Transkriptor

Everybody is getting excited about AI agents, and the one that we're going to be talking about today is actually the world's first open source general AI agent capable of thinking, reasoning, and just generally behaving like a human would.

Now, this project kind of blew up on social media.

It got more than 280,000 views on X for this announcement video that they made.

It also has 2.8 K stars on GitHub at the time of me recording this video.

So let's take a look at what this thing is.

Welcome to plevo deep dive.

I'm Caleb Freezing, your host and you can see here.

So we are here on Cortex A is X page.

That's the name of the company and the open source generalist AI agent that they've built is Suna.

So let's take a look at this reveal video, this launch video to see what all the hype is about.

Welcome.

My name is Marco and this year is the flat in Lisbon we.

Spent the last three weeks building Suna in.

Suna is a.

You can see here he's talking that they actually spent only three weeks building this agent.

Pretty impressive.

General purpose AI agent capable of completing real world tasks.

Guys, say hi.

Suna is the freeze version of a fully AI employee where AI stop just being a tool and start becoming your real teammate.

Suna can reason, plan and take actions across all sort of domains.

It has access to a visual computer where it can write files, execute code, browse the Internet, use the terminal, and much much more.

So it's a bit like Manus in that regard, but again, it's open source.

Did you mention it's open source?

It's open source and you can go on our GitHub and self host it.

Also, it's available today on WWW dot.

Suna dot SO.

For you to try out, let's jump into some.

Examples.

And we'll take a look at the pricing when they're providing their own hosted inference.

But anyways, we'll we'll see what they're going to talk about here with data processing and enhancement.

Here we see Suna working with existing data.

It is a task to enhance a spreadsheet of B to B leads with the CEOs and Head of Sales name Suna approaches this task.

So just like all of these tools, Manus AI Open AI operator, it's able to use the mouse to navigate the page, drag and drop things and make stuff happen.

It's pretty impressive methodically.

It plans out its approach and searches the web and uses LinkedIn to gather the necessary data.

Finally, it enhances the spreadsheet with the newly gathered data and returns the file to the user.

We see the information.

It's able to drag the borders of the cells here in order to actually read what's inside of them.

Has been updated.

We test Suna to go research about itself.

For that, it uses the web browser just like a human would, searching, browsing and clicking, scrolling and more.

It observes the information and later on will use to gather information to create a final research document to represent its finding.

It even goes to its own GitHub page.

I think for an AI, sort of like looking in the mirror, right?

They're researching about themselves to learn more about how they work.

Here's Sunai's task to create a comprehensive report on the Apple stock.

As usual, it starts out by planning its approach and then uses web search and finance APIs to gather information.

Based on the gathering information, it creates visualizations and serves them to the user.

So there we go.

But they've also provided some other examples sort of more deep diving into how the model handles specific tasks that I think are worth taking a look at.

So here you can see it's writing a detailed report about what's happened in the US stock market in the last two weeks, analyzing the SMP 500 trend, and then telling the user what the market is expecting to see in the upcoming weeks.

And this is a report analysis for the bank CFO, or at least that's what they're telling the agent.

So let's take a look and see what it's able to do here.

So you can see the user inputs the prompt and then pretty much right away it starts

researching, starting research, gathering market data.

And So what I'd be curious to know is it's probably using some kind of open source research tool.

And I wonder how good that is compared to Open AI Deep Research or also Gemini 25 Pro Deep Research.

Perhaps it's using the same thing that Perplexity is using.

But I'm just not sure how that would compare in terms of the structure of the data that's being researched and how it's being presented and actually how the research is happening behind the scenes.

So it would definitely be interesting to learn more about that specific aspect because probably a lot of people are going to be using this tool to do a lot of research.

And then finally, it looks like it has generated a report.

Now it's creating APDF version of that report, It's apologizing for an error, it's continuing to apologize and then now it's getting the final report here with an executive summary, detailed market analysis, etcetera, etcetera.

So it looks like it is currently working to create that along with references.

And then there we go.

We've got the report right here in what seems to be MD format, market data format, and then you can export that as APDF.

OK, let's take a look at the next video here where it's going on Crunchbase, Deal Room and TechCrunch, filtering by Series A funding rounds in the SAS finance space, and then building a report with company data, founders and contact info for outbound sales.

This is something that a lot of businesses would find extremely valuable.

And normally you'd actually have to work with a salesperson to get this done and also probably a research assistant as well.

But here you can use Suna to do everything for you.

So again, it's doing a lot of research here in the background to prepare this report.

And then it looks like it's creating a file to store the information.

And is it still doing research?

I need to Scroll down more to see content about Series A funding rounds.

I can see the table.

OK, you can see what it's actually looking at here.

It's going through a bunch of content to acquire all of this information, bunch of lists.

Now it's gathering the data.

Yep, it's going to another website, gathering more data and then the user has to this is, this is a bit surprising because the user has to kind of interject and say, can you give me the list now?

Which typically for an agentic AI solution, it would just do that automatically.

That's the idea that it doesn't just do one thing at a time.

It's able to do multiple things in sequence.

But you know, maybe this is still a work in progress.

And then finally it generates the the report or it's creating the file right now and there it is.

So that went pretty quickly, but you can see here it is again, once again of an MD file, which gives you the option to export as APDF or download the MD file directly.

Let's take a look at the next use case here.

Can you help me research the B to CAI market and show me some graphs of interesting numbers you found?

Browse the web with the browser, please.

OK, so this should give us an example of what it actually does when it's navigating the web.

You can see it's going on Google just like a human would, and there's this little pop up that seems to have stopped it in its tracks.

I'm not sure.

Doesn't seem to have clicked on the pop up.

That would be interesting actually.

What does it do in that situation?

Except all I can see I'm gathering valuable information.

Let me continue my research by looking for more specific information.

I'll navigate to another Interesting.

OK.

It doesn't seem to mention the fact that it's not able to see because of that pop up.

Keep searching here.

All right, It's already creating the visualization.

It is writing Python code to make that happen and it's executing the code in its own terminal.

Very cool.

It had to install some libraries first.

OK, it says I've created the visualizations and then I will run the script to generate the visualizations.

It says I've successfully done that.

It's stuck in a loop here.

Interesting.

Seems like it actually pulled itself out of that loop.

I've successfully created and then it says I see that the visualization was created in a nested directory.

Let me check the contents of that directory.

OK, great.

I've successfully created all the visualizations for our B to CAI market research.

Now I will present the findings and the visualization to you.

So it looks like it didn't get that far.

It wasn't able to actually present the information.

It sort of the user had to step in, OK, so it's just presenting the files here.

So it's not visualizing it immediately in this visualization space.

Instead, it just gives the files to the user, and then it's the user's job to actually open up the visualization file, which is right here as a PNG.

But still the user gets what they wanted, even if it does require a couple of clicks.

Yeah, you've got all these visualizations being shown.

So it was able to do it successfully.

And there were a number of steps in there that probably if you were working with any other tool, any other AI that was just sort of very static and not dynamic and not agentic like this, you'd have to do each of those individually.

OK, now do this.

Now do this, Now do this and it just wouldn't, you wouldn't end up getting the result that you want because again, this is something that really only a genetic AI can pull off.

And I love here too that they're not saying use Suna or subscribe to Suna or download Suna.

They're saying hire Suna today.

So they really are treating this like it's a real agent.

It's almost like a person that you can hire to get the work that you need to get done, done.

If we just click on pricing over here, we can see that they are charging $29.00 a month for

professionals and small teams.

You get only four hours a month out of this, or $199 per month, and you get 40 hours a month.

So let's just do some napkin math here.

Let's say that you were able to find a somewhat competent assistant, maybe someone who just finished high school, and you were gonna pay them 725 per hour, right?

And let's say you want them to be working for 9 hours in a day and you want them to work five days a week and you want them to work four weeks in a month.

So that's roughly how much it's going to cost you for that person.

Now Suna, on the other hand, costs $199 per month, but again, 40 hours in a month.

So you're probably, if you're going to want someone who's going to be working multiple weeks for you, you might want to renew that.

You might want to start a new subscription.

So let's just say 199 is going to get you.

I think that gets you through a week, 40 hours, I think it's, let's say divided by 4.5, yeah, brings us to almost 9 hours in a day for 4 1/2 days.

So you're getting, you know, 1/2 day on Friday, but I think that's probably OK.

So let's say 199 * 4 is what it would really realistically cost you to match the amount of output that you're getting out of an assistant.

So 796.

And we've got this number here, so 796.

So yeah, this is what it costs for a human.

This is what it costs for this AI to be working full time for an entire month.

It'll be interesting though.

I, I'd, I'd, I'd love it if someone made a video about this, doing a side by side comparison,

having this very low level employee who's obviously gonna need a lot of hand holding, a lot of input, a lot of help with problem solving for some of these tasks because many of them are probably going to be outside of their skill set.

If they're young as well, they're not going to be very confident in going ahead and solving problems for themselves.

But I think Suna from Cortex AI is probably going to be the same way, right?

I mean, the examples that they've shown here are very impressive.

It's able to do quite a bit, but at the same time, those are hand picked, right?

They're cherry picked examples.

And obviously they're not going to show us all the times that the model failed, all the times that the model was not able to solve the problem, even though it's agentic and it should be able to overcome those obstacles.

And in fact, we even saw one example, right, of a person who had to, the user had to step in and sort of say, OK, can you share the chart with me?

And then finally, Suna gave them the result that they wanted.

But if that person had stepped away for an hour, let's say they had to go do something and they thought that Suna would be able to accomplish the task.

And it wasn't.

They would come back after an hour and say I wasted so much time because the AI got stuck on this issue.

Now, that's not to say that the same thing wouldn't happen with a human.

Oftentimes when you delegate a task to a low level person, they might get stuck and sometimes you're not available to give them the solution or help coach them through how to solve the problem.

And the difference here, of course, is that a human being can call you.

I don't think Suna has the capability yet at least of calling the user and saying hey I'm facing this issue.

Can you please provide some input so that I can continue?

Although that would be an interesting functionality, this is open source so someone should definitely use that or or build that functionality out.

It's also worth noting that this is the most expensive way to use Suna and that the self hosting option is available.

So if you want to host Suna in the cloud or who knows, maybe you have a bunch of GPU sitting around that you're not using, then you can do it for much cheaper.

And if you want to try Suna out for yourself, all you have to do is head over to suna dot SO slash dashboard and you can take it for a spin.

But that's it for this video.

I really hope that you enjoyed it.

If you did leave a comment down below.

In fact, even if you didn't leave a comment down below and let us know how we can improve these videos going forward.

And also if you would be so kind, hit the like button, hit subscribe and I will catch you in the next one.