drop a link in the description to this podcast.

We'd love to see you there where we share tips and tricks of what is working in chatGPT.

It's a lot easier than a podcast as you can see screenshots, you can share and comment on things that are currently working.

So if this sounds interesting to you, check out the link in the comment.

We'd love to have you in the community.

Thanks for joining me on the OpenAI podcast.

It would mean the world to me if you would rate this podcast wherever you listen to your podcasts and I'll see you tomorrow.