MEETING BRIEF · SUNNY SPADE

Sunny Spade: The AI Production Approach

My findings on the AI side, and the route I think we should take.

— Where I'm at

A lot of this week has been broad exploration to determine what gives us the most control. I've been testing different tools, researching approaches, and trying random things to see what holds up. Where I've gotten so far:

- A style-DNA setup and a multi-model prompt approach, all pointed at one consistent look.
- Some first video passes generated from styleframes plus a face-reference set, feeling out character consistency and a camera push.
- A painterly finishing rig in After Effects: brushstroke and boil, bloom, optical dispersion, soft-focus edges, grain. Reusable on any shot.
- A true-scale 3D blockout of a scene in Cinema 4D with the camera set, to test the 3D-first idea below.

— The route I'd recommend

Here's the pipeline I think fits. A person makes the call at every step, and AI takes the friction out between them.

1. Script. A shot-by-shot written script for the piece. That's the brief. Nothing generates before it. Obviously an AI with all of the vision and language and files and audio fed to the AI to help in the process. AI models usually have features where you have a file and can work along with them highlighting certian things, getting feedback, etc. But you are the one in contol.

2. Boards and animatic, in parallel. A few keyframes of the big story beats get drawn. At the same time, a handful of quick hand-drawn board versions get cut into fast animatics. The speed is the point: real options to look at instead of debating one.

3. Pick the direction. The direction gets chosen: what to take and where to push it.

4. Finish the boards. Depending on the shot, the full storyboard gets finished by hand, or AI takes it about 90% of the way to the chosen direction and the finishing touches happen wherever Ben is comfortable. Either way it lands on a human eye.

5. 3D block the shot. The scene gets built in Cinema 4D or Blender straight from the approved boards: real geometry, real scale, real camera. I have had Claude get me 80-90% of the way there, creating files, organizing the folder structure, creating Cinema 4D template and layer organization, getting all the render settings (frame rate, length, etc) adding the storyboard data into it, artwork, camera, movements, notes, etc into the project so whoever is setting it up knows exactly what is happening putting basic props and models in the scene, camera, etc based on the storyboard sketches.

6. Viewport to styleframe. A clean flat-shaded viewport frame, just dropping in basic colors to different parts of the scene (helps describe to the AI what things are what whats happening) a true 16:9 with the polygon edges showing so depth reads, goes to an image model (Nano Banana 2, GPT Image 2, etc.) to get dressed into a finished styleframe. There's flexibility in how this gets driven. It can match the viewport exactly, or a final frame that someone drew or created can be used as reference, even if it's a different angle or setup: point the AI at that look and tell it to build this scene at this angle with this composition, plus any references for anything else in the scene. Either way, the 3D hands the AI the world space, depth, angles, and scale of everything.

7. Animate in 3D or even 2D. This needs flexibility. The default is to at least animate the camera move in 3D, and depending on the shot, more. For a simple scene, say a camera push with a window opening, the 3D might just be the camera and that window. The other option is to skip building it in 3D and see how much can be directed straight in AI. In general, though, default to animating at least the camera in 3D. I have also animated basic shapes in after effects to tell the ai how to move a certian object in the frame.

8. AI brings it together. The 3D animation and the final styleframe from step 6, as well as any references to particular objects in the scene all get handed to the AI. The styleframe becomes the first frame, and the AI animates it to match the viewport move exactly. Everything left after that is the small stuff, the subtle in-scene details, and there's more control over those the more that goes into the viewport.

— Why the 3D-first part matters

The viewport gives the model the world space, depth, angles, and exact sizing, so frames match shot after shot and we can hold continuity across a whole video instead of one lucky clip. That's the piece that lets us actually control the look while AI does the fast, wide part: brainstorming, boards, and dressing.

— Tooling

The canvas: Coveboard (if its not buggy when we start, still building, or other programs that pull in all the models through their API's). I'd run all of this on one canvas per video, with a master canvas that has all assets, that pulls in every kind of model and carries references, boards, generation, look dev, and animatics through to a cut. That's the tool I've been building, and the animatics on my site are already cut in it. Coveboard is still in progress and I'm finishing the collaboration layer now, so part of my near-term is getting it ready for the team to test on.

Image models for styleframes: Nano Banana 2, GPT Image 2, etc, run off the viewport frames.

Video / motion: I'll pick this through testing and judge it on control, consistency, rights, cost, and reproducibility. That's the written recommendation I owe in September.

— The open question I want us on: character animation

This is the piece I haven't had time to crack yet, and it's the one that matters most since it's a music video and we see the character in most shots.

A couple of directions:

- **Reference-video acting.** Someone records the movement for a scene, even on a phone. Then we tell the AI where the character starts and the path to take, like start front-left, walk behind the center object, end back-right, etc and it uses the reference acting video to exactly product the acting. The question is how well a model actually holds that staging. It needs some tests.
- **Viewport character, then style-match.** Animate the character in the 3D viewport, maybe with a studio doing the high-end animation, then let AI match it into the final look. Most controllable, but much more time consuming and costly.

What I'd suggest: first test someone just acting a scene out and see how much control we get. That tells us fast whether we lean on reference-acting or budget for studio character work. Would love to talk this through.

— Integrations that'd help

- **Photoshop integration:** send boards straight to Photoshop and back, so the finishing pass is frictionless.
- **One-click 3D-to-Coveboard viewport feed:** a button in Cinema 4D and Blender that sends the current viewport as a clean 16:9 at 1K to 2K, flat shades with polygon edges showing. Screenshotting works today, this would just make step 6 instant.

— Timeline

Two things running at once: hitting the benchmark dates, and getting Coveboard ready for the team to actually use.

DATE	WHAT I DELIVER
Now to Aug 3	Lock the benchmark scene, tools, and measurement plan. Get Coveboard usable for team testing.
Aug 3	Benchmark shot breakdown: AI steps, assets, tools, hours, and a fallback plan.
Aug 21	Benchmark assets plus workflow documentation.

Aug 24	AI black-and-white storyboard first pass from the approved treatment and reference package.
Sep 4	Internal animation test and source package done; support the animatic audit.
Sep 11	AI tool recommendation: control, rights and security, cost, reproducibility.
Sep 15 to 18	AI color-propagation pass; support final board approval.
Oct 19 to 23	Support the studio-versus-internal comparison for the path decision.
Nov 6 on	Execute the chosen path's hero assets.

Bottom line

Build a 3D-controlled, human-directed pipeline where AI gives us speed and a ton of range on ideas, boards, and dressing, and we keep shot-by-shot control over the film. Whatever we nail on this one sequence, every chapter after it can inherit.

How I think about AI

I know everyone feels strongly that this stays human. The short version:

1. **The brief comes first.** AI without a brief just gives you options. A brief gives you a direction.

2. **Reference conditions the machine.** I generate against curated reference packs, so consistency is built in instead of hoped for.

3. **Generate wide, kill hard.** Tons gets generated and most of it dies. The kill rate is where the taste lives.

4. **Traditional tools direct it and finish it.** Cinema 4D, After Effects, Photoshop, Illustrator, Figma, and any other design program are where a shot actually gets art-directed and finished, held to the same bar as anything hand-built.

The line: AI never picks the idea, the layout, the cut, or the final grade. It gives us raw material at a volume no team could hit on its own, and every piece of it goes through the same judgment as a frame we made by hand.

So my one-liner: **use AI for speed and for a huge breadth of ideas and iteration, and keep taste and final control fully human.**