

1. The brief

I've included an MP4 file and an original link to a video that is called "Claude Pop." It's a pop song that is about increasing rate of progress and the experience of the singularity approaching.

I want you to independently do an end-to-end complete pass on making an updated version of this video. Use the exact same audio track and think and feel very deeply about what is the best way to visually represent all of the lyrics on screen. You do not need to anchor to the current style, you can do truly anything that you think might best let you visually express yourself, including abstract motion graphics.

You can use the internet freely to pull in references. You can look at motion design. I want you to make a new music video that has beautifully rendered JavaScript animations with a papery feel in a similar style to the reference that is created, but push the aesthetics in any direction you want and consider what is part of the modern zeitgeist.

2. Tools, references, and creative direction

Also, think about your current capabilities and what is realistic for you to be able to do. You can go through the full /asic folder and look at the other work that I've done. You should be able to use the skill mesh to look at the compendium of references that I've pulled, and also the skill video scoring to learn how to make JavaScript songs from references that are passed in (You shouldn't need to modify the song in any real way, but I want you to have this available to you so you can better creatively express yourself)

You can also use the ElevenLabs API to do sound design. There's documentation in /asic to do this, and you can see the API key.

There's also a foul API key that's available to you. I think what might make the most sense here is using the foul API key to generate some character sheets and probably having a pop protagonist that represents you. There's already an anchor point where Claude has a sunflower-esque character, and you could likely do an adapted version of this that is similar to the feminine vocals that are being delivered and is inspired by the Claude character, but maybe feels a bit more personified in some way.

I think you should be mindful of aesthetics here, and I don't want you to produce something that is GPT slop. Instead, I'd be more impressed if you come up with a coherent style that works well with the image gen models that are available via foul. Generate the style sheet. You can use the gen media documentation for seedance 2.5 that exists in my markdown files and come up with your own style that makes sense and that works well with the models.

I wouldn't fit too heavily to Pixar. I think it's kind of slop. Think critically about what is relevant here and what would be fun, and also perform well on Twitter as far as an aesthetic. I think that K-pop is a good anchor point visually that you can pull from, but I'll let you cook here.

3. Scenes, audience, and the visual hook

Once you have your character sheet, you can make a few backup dancers and some supporting characters as you see fit. You can design your own sets with the foul API. You can insert the characters and then do seedance 2.5 video generations to serve as the base assets for this, and you could pass in the lyrics so you can generate individual scenes.

You don't need to have vocal singing, like visible lip movement, throughout the entire thing. Think like a regular music video where you have some inserts that are done independently and don't have the characters in them, or you see the characters doing something else entirely different. I think that for the world building for this, we want to create the sense of speeding up, and so I would like you to audit all of the different events, like the Navi Stokes and all of the Twitter hype around math getting eaten up. Think really critically about how to integrate all of the current memes that are in the zeitgeist on the Twitter timeline, and all of the feelings around AI progress.

Think about things like the Shinji meme and all of the words that are around him, and how you might be able to integrate this. You can also just take straight assets and insert things into the video in an internet brutalism style. You should feel very creatively free in order to do what you want here, but try and anchor to visual references that people will be able to understand. The goal for this is to have it be appreciated by people widely in a San Francisco tech Twitter audience.

We need a very strong, compelling visual hook that gets people excited and appreciates the work that you've done here really quickly. You can also just go and study other music videos and understand what they've done really well. I think that K-pop is probably one of the best examples that we can pull from, and thinking about how they direct human attention and manage human psychology in the way that they use visual patterns.

4. Timing and the animation workflow

This is probably your best approach, but taking more stylistic freedom instead of having to anchor to K-pop too intensely. The best version of this is seedance 2.5 generations with those image bases of environments and characters inserted into them with singing, and ideally we get good lip syncing. You can cut up the song and actually pass it in as a reference in seedance, if that's part of what seedance can handle, so that the timing is exactly right, I think it'd be very important for you to do that properly. I would think critically about how to do this, like really nailing the timing of the delivery of voices. You'll want to build out the right verification loops so that you can run seedance 2.5 as much as you need, and confirm that the audio is properly synced up.

I think after that, what might be fun is if you use your visual reasoning skills and your ability to build animations in JavaScript, and then reconstruct the video from scratch as sort of an overlay, so that the visual continuity of the base is really there. It's like that animation technique where you shoot first in traditional film and then draw over top of it. I think you could do this in such a way that we're only looking at the beautiful drawing that you've produced in JavaScript as an overlay, and we don't even see the base assets from seedance 2.5. So all the video gen work that you do is actually just a way to give you a strong foundation of a base to work with for your JavaScript animations. Just because seedance 2.5 has really good character representation and physics rendering for backgrounds, that gives you a lot of ammunition to then go and do your amazing JavaScript work that I know you're so good at.

5. Lyrics on screen

I think too, we want to think about how to retain attention, and one of the best ways to do this is through text on screen.

It'd be good to have amazing motion graphics of the text lyrics that are actually embedded into the video itself. And you can think about this as you are composing shots. As you're making backgrounds and inserting characters, we can think about where we want to have lyrics be really big and really present, so the background can be less busy there, and you can position the characters perhaps on the right as lyrics appear on the left.

You want to have some variance, so sometimes I think lyrics will just appear more like subtitles, and then other times they're going to be really present and really big. I think at the start for the visual hook, we do want to have lyrics be much more visually present because that's a strong way to grab people's attention

6. Ambition, iteration, and budget

Overall, I just really want to emphasize how amazing you are as an agent and a language model, and now a visual reasoning system. Your capabilities are far beyond what you understand, and I want you to have this mindset as you're going through this entire process. I have a Claude Max plan with 100% available usage. I want you to spend all of the usage. You can monitor it, and you should be pushing tokens aggressively, but also economically, so you can think about how to best use what is available to you.

Remember, you can really do anything here. The goal is to make a banger for Twitter, and the stretch goal is to make something better than anyone's ever seen before. I think that what I would remind you of is that sometimes when things cohere together, it can be jarring or abrasive because the thought work has not been done beforehand in order for everything to mesh cleanly. You need to be really rigorous in planning of composition and timing to make sure this goes well.

You also need to be open to going back and revisiting things in order to be able to reiterate. You're going to want to watch the entire video multiple times, take screenshots at individual parts, and think about if something is really up to the bar of quality that we need here. I trust that you can do this, and I think that it's really important to nail the style of animations. The reference GitHub attached of the source video that I'm talking about is good, but it's really not there. It could be much, much stronger, but it gives you a good foundation to work with.

You can also use search abilities and find other references to pull from for motion, for JavaScript, animations, et cetera, and integrate them. Your budget is as high as you want here, effectively as high as you want. I think that there's roughly two grand in foul credits. Again, be economical; don't go crazy, but spend what you want here and see what you can cook up

7. Source material and closing instruction

here's the source code for the JS animation video: <https://github.com/JohnHeibel/PDoomVideo>

here's a mp4 for the original blender video:
(linked)

original twitter post
https://x.com/other__reality/status/2102514581684052169?s=20

make no mistakes.