

SketchDynamics: Exploring Free-Form Sketches for Dynamic Intent Expression in Animation Generation

Boyu Li

Division of Arts and Machine Creativity
The Hong Kong University of Science and Technology
Hong Kong SAR, China
blibr@connect.ust.hk

Zeyu Wang

Computational Media and Arts
Hong Kong University of Science and Technology
(Guangzhou)
Guangzhou, China
Computer Science and Engineering
Hong Kong University of Science and Technology
Hong Kong, China
zeyuwang@ust.hk

Lin-Ping Yuan

Department of Computer Science and Engineering
The Hong Kong University of Science and Technology
Hong Kong SAR, China
yuanlp@cse.ust.hk

Hongbo Fu*

Division of Arts and Machine Creativity
Hong Kong University of Science and Technology
Hong Kong, China
fuplus@gmail.com

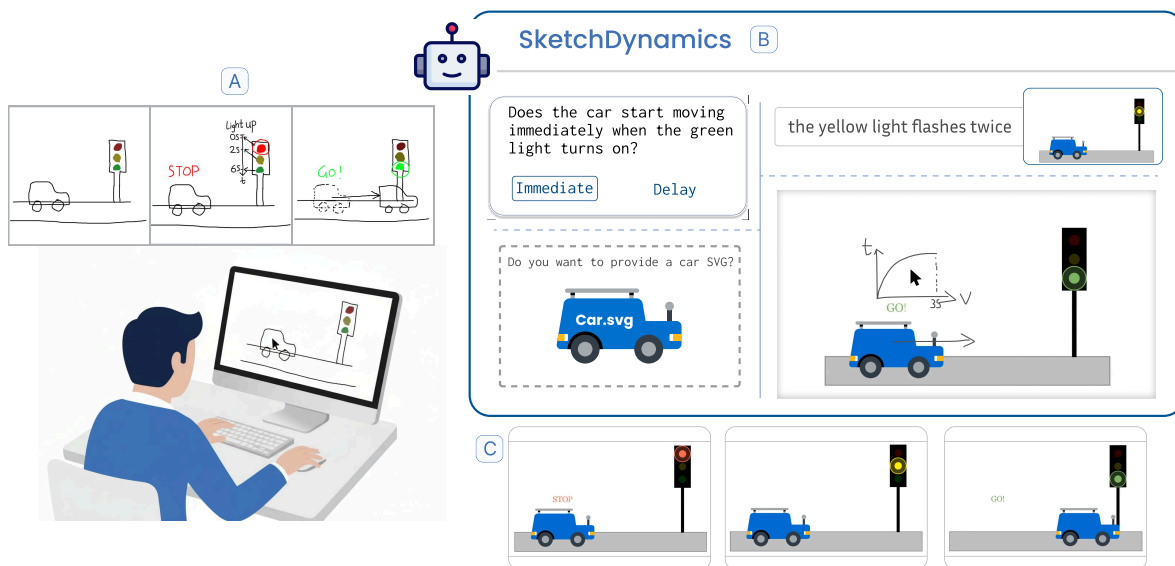


Figure 1: Example interaction with SketchDynamics. (A) A user sketches a storyboard of a traffic-light scene (red stop, yellow flashing, green go). (B-Left) The system proactively raises clarification questions (e.g., whether the car starts moving immediately or after a delay, and whether a car SVG should be provided), and the user responds by selecting options or uploading assets. (B-Right) The user further refines the generated frames by specifying that the yellow light flashes twice and by sketching a velocity-time curve to indicate changes in the car's speed. (C) At the bottom, the final video clips illustrate the produced animation sequence.

*Corresponding author.



This work is licensed under a Creative Commons Attribution 4.0 International License.
CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2278-3/26/04
<https://doi.org/10.1145/3772318.3791071>

Abstract

Sketching provides an intuitive way to convey dynamic intent in animation authoring (i.e., how elements change over time and space), making it a natural medium for automatic content creation. Yet existing approaches often constrain sketches to fixed command tokens or predefined visual forms, overlooking their free-form nature and the central role of humans in shaping intention. To address this,

we introduce an interaction paradigm where users convey dynamic intent to a vision–language model via free-form sketching, instantiated here in a sketch storyboard to motion graphics workflow. We implement an interface and improve it through a three-stage study with 24 participants. The study shows how sketches convey motion with minimal input, how their inherent ambiguity requires users to be involved for clarification, and how sketches can visually guide video refinement. Our findings reveal the potential of sketch–AI interaction to bridge the gap between intention and outcome, and demonstrate its applicability to 3D animation and video generation.

CCS Concepts

• **Human-centered computing** → *Interactive systems and tools*.

Keywords

Free-Form Sketch, Dynamic Intention, Creativity Support

ACM Reference Format:

Boyu Li, Lin-Ping Yuan, Zeyu Wang, and Hongbo Fu. 2026. SketchDynamics: Exploring Free-Form Sketches for Dynamic Intent Expression in Animation Generation. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3772318.3791071>

1 Introduction

Sketching is one of the most accessible and intuitive ways for users to visually convey dynamic authoring intent, i.e., how they wish elements to change over time and space, particularly in expressive content domains such as film and animation [4, 7]. Traditionally, sketches are used in two complementary ways to express dynamic intent. The first is as storyboards [27]. From early Disney animation [51] to contemporary film production [52], sketches are combined with scripts to depict keyframes and narrative flow. The second is as free-form annotations, where marks such as arrows, dashed lines, and curves are used to suggest temporal evolution or spatial transformations [13]. Together, these practices enable people to represent sequences of motion and interaction without requiring full implementation, offering a lightweight yet powerful foundation for subsequent production.

Building on these practices, research has explored how to interpret users' authoring intent from sketches and translate that intent into final content. For example, generative models can add and remove objects in a video based on a sketch drawn on a single video frame [44, 45], or generate an entire video sequence from a storyboard sketch [43, 75]. Beyond generation, sketches have also been explored for incremental authoring [14, 31, 34], where users construct animations step by step and each stroke is recognized [66] or matched to a predefined command [14], triggering corresponding system actions (e.g., appearing, translating, or morphing) that gradually accumulate into a complete animated result. However, these approaches often constrain the expressive potential of sketches, limiting them either to static layout representations or to predefined visual forms that encode specific animation commands. As a result, sketches shift from being an approach for open-ended creativity to functioning as rigid control tokens, missing the opportunity to harness their flexibility and spontaneity for expressing free-form dynamic intent, such as motion and interaction [63].

We aim to investigate how free-form sketches can effectively capture dynamic intent and be leveraged as a foundation for generating dynamic content. Fortunately, recent advances in vision–language models (VLMs) have introduced powerful visual understanding capabilities, including the ability to interpret abstract and informal sketches. Such capabilities have been demonstrated in recent works such as CodeShaping [69], though primarily in the domain of code generation. We select motion graphics as our starting point, as they offer accessible yet expressive animations with clear motion dynamics, making them a practical setting for investigating the alignment between sketches and generated animations. Specifically, we focus on explainer-style motion graphics with the roles of teaching aid and visual discourse [11], which are widely used across diverse domains and age groups [29], prioritizing the clarity of motion over visual realism to effectively communicate concepts [24].

Building on this foundation, we conduct a three-stage study to explore how sketch-based interactions can drive dynamic content creation. In the first stage, to study sketch expressiveness and VLM interpretation, we provide a unified web interface where users can draw free-form sketch storyboards to generate animated videos through code-based rendering. The results highlight the strong potential of free-form sketches to convey animation intent with minimal, intuitive input (see Fig. 5 for a collection of such examples), though their inherent ambiguity often leaves user intention underspecified. To address this, the second stage introduces a clarification cue that adaptively matches sketch ambiguity level, offering lightweight interface interactions for users to intervene in the VLM's interpretation. Findings suggest that clarification turns sketch ambiguity into a resource for intent articulation, but some intentions only become concrete when users view the generated output and refine it accordingly. Therefore, in the last stage, we develop an intuitive visual editing approach that allows users to express their intent on top of the generated result directly. The refinement mechanism closes the gap between system output and user intent, allowing precise control with low effort.

In general, rather than pursuing high-quality video results, our contribution lies in reactivating the expressive freedom of sketches, which are casual and intuitive drawings that anyone can produce to freely convey dynamic intent without any constraints. Here are our contributions:

- We propose a sketch-based dynamic intent expression paradigm where sketches are not predefined forms or fixed commands, but as open-ended prompts, in which sketch ambiguity, VLM interpretation, and user intervention collaboratively shape dynamic content.
- We implement SketchDynamics, a proof-of-concept instantiation of this paradigm, using free-form sketch storyboards to generate explainer-style motion graphics, supported by an adaptive clarification interface and multimodal input for iterative refinement.
- Through a three-stage user study, we validate the effectiveness of this paradigm. We demonstrate its capacity to support diverse free-form sketches, enabling users to articulate precise dynamic intent with minimal interaction.

2 Related Work

2.1 Understand Authoring Intent in Sketch

Sketching has long been valued for its ability to convey authoring intent through just a few strokes, making it a widely adopted practice in early-stage ideation across fields such as architecture [38, 72], art, animation, and user interface design [34]. Numerous HCI researchers have explored how sketches can produce meaningful results for rapid prototyping, emphasizing their immediacy and fluidity as a medium [10].

A major part of this research treats sketches or single strokes as predefined symbols or commands [55, 66, 71]. With simple stroke recognition methods [19] like the \$1 recognizer [66], the system could classify sketches and trigger specific operations. A seminal example is SILK [34], which allowed interface designers to hand-draw UI widgets and manipulate them with drawn strokes. Lineogrammer [73] extended this idea by beautifying sketches into clean diagrams and introducing more flexible action sketches, such as joining lines or lasso selection, enabling intuitive manipulation without tool switching. Using a similar technical pipeline, K-Sketch [14] introduces an interface that enables novice users to create 2D animations through informal sketches. These interfaces provided immediate feedback by recognizing drawn shapes as components and commands, preserving the speed and informality of sketching. However, as shown in Fig 2, these systems interpret sketches through rigid one-to-one mappings, translating recognized symbols (e.g., arrows, lassos) directly into predefined graphical transformations (e.g., translation, rotation) or specific animation effects (e.g., morphing, fading). To support such mappings, designers needed to predefine gesture categories and their corresponding actions, meaning that users had to learn and conform to the system’s sketch–action vocabulary. As a result, while interaction felt informal and fluid, the semantic space of motion intent remains constrained to what had been explicitly encoded, preventing sketches from expressing abstract, causal, or contextual forms of motion intent.

Therefore, some research has moved beyond hand-coded sketch mapping, learning to infer higher-level intent behind a sketch [26]. Instead of asking “what symbol was drawn?”, these systems learn to ask “what does the user mean by this drawing?” In user interface design, Swire [22] exemplifies this shift by training a neural network on thousands of drawn UIs: it can interpret a free-form sketch of a UI layout in terms of its semantic similarity to existing interface examples. Samuelsson et al. [56] conducted an elicitation study with software engineers to understand how sketches might express code editing commands. Building on such insights, Code Shaping [69] introduced an interactive environment where developers draw directly over their source code to invoke edits. Behind this is the ability of current vision–language models: rather than simple classification, such models can understand quite abstract sketches, interpret them, and generate new content.

However, VLMs cannot always understand what users want to express in free-form drawings. Beyond the limitations of current algorithms for sketch understanding, an important reason is that sketches inherently carry ambiguity [63]. Some prior work has addressed this by imposing additional constraints to guide the AI’s interpretation. For example, Code Shaping ultimately introduced

“*command brushes*,” a restricted set of sketch shapes each mapped to a specific edit operation, in order to improve recognition reliability. While this approach improved accuracy, it effectively reintroduced a symbolic shorthand, thereby limiting the expressive freedom of sketching. In contrast, we view ambiguity not as a flaw to be eliminated but as a feature that can drive creative workflows [17]. Our approach further explores engaging users in co-interpretation, where the system proposes an interpretation and the user iteratively refines or corrects it. This paradigm is especially well-suited to animation and dynamic media authoring, where sketches are inherently provisional and iterative refinement aligns with how creators naturally explore temporal ideas [13].

2.2 Sketch as a Generation Prompt

Sketches have been widely explored as prompts for generative algorithms due to their intuitive and expressive nature. Since sketches provide a quick and intuitive way to communicate spatial structure and layout, they have been widely used to guide the generation of static content such as images [9, 41, 67, 74], 3D models [48, 49], websites [2, 58]. Building on these applications, recent work [25, 32, 43, 45, 75] has begun to explore the use of sketches in dynamic content generation, where conveying temporal and semantic change becomes essential. In pixel-based video generation, sketches can be drawn directly onto images to indicate motion trajectories or directional flow, serving as intuitive cues for generating animated sequences [32]. Alternatively, sketches are used as masks or editing guides, allowing users to specify which regions of a video should change and how those transformations evolve over time [45, 61, 68].

While sketch cues overlaid on images are effective for indicating short-term motion or localized edits, recent works [12, 16, 18, 35] have begun exploring sketch storyboards to support longer-range structure and narrative purely through sketch. Traditionally used in early-stage film and animation ideation, sketch storyboards offer a sequence-based representation of motion and narrative intent, making them a natural fit for guiding generative algorithms [20]. Different systems adopt sketch storyboards in domain-specific ways. SketchVideo[43] treats sketches as static illustrations of each keyframe’s visual structure, using carefully drawn outlines to guide photorealistic image generation. Sketch2Anim[75], on the other hand, relies on a predefined sketch format, typically a stick figure combined with a trajectory, to represent skeletal movement across frames. In both cases, the sketch input is tightly constrained and tailored to the specific task.

While effective for domain-specific tasks, these constrained sketch representations fall short for motion graphics video creation, where compositions often feature multiple vector-based elements moving in abstract, stylized, or non-physical ways [5, 33]. Such sketch formats impose rigid structural or pose constraints, limiting users’ ability to articulate complex and unconventional animation ideas [69]. To overcome these limitations, we adopt free-form sketch storyboards that support open-ended drawing without constraints, enabling richer expression of spatial, temporal, and stylistic intent. This flexibility allows sketches to function not as fixed blueprints but as adaptable cues that guide generation. Building on this foundation, our system further supports an interactive, iterative authoring

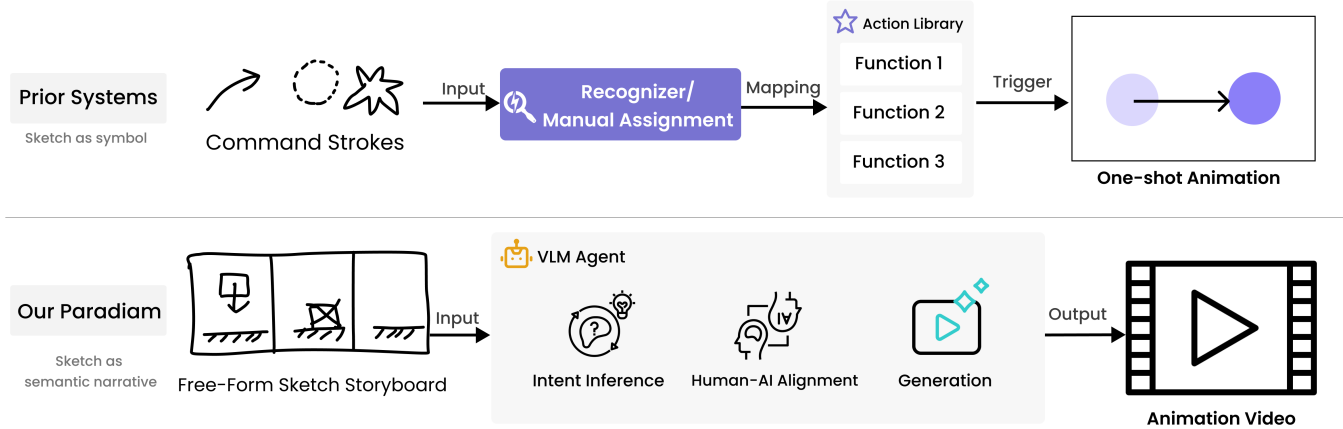


Figure 2: Prior systems [14, 29–31] treat sketches as symbols to be mapped to animation commands (Top). In contrast, we treat sketches as semantic expressions of motion intent (Bottom), leveraging VLM-based commonsense reasoning and human-AI clarification to interpret why and how motion happens.

process, allowing users to refine their animations over multiple rounds rather than relying on one-shot generation.

2.3 Vector Animation Authoring for Casual Creators

Creating motion graphics or vector animations traditionally requires substantial expertise and professional tools such as Adobe After Effects [15]. HCI research has long sought to lower this barrier by developing more intuitive authoring paradigms. Early sketch-driven systems such as Draco [30] and Kitty [29] let users draw scenes and motion cues directly on the canvas, inferring causal animations without requiring keyframe editing, while SketchStory [36] applied similar ideas to animated data storytelling. These works emphasize continuous, pen-on-canvas interaction over traditional timelines. Building on this, more recent end-to-end systems scaffold the entire workflow for novices: Katika [24] guides users from storyboards or scripts to finished videos via libraries of pre-designed animations, and DancingBoard [8] streamlines motion-comic production with guided steps and automated animation suggestions. Some works focus on more specialized domains of vector animation. In data visualization, tools such as Data Illustrator [47] and Data Animator [60] enable designers to create animated charts as if drawing static graphics, automatically interpolating between states. For text animation, works like The Kinetic Typography Engine [37] and TextAlive [28] demonstrated powerful systems where users could algorithmically define and control motion. Piet [57] further extends this direction by providing targeted interfaces that facilitate color authoring in motion graphics videos. Most recently, AI-driven authoring has emerged: Keyframer [62] and AnyAni [54] use large language models to translate natural-language prompts into editable vector animations, supporting iterative refinement through text or direct manipulation.

Building on this trajectory, researchers have begun to explore how vision-language models (VLMs) can support motion graphics generation. MoVer [50] introduces a motion verification DSL

that encodes spatio-temporal properties of animations in first-order logic, demonstrating how formal reasoning can be applied to ensure correctness of motion behaviors. LogoMotion [46] analyzes vector graphics (SVGs) to automatically suggest motion designs, showing how pre-existing assets can be transformed into animations. These works highlight the potential of AI models to reason about both semantics and temporal logic in vector-based media. However, a gap remains. Existing systems have largely focused on short, self-contained examples with relatively few animated elements, often suited for toy demos or tightly scoped use cases. In contrast, many real motion graphics videos—even simple explainers—already involve multiple objects evolving over time, which introduces additional complexity for authoring and interpretation. Supporting such scenarios requires more direct and flexible visual guidance, enabling creators to express dynamic intent without relying solely on pre-defined templates or textual prompts. In this paper, we explore free-form sketches as a visual guidance mechanism. Sketches offer an intuitive way to externalize spatial-temporal ideas, while VLMs can interpret these abstractions into vector animations. This combination bridges informal creative expression with structured animation output, extending bi-directional editing [21] to non-programmers.

3 SketchDynamics

Free-form sketch storyboards are often used for quick idea exploration in dynamic content creation, and they also hold promise as input for generation. Yet it remains unclear how such informal sketches can reliably convey a creator’s intent across the authoring process. To explore this, we developed a proof-of-concept system, SketchDynamics, which examines how free-form sketches can function as a medium for expressing animation ideas. With this system, users sketch simple storyboards, AI interprets the drawings, translates them into vector-based animation code, and renders them into video. To situate the study in a prototyping context that mirrors early creative practice, we did not require participants to have advanced drawing skills, encouraging them to use simple and intuitive

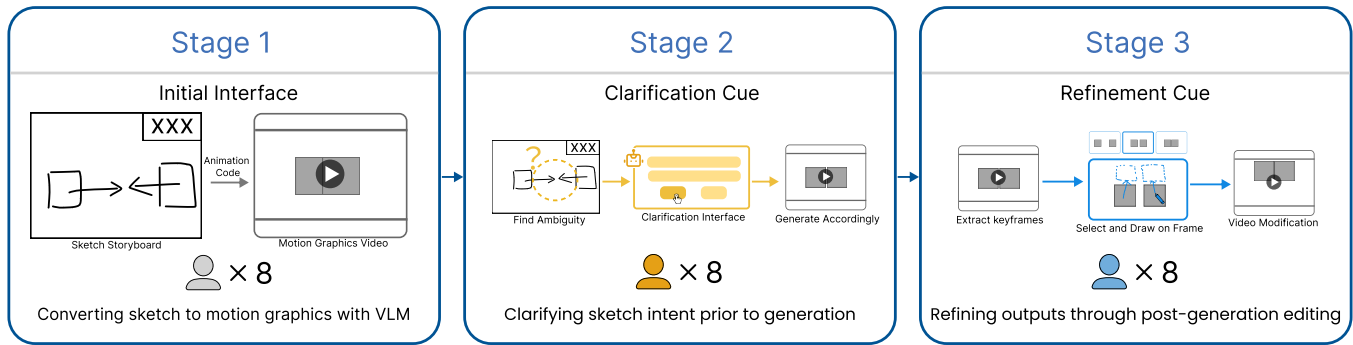


Figure 3: Study illustration showing the three stages of our investigation: (1) sketch input and direct generation, (2) clarification through disambiguation cues, and (3) iterative refinement with contextual editing. Each stage involved eight new participants in the study.

sketches within just a few minutes as a way of communicating their animation intentions.

We design the system iteratively through the following three stages (see Fig 3):

Stage 1: Understanding sketch and interpretation. In this stage, we implemented a unified interface that understood sketch storyboards from users and transformed them into motion graphic videos (Sec. 4.1). We explored how participants naturally expressed animation intent through free-form sketches and how these sketches were interpreted when directly converted into videos. The study findings revealed that participants tended to use highly diverse and abstract sketches for quick prototyping. However, even though the VLM demonstrated impressive capacity to interpret many inputs, it still often failed when confronted with ambiguous sketches.

Stage 2: Investigating ambiguity in sketch interpretation. Through empirical use in the first stage, we identified recurring cases where user sketches were too abstract or ambiguous for the system to interpret reliably (e.g., an arrow could indicate either translation or rotation). This led us to investigate a lightweight way for enabling users to intervene in the machine’s interpretation process. In this stage, we proposed a clarification cue mechanism (see Sec. 5.1), which separated ambiguity into four levels and provided adaptive interfaces for user input. Our study showed that participants generally considered clarification requests reasonable and helpful, noting that such prompts not only reduced misinterpretation but also assisted them in further shaping their intent.

Stage 3: Exploring more intuitive iterative editing. Not all ambiguities could be resolved in the early stage, as some sketches were too underspecified to be faithfully interpreted until a video was generated to provide additional context. In this stage, we investigated how users could refine the output when it did not meet their expectations, introducing a frame-based interaction method (see Sec. 6.1) that combined keyframe extraction and annotation to offer clearer guidance toward the desired result. Our study showed that the refinement strategy was efficient and intuitive, allowing participants to annotate video frames to express intent with minimal effort.

4 Stage 1: Initial Sketch to Animation Authoring

We began with a user study to explore how participants sketched animation ideas and how these sketches were translated into preliminary video outputs. Rather than optimizing for output quality, this baseline stage focused on observing sketching behaviors and interpretive patterns, laying the foundation for identifying issues and guiding subsequent investigations. We chose this approach over an artifact-based analysis of existing online storyboards [14], as the uncontrolled nature of such materials (e.g., varying time and quality) would have limited the validity and comparability of our observations.

4.1 Initial Implementation

To begin with, we designed a unified web interface where users can directly explore how their free-form sketches translate into motion graphics. The interface consisted of three connected views (Fig.4). In the sketch view (Fig. 4-A), users drew on a central canvas using freehand strokes. The toolbar at the top provided lightweight editing functions such as undo, clear, pen, eraser, color, and pen size. Sketches could be exported or imported in multiple formats (PNG, SVG, or JSON data), and entire storyboards could be saved or reloaded. A sidebar displayed all sketches as page thumbnails, allowing users to navigate between frames or remove sketches. Once sketches were created, they were organized in the storyboard view (Fig. 4-B). The interface displayed all sketches in sequence, giving users an overview of how their drawings formed a storyboard. Each sketch could optionally be annotated with short text notes, but the interface emphasized sketches as the primary mode of input. When the storyboard was complete, the user entered the AI video view (Fig. 4-C). Here, a single click on [Generate Script] sent the storyboard to the model, which returned executable code. By pressing [Render Preview], the code was compiled into video and displayed in the built-in player, allowing the user to immediately check the generated result.

Behind the interface, we carefully engineered the sketch-to-video translation process to support this interaction. The first step was ensuring that the VLM could interpret sketches and generate corresponding animation code. To this end, we designed prompts with paired sketch descriptions, code examples, and formatting rules,

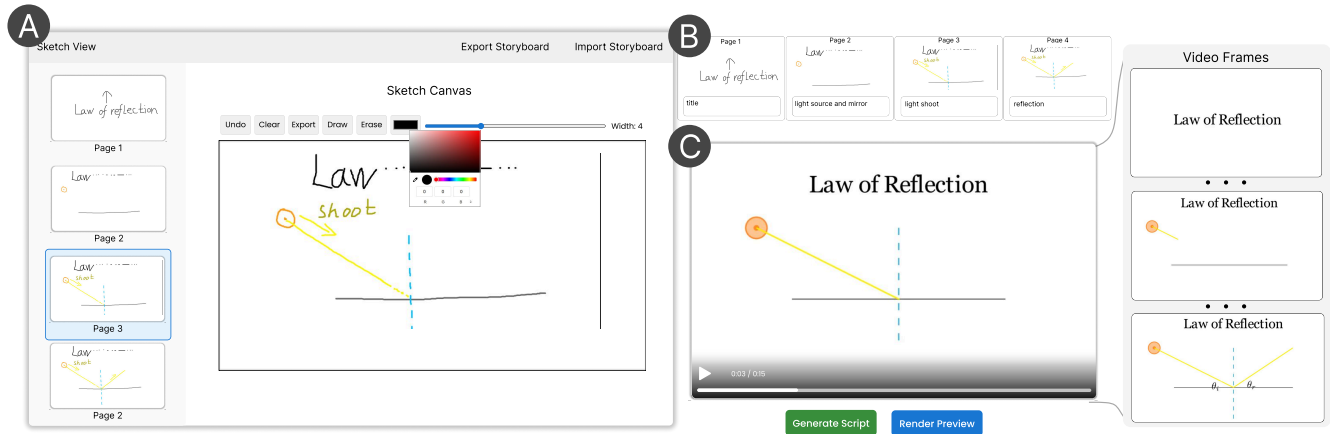


Figure 4: Interface design for the first Stage. (A) The user sketches a frame-level storyboard to explain the law of reflection, drawing incident and reflected light rays on a mirror surface. (B) adds brief notes at the bottom to specify the cues, and (C) generates a preview video from the storyboard.

asking the model to output executable scripts. A central challenge was how to represent the storyboard so the model could reason across sketches effectively. We tested two input strategies: (1) feeding each sketch individually with its script, and (2) compositing all sketches and text scripts into a single storyboard input. Results indicated that the second approach produced more coherent and temporally consistent animations, likely because the model could better reason across sequential sketches. This strategy was therefore adopted in subsequent development. The second step was to render the generated code into video. We instructed the VLM to output executable Manim [1] Python code, which was then compiled into animations. Manim, a declarative animation library widely used for educational explainer videos, provided the rendering backend. This code-first design not only enabled fine-grained control at the script level (e.g., timing, trajectories, styles) but also ensured that animations were built from scalable vector-based primitives, making the video extensible and well-suited for future iterative editing.

4.2 Study Design

Participants. Consistent with previous research [69], we recruited eight participants (P1–P8; five male, three female) from a local university, with an average age of 26.2 years (SD = 4.2). All participants were right-handed and had prior exposure to motion graphics as viewers. Five participants reported previous experience in creating motion graphics using tools such as After Effects [15] or CapCut [6], while the remaining three had only basic video editing experience, primarily limited to simple cutting and sequencing. We intentionally recruited participants with varying levels of expertise to capture perspectives from both experienced and novice creators, while ensuring that all had sufficient familiarity with motion graphics as consumers to provide informed feedback on the generated results.

Tasks. Since our system was designed to support storyboard creation from scratch, we did not impose a narrowly constrained task

or theme. Instead, we asked participants to imagine creating a short explainer video for a concept of their choice. To keep the task manageable while still showcasing the workflow, we provided three guidelines: (1) the video should illustrate a concept that can be reasonably explained through basic shapes, (2) participants should avoid overly complex shapes or animation techniques, and (3) the storyboard should be concise, consisting of around four sketches. Specifically, we limited concepts to logical explanations to evaluate functional accuracy; We restricted visuals to basic shapes to abstract away illustrative details, allowing participants to focus their cognitive effort on articulating dynamic logic (e.g., causality, flow) rather than static aesthetics; and we focused on fundamental animation primitives (e.g., displacement, deformation) rather than high-dimensional tasks like character animation, as their complexity makes it difficult to verify the specific animation intent. We conducted the experiment on a desktop interface with mouse-based interaction, as our focus was not on the precision of sketching.

We restricted visuals to basic shapes to abstract away illustrative details, allowing participants to focus their cognitive effort on articulating dynamic logic (e.g., causality, flow) rather than static aesthetics

Procedure. Participants began the experiment by completing a survey to gather information on their prior experience with video creation and sketching (Appendix. C). Next, they were given a brief overview of the project and shown the core features of the interface. Some example cases were presented to illustrate the expected input and outcome. We specifically presented examples of free-form sketches and encouraged participants to use more flexible sketching methods to express their ideas, with less emphasis on scripting. Following the introduction, participants had a 10-minute exploration phase, where they could experiment with the sketching tools, test the generation of animation previews, and become familiar with the interface’s capabilities. Participants then moved on to the creation phase, where each participant was tasked with creating three separate videos, each representing different content. This meant that

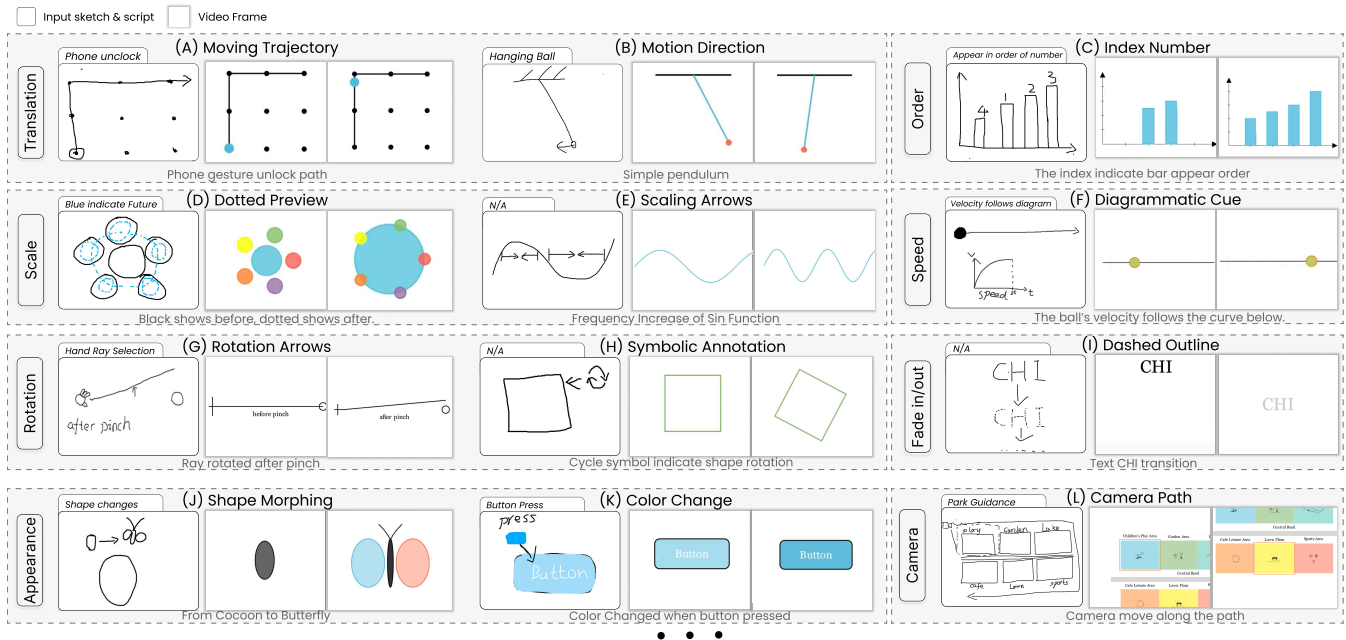


Figure 5: Selected sketches with scripts and animation excerpts from participants' storyboards in Stage 1. Above the sketch is the script (N/A means no script); the gray text below describes the animation. These examples illustrate common ways sketches were used to convey animation intentions (e.g., translation, scaling, rotation, appearance). The categories shown here are not meant as a strict taxonomy but as illustrative examples; in practice, participants' free-form drawings were far more varied and often mixed multiple notations within a single sketch. Some animations shown are refined versions to better reflect participants' intended outcomes.

they needed to create three distinct sketch storyboards and review three generated videos, completing the task in approximately 30 minutes. Finally, the session concluded with a semi-structured interview and a questionnaire, where participants provided feedback on their experience in translating ideas into sketches, as well as their thoughts on the generated videos. The similarity between participants' intended animations and the generated outputs was analyzed using both qualitative feedback and the Alignment metric from the questionnaire.

4.3 Results

Across all sessions, we collected 24 attempts. Not all were successful: five outputs were judged by participants as failures, where the generated video significantly diverged from their intended idea. As a first-stage probe, such outcomes are acceptable and highlight the current gap between sketch expressiveness and model interpretation. Participants appreciated that they were allowed to attempt animations across a wide variety of domains, rather than being constrained to a specific area. The topics participants chose to explore were diverse, covering but not limited to VR interaction, data visualization, phone interaction, traffic rules, celestial motion, algorithm pipelines, mathematical concepts, physical rules, and logo animation. The unsuccessful cases provide further insight into the current limitations of sketch-based intent expression. These included a block colliding back and forth with a wall (P2), a maze-walking path (P2), a rainfall animation (P5), a bouncing ball (P7),

and a grid-warping effect (P8). Based on different animation intents and sketch types, we selectively demonstrate some results from the user in Fig 5.

Expanding the Role of Sketches. In our study, sketches served not only as a way to specify layout or static elements, but also as a medium for expressing animation intent. When attempting to represent dynamics, participants resorted to a wide variety of abstract conventions such as arrows (Fig. 5-B), onion skinning for future states (Fig. 5-D), or index numbers for sequencing (Fig. 5-C). However, these conventions were highly diverse and often idiosyncratic, shaped by individual habits and prior experiences. For instance, P3 emphasized that *"everyone has their own way of drawing arrows, mine is from engineering class"*, while P6 noted that *"sometimes I just use circles, sometimes a line, it depends on what comes to mind first"*. This diversity meant that the same animation could be represented in multiple ways, and similar sketches could point to different intentions. As P4 reflected, *"the arrow could mean movement, or just attention, it depends on the context"*. Indeed, 7 out of 8 participants explicitly remarked that they have their own strategies to depict movement, often *"no regulation"* (P2) or *"quite random"* (P5). This means that free-form sketching opened up a broad expressive space, but one that was highly individualized and semantically unstable, requiring interpretation and negotiation when used to communicate temporal behaviors.

Sketch Interpretation Beyond Expectations. Participants often experimented with sketches that were highly abstract or loosely defined, yet many of these were still interpreted meaningfully by the model. Several participants created inventive sketch forms to convey intent—for example, using gesture line or symbolic marks to indicate motion or transformation. Despite this, 6 out of 8 participants admitted that, at the beginning, they doubted whether such free-form drawings could be understood by the machine. P8 further reflected that some sketches were inherently ambiguous, “*different people could read it differently.*” While some outputs indeed diverged from their expectations, participants often adapted by redrawing their sketches or supplementing them with short text scripts to clarify intent. Interestingly, in 7 cases, participants reported being positively surprised when the system successfully captured the essence of their sketches in ways they had not anticipated. Although the user ratings (Fig. 12) indicate that most participants felt the animations were not a perfect match to their intended outcome, they still perceived the results as capturing their intent to some extent and described them as “*unexpected but meaningful*”.

Semantics over geometry. A recurring theme was that the VLM often prioritized semantic intent rather than literal geometric fidelity. Instead of replicating sketches exactly, the model tended to generate canonical or cleaned-up versions of the intended motion. For instance, when P6 drew a wobbly sine curve, the output appeared as a smooth sinusoid. P6 described this as “*beautified*”, appreciating that the messy strokes became clearer in the final animation. Similarly, P9 noted that their “*rough arrows*” were rendered as consistent trajectories, which they found “*nicer than I expected.*” However, this abstraction also created moments of mismatch. P4, who attempted to specify a very precise trajectory, complained that the output “*ignored the exact angle I drew.*” Likewise, P11 commented that “*it feels like it understood the idea but not the detail.*” In total, 7 of 8 participants remarked that the system often generalized their sketches to capture meaning rather than form, with reactions ranging from surprise and delight to mild frustration. Several (e.g., P3, P5) responded by redrawing shapes or adding text notes to enforce precision.

Balancing detail and abstraction in sketching. Participants varied in how much detail they invested in their sketch storyboards. At one extreme, some treated the canvas almost like a frame-by-frame storyboard, drawing each scene carefully to make their intent explicit. While this approach produced clearer outputs, participants also described it as “*tedious*” (P3) and “*too much effort for a short video*” (P7). In contrast, many preferred using abstract marks that conveyed animation intent more quickly. As P6 explained, this was “*the fastest way to show the idea,*” even though it sometimes made the system’s interpretation less predictable. P8 similarly noted that abstract cues “*save time but leave ambiguity.*” This trade-off reflects a central tension: detailed sketches improve clarity but increase workload, while abstract sketches reduce effort but demand stronger interpretive support from the model. Reflecting this tension, P5 emphasized that although they were willing to rely on abstract sketches, they expected the system to “*just understand the sketch*” without needing additional textual specification.

5 Stage 2: Clarification Guide Interpretation

Findings from the first stage revealed two key challenges in sketch-based authoring. The first concerned the gap between a user’s intended animation concept and the sketch they produced. Participants often relied on highly abstract marks inspired by everyday experience or improvised on the spot; while efficient for the drawer, such sketches were difficult for others to interpret and often left their meaning ambiguous. The second challenge arose between sketch input and machine interpretation, where the limitations of current AI reasoning led to outputs diverging from participants’ expectations. For example, in Fig. 6-B, the VLM might randomly decide whether to include the annotated timeline in the final animation, resulting in an output that does not match the user’s expectations. Similarly, in the case of Fig. 6-D, the VLM might generate an arbitrary curve, whereas the user may have intended to create a visualization based on real data. To address these issues, the second stage aimed to involve users more directly in the interpretation process, allowing them to understand and influence how sketches were translated into animations. Crucially, this was designed without burdening users to provide lengthy textual descriptions. Instead, we developed lightweight mechanisms that let users guide and refine the system’s understanding interactively.

5.1 Clarification Cue

Prior work has highlighted that sketches are inherently ambiguous [59, 64] and that such ambiguity can be a productive resource in design interpretation [17]. Our observations in the first stage confirmed this tendency in the context of animation authoring: participants often relied on abstract arrows, shorthand marks, or improvised symbols to convey motion (Fig. 6), which were efficient for them but difficult for others, or the VLM, to interpret. Even when the intended action was clear, essential details such as duration, scale, or speed were frequently left unspecified. These findings illustrate that ambiguity in sketching arises not only from multiple possible interpretations but also from underspecified parameters and context-dependent symbols, all of which can cause model outputs to diverge from user expectations.

Rather than treating ambiguity as an error to be eliminated, we approached it as a variable condition to be managed. Our design philosophy was to preserve the freedom of freehand sketching while offering just-in-time interventions to clarify intent when needed. This philosophy echoes early ideas of free-form sketch beautification [23], supporting multiple candidates for ambiguity and requesting the user’s interaction. To this end, we adopted a layered prompting strategy that adaptively maps the degree of ambiguity to different levels of system intervention. This draws on the principle of progressive disclosure: users only encounter as much clarification effort as the ambiguity of their sketches demands, thereby minimizing disruption while maintaining interpretability. Concretely, based on previous research [3, 14] about sketch ambiguity and empirical results from the first stage, we developed four complementary mechanisms that align with increasing levels of sketch ambiguity, ranging from low to very high, each offering progressively richer forms of clarification, as shown in Fig. 6.

- **Quick confirm.** For low-uncertainty cases such as uncertain strokes or vague geometry [64], the system generates a primary guess

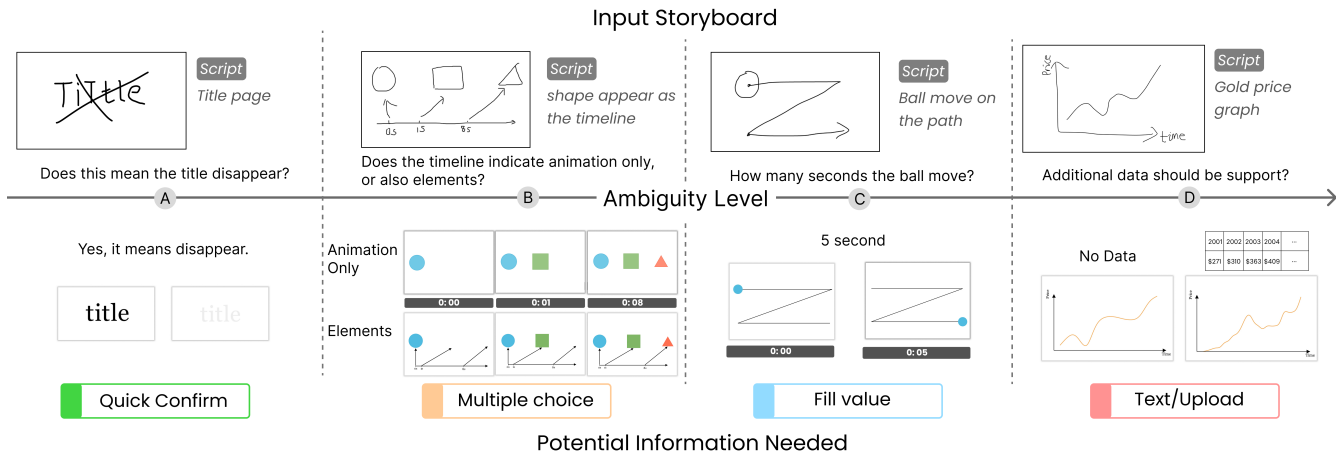


Figure 6: Illustration of ambiguity clarification. From left to right, the examples show sketch ambiguities ranging from low to high. Top: different types of ambiguities in storyboards, Bottom: the corresponding types of information needed to resolve them, mapped to our clarification strategies.

and asks the user for yes/no confirmation (e.g., “Should this line be used as a motion path?”).

- **Multiple choice.** When a sketch could plausibly mean several different things [17], the system presents alternative previews and asks the user to pick the intended one (e.g., a curved arrow could be shown as either “rotation” or “decorative arrow”).
- **Fill value.** For actions that are recognized but lack specific parameters, the interface asks for numeric or scalar values. While many aspects of the video could in principle be parameterized, we prompt only for those with a significant impact on the outcome. For example [14], the system may ask how long a ball should take to traverse a path, while leaving minor details such as fade-in duration to sensible defaults.
- **Text or upload resources.** In cases where sketches are highly abstract or symbolic [17, 63], the system allows users to provide free-form clarification, either as a short text note or by uploading an external reference asset.

After completing a sketch storyboard, users proceed to generation in the AI video view. When they press [Generate Script], if the VLM confidently interprets the sketches, the interface produces the animation code and renders a preview video directly. When ambiguities are detected, however, the interface does not produce arbitrary results; instead, it pops up a clarification panel. For example, if a curved arrow could be read as either rotation or decoration, the interface presents two animated previews for the user to select the intended one. In another case, if a user’s symbol is too abstract, they can upload a custom SVG asset (e.g., a heart icon to represent an object appearing with a “pop” effect). For implementation, as shown in Fig 7, the clarification mechanism is supported through prompt engineering and contextual memory. The VLM is explicitly instructed to request clarifications only when necessary and only on aspects that critically affect the generated animation, avoiding excessive or redundant queries. When the user provides input, such as selecting between preview options or uploading an external SVG asset, these responses are added as a supplementary context for

regeneration. In addition, the system maintains an *intention disambiguation memory*, which records user-provided clarifications and automatically reuses them when the same sketch is edited and regenerated later.

5.2 Study Design

We recruited eight new participants (P1-P8; six male, two female) from the same university, with an average age of 25.4 years (SD = 2.9). All participants had prior experience in creating videos, though none were experts. The study design followed the same task guidelines and procedures as in Stage 1, with the addition of the new clarification mechanism. When a cue requires external resources (e.g., SVG files), we allowed users to directly search online to retrieve the corresponding assets. During the study, we manually recorded participants’ sketches, the generated results, and the clarification cues that were triggered throughout the process, and collected user feedback.

5.3 Results

We collected a total of 87 clarification cues across the 24 creation attempts. In addition, participants were free to revise their sketches and re-generate videos within the same attempt; for analysis, we counted all cues triggered during that cycle under a single attempt. This occasionally inflates the count for sketches that went through several iterations, but more accurately reflects the total clarification effort needed to reach a satisfactory result. In two cases (T5 and T10), participants ultimately judged the final output as misaligned with their expectations, both involving physics-inspired animations such as bouncing or collision. Overall, participants expressed a positive attitude toward the clarification cues. Rather than perceiving them as interruptions, several described them as “helpful checks” (P3) or “a way to steer the system back on track” (P7). Even when additional questions extended the workflow, participants valued

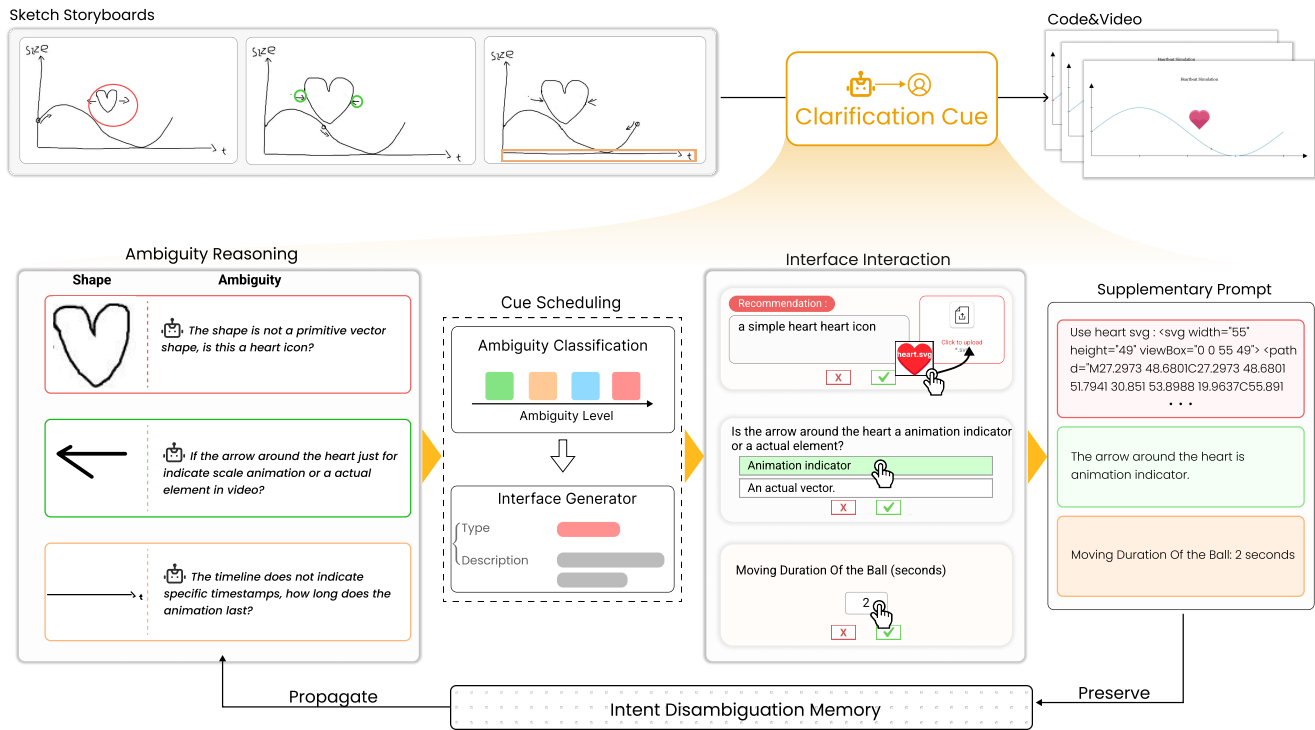


Figure 7: Example workflow of the clarification cue mechanism. The system detects ambiguities in sketches, classifies them by level, prepares them as JSON for the interface generator, and prompts users with lightweight clarifications (e.g., confirmation, choice, value input, or asset upload). User responses are used as a supplementary prompt for regeneration and stored in an intent disambiguation memory to avoid repeated queries in future generations.

the cues as opportunities to make their intent explicit and avoid wasted generations.

Overall Cue Usage. Figure 8 summarizes the distribution of cue types per attempt in the form of a heatmap. The pattern shows that multiple choice cues were triggered most frequently, reflecting the prevalence of sketches with more than one plausible interpretation, which is the most common ambiguity. Quick confirm cues were relatively rare, as the model often resolved low-ambiguity cases confidently without further questioning. Fill value cues also appeared infrequently because the system was designed to ask only for parameters with a major impact on the outcome, and that case is relatively uncommon. By contrast, text/upload cues occurred at a moderate level, indicating that participants frequently drew highly abstract or symbolic sketches (e.g., rough icons) that required extra resources or textual hints. In these cases, the prototype demonstrated its ability to reason semantically about the sketch and prompt the user for relevant supplementary information.

The heatmap also reveals strong variation across attempts. We did not include detailed per-attempt breakdowns of cue types in the paper, as these would be too granular, but the heatmap variation and density of cue usage across the study are meaningful. Some creations, such as T5 and T13, required a dense sequence of cues (up to 5–6 within one attempt), while others (e.g., T7, T14) proceeded with almost no intervention. This disparity highlights how cue

demand was closely tied to the degree of abstraction in the sketch. For instance, P3 commented that “I only needed help when I drew something unusual,” whereas P8 described the cue prompts as “a relief when I was not sure if the sketch was understandable.”

User Reactions to Cue. Participants generally responded positively to the clarification cues, with 7 out of 8 noting that they helped avoid misleading outputs and made them more aware of the ambiguity in their sketches. Quick confirm cues were described as “lightweight” and easy to dismiss, while multiple choice cues were seen as the most necessary since arrows and curves often admitted multiple interpretations. For example, in T10 a participant drew a looping arrow and appreciated being asked to choose between “rotation” and “decorative arrow.” Fill value cues divided opinions: some found specifying duration or speed helpful, while others felt it broke the free-flow nature of sketching. Text/Upload cues, though requiring more effort, were valued when rough sketches stood in for specific icons. In T23, a participant drew a rough star sketch, which triggered a clarification cue suggesting the upload of a pentagram icon. P8 remarked that they were “surprised” the system could interpret such an abstract mark as a symbolic request for a sparkle effect. Overall, all participants felt that the clarification cues made sense in resolving sketch ambiguities. According to the user rating (Fig. 12), compared with stage 1, clarification cues generally align the user’s intention with the final animations, and participants did




Figure 8: Heatmap of clarification cue usage across 24 tasks (T1 to T24). The y-axis displays four types of clarification cues: Quick confirm, Multiple choice, Fill value, and Text/Upload. Each block in the grid represents the total Cue Count (frequency of use) for a specific cue type within a specific task. Color intensity corresponds to this count, with the scale shown in the legend on the right (darker orange indicates a higher count, up to 3).

not find these reasonable checks to be significantly disruptive to their creative process.

The Result after Cue. After clarification, 19 out of 24 attempts produced outputs that participants rated as closer to their intended animation, compared to the initial generation. In 7 attempts, users explicitly described the cue as the turning point that made the result usable (e.g., P6: “*the timing became right only after I entered the value*”). Fill value cues directly improved temporal qualities: in T13 and T21, entering a duration smoothed otherwise abrupt motions. Text/Upload cues often corrected semantic mismatches: in T23, replacing a hand-drawn star with an SVG asset transformed the effect into what P8 called “*finally what I had in mind*.” By contrast, Multiple choice cues mainly resolved directional intent; in T10, choosing “*rotation*” instead of “*arrow symbol*” shifted the system output to match the participant’s storyboard. Failures still occurred: 2 attempts involving physical dynamics remained unsatisfactory even after cues, suggesting limits in the model’s reasoning rather than the cue mechanism. Taken together, the logs show that cues did not simply add extra interaction but often directly determined whether a sketch led to a usable animation.

Workflow Integration. Participants emphasized that the clarification cues were most valuable when integrated smoothly into the overall authoring workflow, rather than appearing as separate steps. In practice, cues were triggered inline during generation and resolved within a few clicks, which participants described as “*not disruptive*” (P3) and “*just part of the flow*” (P6). Several participants noted that without these cues, the system might “*guess wrong*” (P5) or return outputs that “*look nothing like what I wanted*” (P2). By contrast, lightweight confirmations or multiple-choice previews felt like natural checkpoints that improved trust in the system. Importantly, participants highlighted that the cues helped them understand how the model interpreted their sketch, giving them a sense of control. As P8 put it, “*I can see what the AI thinks, and fix it right away*.” Overall, rather than slowing down creation, clarification cues were seen as seamlessly woven into the workflow, balancing automation with user agency.

6 Stage 3: Refinement with Visual Context

In the second stage, we introduced clarification cues to help disambiguate sketches during the initial authoring process. However, once a video was generated, users who wanted to make changes could only revise their sketch storyboard and trigger a full re-generation. This workflow was inefficient and often unpredictable: re-generation could yield results that diverged widely from the original, even if the user only intended to adjust a small local detail. For example, participants described situations where they were satisfied with most of the video but wanted to tweak a single motion, only to receive a completely different animation after re-generation. To address this limitation, Stage 3 explores how users can visually express their intent for modification directly on the generated video, refining the output step by step with iterative local editing.

6.1 Refinement Cue

The refinement workflow begins once an initial video has been generated. Instead of redrawing or modifying the entire sketch storyboard to produce a new result, users can refine the output directly within the video itself. As shown in Fig. 9, the interface processes the video into a set of keyframe images, which act as semantic anchors that provide stable, context-rich entry points for targeted edits. Users can select a target keyframe in which they want to modify the related animation, and express refinement intention in two complementary ways. They may sketch directly on the paused frame to indicate spatial or motion adjustments (e.g., extending a trajectory line to lengthen an object’s movement). Alternatively, they may issue concise textual refinements such as “*fade slower*” or “*loop twice*.” To enable refinement, we implement two key steps. By combining the immediacy of visual sketching with the precision of lightweight textual input, this approach supports a hybrid refinement workflow that flexibly captures both spatial and temporal aspects of animation.

To enable such refinement interaction, we implement two key steps. First, to provide interpretable keyframes as anchors for users to apply intentions, we use the VLM to analyze the initial animation code, detect salient timestamps, and expose them as keyframes.

Second, to provide full context for localized video editing, we supply the VLM with user refinements (sketched on a keyframe or expressed in text), the timestamp of the keyframe, all extracted keyframes, and the initial animation to be modified. The VLM then updates only the relevant portion of the code, and the revised code is rendered into an updated preview, ensuring efficient and focused iteration.

6.2 Study Design

We recruited eight new participants (P1–P8; four male, four female) from the same university, with an average age of 25.8 years (SD = 3.7). None had participated in the previous stages. At this stage, participants were given access to the full set of mechanisms introduced in the paper, including both clarification cues and refinement cues.

The task design followed the same open-ended format as in earlier stages: participants were asked to create a short explainer-style animation of a concept of their choice using free-form sketching. However, to emphasize iterative creation and refinement, each participant was instructed to generate an initial video and then improve it through at least one round of refinement using the new mechanism. Since the refinement process required more time than earlier tasks, participants were asked to complete one animation project within the 30-minute session.

6.3 Results

Overall Use of Refinement Cues. Across the eight participants, we collected 12 edited videos with a total of 55 refinement operations. On average, participants performed 4.6 refinements per task, though the distribution was skewed: some participants stopped after one or two small adjustments (P2, P5), while others explored more iterative editing cycles with up to eight refinements (P7). Most refinements (36/55) were sketch-based, while the others were text-based. This division reflected a functional split: spatial animations (e.g., paths, shape sizes, directions) were usually expressed visually, while temporal refinements (e.g., timing, repetition) were handled with short text cues.

Efficiency and Locality. Participants emphasized the efficiency of directly correcting generated results instead of redrawing entire storyboards. In 10 of the 12 final outputs, users reported that the unaffected portions of the video remained stable, which they considered crucial to maintaining creative momentum. In the user ratings (Fig. 12), although completing animations in this stage took longer than in the previous stages, participants did not report a corresponding increase in perceived effort. This suggests that the interface interaction is efficient even with the added steps, and the improved animation outcomes justified the additional time. In interviews, six participants noted that this locality reduced frustration: “I don’t have to start over again—just fix the part I don’t like” (P4). Compared to Stage 1, where small sketch changes could lead to unrelated alterations, Stage 3’s refinement preserved coherence and increased trust in the system.

Preferred Strategies. Different editing strategies were observed. Some participants (P3, P5, P7) intervened early, pausing the video

within the first few seconds to adjust motions before they propagated. Others (P1, P2, P8) preferred to watch a complete draft first and then make targeted corrections. These choices also influenced how edits were sequenced. For instance, as shown in Fig. 10, P8 ensured that the overall structure of the neural network animation was correct before refining details such as layer-wise flow and color changes, describing the process as “*just small fixes once it was understood.*” In contrast, P5 followed a step-by-step approach, editing the rocket’s trajectory and fall before addressing the explosion and pacing, noting a preference to go “*scene by scene.*” Strategy selection was shaped not only by personal preference but also by the nature of the task and the system’s prior knowledge. In P8’s case, the VLM correctly recognized the rough sketch as a fully connected network and applied concepts such as backpropagation, which made the creation process more automatic and reduced the need for low-level specification. As P8 remarked, it felt like the system “*already knows what I mean,*” which shaped his choice to focus on high-level adjustments rather than redrawing details.

User Perception of Control. Seven out of eight participants reported that refinement cues gave them a stronger sense of control than in previous stages. They valued being able to “*tell the AI exactly what’s wrong, instead of hoping it guesses right*” (P7). This contributed to a perceived shift from one-shot generation to an iterative, collaborative process. However, two participants also tested the system with higher-level semantic requests (e.g., “*make this look more natural*”), which were not successfully interpreted. These limitations revealed boundaries of the refinement mechanism: while effective for concrete spatial-temporal edits, it struggled with abstract stylistic or semantic adjustments. Nonetheless, participants widely agreed that refinement cues transformed the interaction into a more predictable and empowering authoring experience.

7 Generalizability and Extensibility

Beyond motion graphics, our concept can be extended to other forms of sketch-based dynamic content. Here we present a design outlook that builds on our prototype to envision alternative styles of output. While we have not implemented a complete system for these scenarios, we illustrate how the proposed pipeline, which combines vision–language model interpretation with clarification and refinement cues, can be organically integrated with other AI models to explore broader possibilities for sketch-driven creation.

7.1 Video Generation

Currently, video generation has garnered significant attention [40], with recent work such as SketchVideo [43] leveraging sketch storyboards to produce videos. However, that approach primarily relies on detailed keyframe layouts, which demand considerable drawing expertise and are often accompanied by extensive textual scripts. In contrast, our approach emphasizes minimal sketches, accessible to a broad range of users, for expressing abstract concepts, thereby lowering the entry barrier. Through the integration of these sketches with our proposed mechanisms, we enrich their interpretive depth and provide stronger support for video generation. At this stage, our demonstration is still conceptual and involves partial manual implementation: we primarily use our prototype to generate prompts for external image and video generation models, with the core

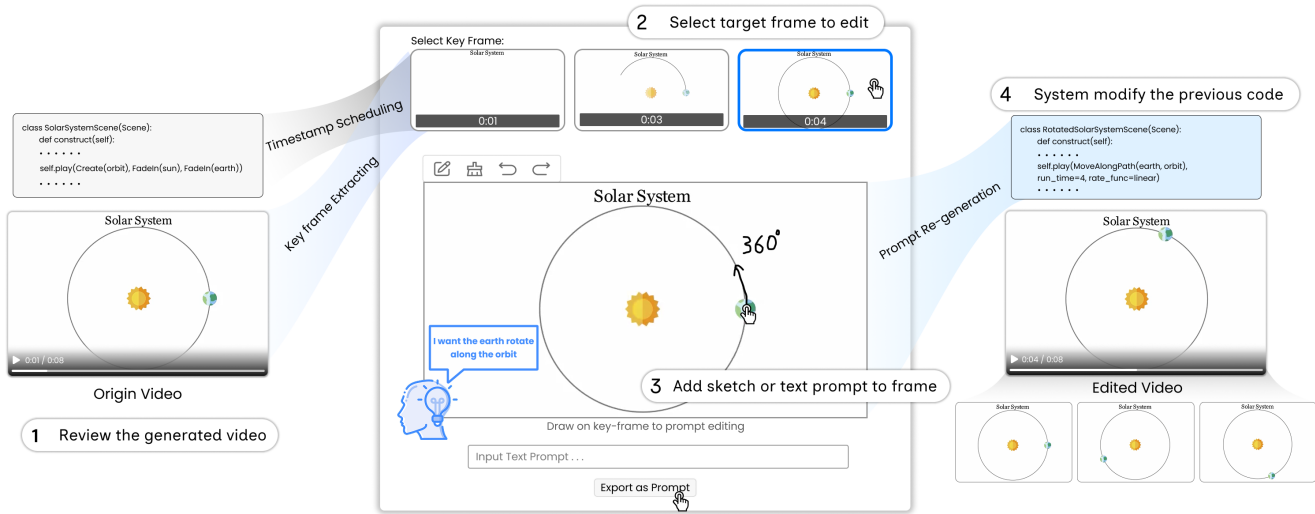


Figure 9: Example of the refinement cue workflow. In this case, the user intends the Earth to move along its orbit. The process begins with (1) reviewing the initial video. Next, (2) the system automatically extracts a set of representative keyframes, from which the user selects the one most relevant to their intended edit. (3) On the chosen keyframe, the user adds an arrow to indicate the desired motion. Finally, (4) the system updates the underlying code accordingly and regenerates the video, producing an edited version where the Earth rotates along the orbit.

contribution lying in the translation of user intent into effective prompts (see Fig. 11-A). Compared with commercial systems such as Higgsfield’s “draw-to-video” feature¹, which generates videos by annotating over existing images, our concept enables users to communicate their intent before generation through iterative, back-and-forth interaction, even when the initial input is an abstract sketch drawn from scratch.

7.2 Dynamic 3D Scene

Beyond 2D motion graphics, 3D dynamic scenes are widely applied in domains such as games and explanatory videos [39]. To illustrate this extensibility, we extended our concept to 3D animation within Unity. As shown in Fig. 11-B, the process begins with a sketch storyboard in which the user draws two cubes and a wall, using arrows to indicate motion and collision. This example demonstrates how the proposed mechanisms can assist the generation process by supplementing sketches with detailed attributes. In practice, we envision integrating our pipeline with Unity’s Model Context Protocol (MCP), enabling the VLM to interact directly with the Unity Editor by issuing function calls or generating C# scripts after interpreting the sketch. Such a workflow would allow storyboard concepts to be rapidly transformed into executable 3D scenes.

8 Discussion

We began with the question of how humans can communicate dynamic intent to AI through free-form sketches. Across three stages, we progressively shaped this process through our proposed design interventions, first by probing expressiveness, then by supporting clarification, and finally by enabling refinement. This revealed both

the potential and the limits of sketch-based interaction. In the discussion, we connect these results to broader themes and outline implications for future research.

8.1 Synergistic Cue Integration

Our integrated system reveals that clarification and refinement are not merely sequential steps, but complementary mechanisms that support the progressive formalization of intent. We observe a distinct division of labor: Clarification Cues (Stage 2) act as a semantic guardrail before generation, resolving high-level structural ambiguities (e.g., distinguishing a trajectory path from a shape outline) to ensure a correct foundation. In contrast, Refinement Cues (Stage 3) address perceptual imprecisions after generation, enabling users to tune local sensory details (e.g., timing or color) that are often difficult to articulate upfront. This “disambiguate-then-refine” synergy enables a paradigm of “lazy specification”: users can start with low-effort, abstract sketches, relying on the system to proactively align high-level intent while deferring precise adjustments to the post-generation phase. This shifts animation authoring from a high-stakes, one-shot task to a fluid, collaborative dialogue where intent is shaped iteratively.

8.2 Shaping Intent with AI

In our study, we observed that users often do not have a fully-formed animation intent at the beginning of the authoring process. Instead, their intent gradually evolves through interaction and iteration. Similar observations have been made in other domains involving generative AI [70], where the core challenge lies in supporting users in shaping their ideas fluidly rather than expecting them to articulate a complete intent upfront. While some existing

¹<https://higgsfield.ai/flow/draw>

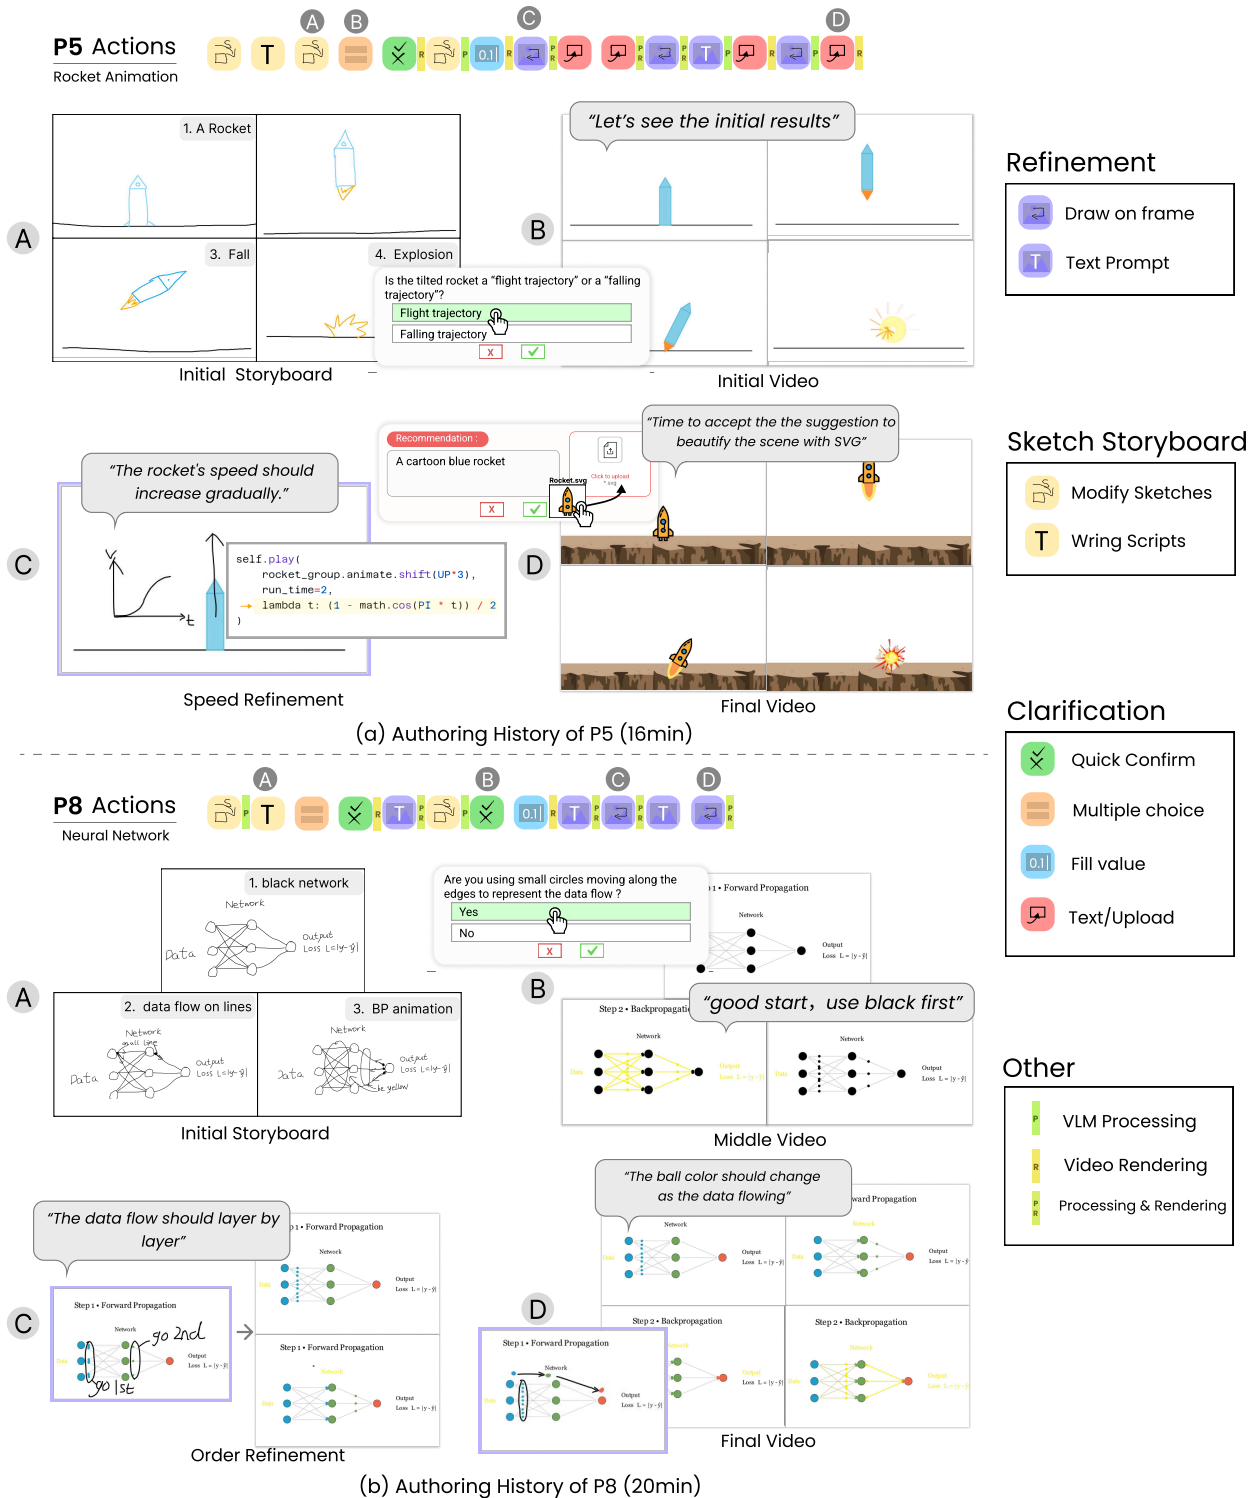


Figure 10: Authoring process of two participants. The actions are sequentially listed (only resolved clarification cues are listed), and representative steps are selected for illustration.

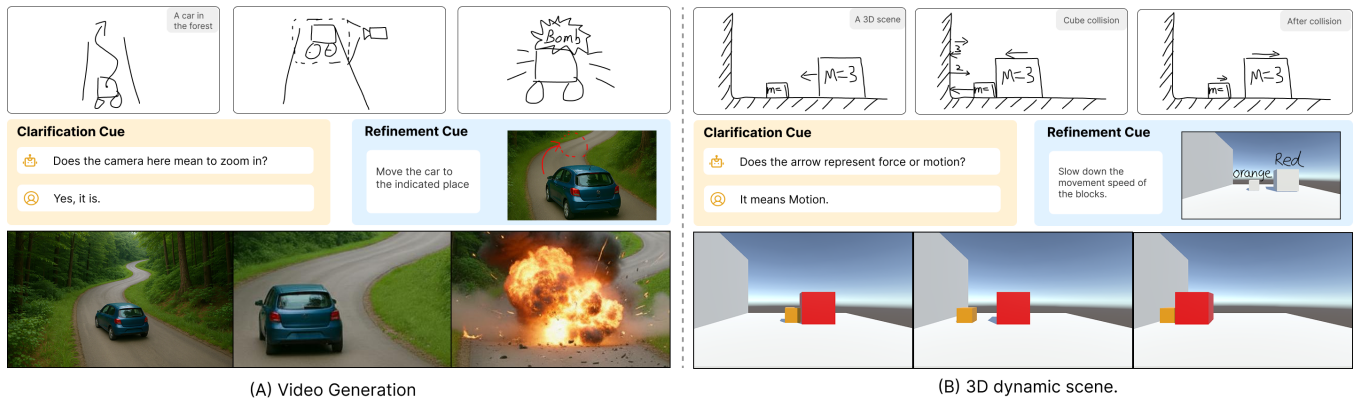


Figure 11: Demonstration of extensibility of SketchDynamics. (A) Video generation example of a car driving along a forest road and exploding. (B) 3D dynamic scene example of block collision simulation. Top: user-drawn sketch storyboard. Middle: interactive clarification and refinement cues. Bottom: corresponding output video frames.

systems attempt to address this by increasing output diversity, they often lack precise alignment with user intentions due to limited vision-language understanding capabilities. Other systems rely on complex, multi-step prompt engineering workflows to help users refine their ideas, but these often interrupt the creative flow.

We argue that sketching provides a lightweight and effective medium for expressing and gradually refining animation intent. Prior systems, such as K-Sketch [14], have highlighted that traditional animation tools are often too complex for novices due to their emphasis on frame-by-frame control. In contrast, simple sketches can allow users to communicate temporal dynamics without extensive effort, making the process significantly more accessible. In our exploration, users can either draw rough sketches from scratch or annotate on top of generated results. This supports a more fluid, back-and-forth interaction with the VLM understanding and generation. For example, during Stage 3 of our study, users often began sketching without a clear goal in mind. It was only when presented with clarification cues—such as candidate suggestions or disambiguation options—that they began to consciously shape their intent. This observation echoes findings from recent work like AnyAni [54], where novice users struggled to articulate animation needs and required iterative refinement through system feedback.

Crucially, we avoid a traditional end-to-end generation pipeline in which the AI produces a complete result in one shot, leaving little room for user intervention. Instead, we preserve the opportunity for users to continually influence and reinterpret how the AI understands their sketches. This not only allows users to refine the output but also helps them reflect on and reshape their own ideas. Through this iterative, co-creative loop, the final animation gradually emerges—balancing the strengths of both human intuition and AI generation.

8.3 Sketching Strategies

Our study reveals that participants naturally employed sketches at different levels of abstraction. On the one hand, *high-level sketches* were frequently used to represent rough shapes, motion directions, or symbolic marks. On the other hand, users occasionally turned to

low-level sketches, such as precise motion paths or data diagrams, when finer control was needed. Because our system relies primarily on the reasoning capabilities of VLMs, it was more effective at interpreting high-level sketches that conveyed semantic intent, while its ability to handle low-level details was less consistent. In practice, this meant that abstract annotations were readily understood, substantially lowering the entry barrier for novice users [13, 14], whereas capturing accurate shapes or trajectories often depended on chance alignment with the model’s visual interpretation. More generally, the two levels demand different forms of technical support. Low-level sketches require specialized algorithms to preserve geometric fidelity and temporal precision, as in ControlNet [74], whereas high-level semantics benefit more from the broad reasoning capacities of VLMs [46]. Future work should explore hybrid authoring workflows that fluidly integrate both modalities. Adaptive systems that allow users to begin with high-level, accessible annotations and selectively refine with low-level precision sketches could enable smooth transitions between rapid ideation and detailed specification, aligning authoring effort with creative goals.

In our analysis, we found participants used two main ways to sketch dynamic intent. The first was in-image motion annotations, where they drew arrows, motion curves, or ghosted outlines directly onto a single frame. This was a concise way to define a specific trajectory or transformation. The second was across-frame sequencing, the classic storyboard method where they conveyed motion by drawing successive keyframes. This approach naturally encoded temporal evolution and was better for showing complex compositional changes. This duality reflects a clear trade-off between efficiency and precision. In-image annotations were fast and preferred for simple translations or rotations. Sequencing, while more work, offered far more control over timing and complex interactions. Interestingly, each mode created unique problems for the VLM. A frame crowded with in-image annotations could become a mess of conflicting cues. Yet with sequencing, if the temporal gap between two frames was too large, the model might misinterpret a single moving element as two completely separate objects.

8.4 Sketch Inherent Uncertainty

VLMs demonstrate impressive multimodal reasoning capabilities, yet they still struggle to interpret sketches in ways that align with creator intent—especially when users [53], particularly novices, are limited to quick, abstract drawings due to time constraints [42]. This becomes a primary source of ambiguity in sketch-based inputs. In the context of dynamic content creation, this challenge is further amplified: static sketches are often underspecified for expressing animation, lacking temporal information, clear motion styles, or distinctions between annotations and scene elements. These limitations make animation sketches uniquely ambiguous.

While the model’s ability to detect such ambiguity inevitably depends on its reasoning capacity and the quality of prompt engineering, our findings suggest that ambiguity should not be viewed merely as a technical flaw, but rather as a generative opportunity. In our Stage 2 study, many ambiguities surfaced by the VLM aligned with users’ own intuitions, helping reveal aspects of intent that were not fully formed. In other cases, mismatches between model and user interpretations prompted valuable reflection and clarification. There were also instances where the model failed to detect ambiguity that seemed obvious to the user. Rather than dismissing these discrepancies as failure, we frame them as moments of negotiation—opportunities for users to refine their intent and guide the model’s interpretation. This reflects a shift from deterministic automation to an interactive, human-in-the-loop workflow. Importantly, we did not fine-tune a task-specific VLM, but instead explored how general-purpose models might already support this kind of collaborative meaning-making. Future work could extend this foundation by incorporating sketch-specific datasets or animation-aware reasoning modules to improve the alignment between model inference and user intent.

8.5 Sketch Parsing

In this study, we employed the powerful visual understanding capabilities of VLMs to parse users’ free-hand sketches. However, due to the limitations of current general-purpose VLMs, sketches had to be exported as images and analyzed at the pixel level, rather than being further parsed in the vector space. As a result, the VLM could provide an overall semantic interpretation, but was unable to fully leverage the structural and temporal information embedded in stroke data. While prior work has explored stroke-level analysis [66], existing methods are limited to basic classification and predefined symbols, which undermines support for true free-form authoring. There are also attempts that have explored feeding SVGs or stroke sequences into LLMs for interpretation, but purely textual reasoning often struggles to capture 2D spatial information [65]. Consequently, our system could not effectively explain the specific meaning of a stroke (e.g., strokes that function as static layouts, dynamic annotations). In free-hand drawing scenarios, this complexity became even more pronounced, as the stroke may play multiple roles and can not be easily classified. Our findings point toward the need for a cross-modal VLM specifically designed for sketch parsing in future work. Compared with pixel-only input, integrating vector- and stroke-level parsing could unlock richer opportunities, as stroke sequences data inherently encode order, timing, and grouping cues essential for animation authoring. For

example, if a stroke is drawn at the end, it might be an animation annotation. This also opens new possibilities for sketch-based authoring systems, where the systems can reason more naturally about the sketching process and provide responsive feedback.

8.6 Limitations

This work has several limitations. 1) *Participant diversity*: The participant pool was relatively homogeneous, consisting mainly of university students with some background in graphics or video, and thus may not fully capture perspectives from children, complete novices, or professional animators. 2) *Real-time interaction*: The system does not yet provide instant feedback at the stroke level, as sketches need to be translated into code and rendered into video, which typically takes several seconds. 3) *Model dependency*: The approach currently builds on general-purpose VLMs, whose performance is shaped by their training data and reasoning mechanisms, and future advances in these models may influence the outcomes. 4) *Task scope*: The study focused on short storyboards and motion graphics (tens of seconds), given our emphasis on sketching and the higher cost of generating longer videos. Extending the approach to longer or more complex animations is a promising avenue for future work. 5) *Expressive range*: The current implementation is constrained by the capabilities of the Manim library, which primarily supports procedural 2D vector animations and simple transformations (e.g., translation, rotation, scaling, and color changes). It does not yet accommodate more expressive animation forms such as character motion, fluid deformations, or physics-based dynamics.

9 Conclusion

This paper explores how free-form sketches can serve as an intuitive medium for expressing animation intent in the automatic creation of dynamic content. Using simple motion graphic scenarios as a starting point, we investigate the alignment between user intentions and generated results. Our overarching goal is to enable novice users to communicate animation ideas through minimal, informal inputs and to have the machine effectively interpret and realize these intentions. We improve the process of simple static sketches to convey dynamic intent across three stages: In the first stage, we support a basic interface to transform free-form sketch storyboards to motion graphics with the assistance of VLM. Observing the ambiguity and misalignment in this initial interaction, the second stage introduces clarification mechanisms that surface and resolve uncertainty in the sketch interpretation process. Building on this clarified understanding, the third stage incorporates intuitive visual post-editing tools, enabling users to close the gap between user intention and the generated outcome. Together, these stages form a continuous feedback loop that transforms VLMs from one-shot generators into interactive collaborators, capable of handling uncertainty, supporting intent refinement, and empowering users to co-construct dynamic content through sketch-based interaction.

Acknowledgments

We thank the anonymous reviewers for their helpful comments and suggestions. We also thank the participants who contributed their time to the user study. This work was partially supported

by the National Natural Science Foundation of China (Project No. 62502410).

References

- [1] 3b1b. 2025. Manim: Animation Engine for Explanatory Math Videos. <https://www.manim.community/>.
- [2] Minhazul Arefin, Kazi Mojammel Hossen, and Robin Khan. 2023. pic2code: Generating HTML Code from Handwritten Picture. In *2023 26th International Conference on Computer and Information Technology (ICCIT)*. 1–6. <https://doi.org/10.1109/ICCIT60459.2023.10441098>
- [3] Danilo Avola, Maria Chiara Caschera, Fernando Ferri, and Patrizia Grifoni. 2010. Classifying and Resolving Ambiguities in Sketch-Based Interaction. *International Journal of Virtual Technology and Multimedia* Vol.1 (03 2010), 104–139. <https://doi.org/10.1504/IJVTM.2010.032056>
- [4] Sukanya Bhattacharjee and Parag Chaudhuri. 2020. A Survey on Sketch Based Content Creation: from the Desktop to Virtual and Augmented Reality. *Computer Graphics Forum* 39, 2, 757–780. <https://doi.org/10.1111/cgf.14024> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14024>
- [5] Ron Brinkmann. 2008. *The Art and Science of Digital Compositing: Techniques for Visual Effects, Animation and Motion Graphics*. Morgan Kaufmann.
- [6] CapCut. 2025. Capcut | All-In-One Video Editor & Graphic Design Tool. <https://https://www.capcut.com/>.
- [7] Senthil Chandrasegaran, Devarajan Ramanujan, and Niklas Elmqvist. 2018. How Do Sketching and Non-Sketching Actions Convey Design Intent?. In *Proceedings of the 2018 Designing Interactive Systems Conference* (Hong Kong, China) (DIS '18). Association for Computing Machinery, New York, NY, USA, 373–385. <https://doi.org/10.1145/3196709.3196723>
- [8] Longfei Chen, Shengxin Li, Ziang Li, and Quan Li. 2025. DancingBoard: Streamlining the Creation of Motion Comics to Enhance Narratives. In *Proceedings of the 30th International Conference on Intelligent User Interfaces (IUI '25)*. Association for Computing Machinery, New York, NY, USA, 477–503. <https://doi.org/10.1145/3708359.3712167>
- [9] Shu-Yu Chen, Wanchao Su, Lin Gao, Shihong Xia, and Hongbo Fu. 2020. Deep-facereading: Deep Generation of Face Images From Sketches. *ACM Trans. Graph.* 39, 4, Article 72 (Aug. 2020), 16 pages. <https://doi.org/10.1145/3386569.3392386>
- [10] Yilan Chen, Kin Chung Kwan, and Hongbo Fu. 2023. Autocompletion of repetitive stroking with image guidance. *Computational Visual Media* 9, 3 (2023), 581–596. <https://doi.org/10.1007/s41095-022-0288-2>
- [11] Fanny Chevalier, Nathalie Henry Riche, Catherine Plaisant, Amira Chalbi, and Christophe Hurter. 2016. Animations 25 Years Later: New Roles and Opportunities. In *Proceedings of the International Working Conference on Advanced Visual Interfaces* (Bari, Italy) (AVI '16). Association for Computing Machinery, New York, NY, USA, 280–287. <https://doi.org/10.1145/2909132.2909255>
- [12] James Davis, Maneesh Agrawala, Erika Chuang, Zoran Popović, and David Salesin. 2006. A Sketching Interface for Articulated Figure Animation. In *ACM SIGGRAPH 2006 Courses* (Boston, Massachusetts) (SIGGRAPH '06). Association for Computing Machinery, New York, NY, USA, 15–es. <https://doi.org/10.1145/1185657.1185776>
- [13] Richard Davis and James Landay. 2004. Informal Animation Sketching: Requirements and design. *AAAI Fall Symposium - Technical Report* (01 2004).
- [14] Richard C. Davis, Brien Colwell, and James A. Landay. 2008. K-Sketch: A 'Kinetic' Sketch Pad for Novice Animators. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) (CHI '08). Association for Computing Machinery, New York, NY, USA, 413–422. <https://doi.org/10.1145/1357054.1357122>
- [15] Adobe After Effect. 2025. Adobe After Effects - Motion graphics software. <https://www.adobe.com/products/aftereffects.html>.
- [16] Maxime Garcia, Remi Ronfard, and Marie-Paule Cani. 2019. Spatial Motion Doodles: Sketching Animation in VR Using Hand Gestures and Laban Motion Analysis (MIG '19). Association for Computing Machinery, New York, NY, USA, Article 10, 10 pages. <https://doi.org/10.1145/3359566.3360061>
- [17] William W. Gaver, Jacob Beaver, and Steve Benford. 2003. Ambiguity as A Resource for Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, Florida, USA) (CHI '03). Association for Computing Machinery, New York, NY, USA, 233–240. <https://doi.org/10.1145/642611.642653>
- [18] Martin Guay. 2015. Sketching Free-Form Poses and Motions for Expressive 3D Character Animation. *HAL 2015* (2015). <http://dml.mathdoc.fr/item/NNT:2015GREAM016>
- [19] Tracy Hammond, Brian Eoff, Brandon Paulson, Aaron Wolin, Katie Dahmen, Joshua Johnston, and Pankaj Rajan. 2008. Free-Sketch Recognition: Putting the CHI in Sketching. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems* (Florence, Italy) (CHI EA '08). Association for Computing Machinery, New York, NY, USA, 3027–3032. <https://doi.org/10.1145/1358628.1358802>
- [20] John Hart. 2013. *The Art of the Storyboard: A filmmaker's introduction*. Routledge.
- [21] Brian Hempel, Justin Lubin, and Ravi Chugh. 2019. Sketch-n-Sketch: Output-Directed Programming for SVG. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 281–292. <https://doi.org/10.1145/3332165.3347925>
- [22] Forrest Huang, John F. Canny, and Jeffrey Nichols. 2019. Swire: Sketch-based User Interface Retrieval. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/3290605.3300334>
- [23] Takeo Igarashi, Satoshi Matsuoka, Sachiko Kawachiya, and Hidehiko Tanaka. 1997. Interactive Beautification: A Technique for Rapid Geometric Design. In *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology* (Banff, Alberta, Canada) (UIST '97). Association for Computing Machinery, New York, NY, USA, 105–114. <https://doi.org/10.1145/263407.263525>
- [24] Amir Jahanlou and Parmit K Chilana. 2022. Katika: An End-to-End System for Authoring Amateur Explainer Motion Graphics Videos. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 502, 14 pages. <https://doi.org/10.1145/3491102.3517741>
- [25] Hao Jin, Hengyuan Chang, Xiaoxuan Xie, Zhengyang Wang, Xusheng Du, Shaojun Hu, and Haoran Xie. 2024. Sketch-Guided Motion Diffusion for Stylized Cinemagraph Synthesis. arXiv:2412.00638 [cs.CV] <https://arxiv.org/abs/2412.00638>
- [26] Bo Kang, Joseph J. LaViola Jr., and Pamela Wisniewski. 2017. Structured Input Improves Usability and Precision for Solving Geometry-based Algebraic Problems. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 4692–4702. <https://doi.org/10.1145/3025453.3025468>
- [27] Jun Kato, Kenta Hara, and Nao Hirasawa. 2024. Griffith: A Storyboarding Tool Designed with Japanese Animation Professionals. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 233, 14 pages. <https://doi.org/10.1145/3613904.3642121>
- [28] Jun Kato, Tomoyasu Nakano, and Masataka Goto. 2015. TextAlive: Integrated Design Environment for Kinetic Typography. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 3403–3412. <https://doi.org/10.1145/2702123.2702140>
- [29] Rubaiat Habib Kazi, Fanny Chevalier, Tovi Grossman, and George Fitzmaurice. 2014. Kitty: Sketching Dynamic and Interactive Illustrations. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (Honolulu, Hawaii, USA) (UIST '14). Association for Computing Machinery, New York, NY, USA, 395–405. <https://doi.org/10.1145/2642918.2647375>
- [30] Rubaiat Habib Kazi, Fanny Chevalier, Tovi Grossman, Shengdong Zhao, and George Fitzmaurice. 2014. DRACO: Sketching Animated Drawings with Kinetic Textures. In *ACM SIGGRAPH 2014 Studio* (Vancouver, Canada) (SIGGRAPH '14). Association for Computing Machinery, New York, NY, USA, Article 34, 1 pages. <https://doi.org/10.1145/2619195.2656322>
- [31] Rubaiat Habib Kazi, Tovi Grossman, Nobuyuki Umetani, and George Fitzmaurice. 2016. Motion Amplifiers: Sketching Dynamic Illustrations Using the Principles of 2D Animation. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 4599–4609. <https://doi.org/10.1145/2858036.2858386>
- [32] KLING. 2025. KLING AI: Next-Generation AI Creative Studio. <https://klingai.com/>.
- [33] Jon Krasner. 2013. *Motion Graphic Design: Applied History and Aesthetics*. Routledge.
- [34] James A. Landay and Brad A. Myers. 1995. Interactive Sketching for the Early Stages of User Interface Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '95). ACM Press/Addison-Wesley Publishing Co., USA, 43–50. <https://doi.org/10.1145/223904.223910>
- [35] James A. Landay and Brad A. Myers. 1996. Sketching Storyboards to Illustrate Interface Behaviors. In *Conference Companion on Human Factors in Computing Systems* (Vancouver, British Columbia, Canada) (CHI '96). Association for Computing Machinery, New York, NY, USA, 193–194. <https://doi.org/10.1145/257089.257257>
- [36] Bongshin Lee, Rubaiat Habib Kazi, and Greg Smith. 2013. SketchStory: Telling More Engaging Stories with Data through Freeform Sketching. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (Dec. 2013), 2416–2425. <https://doi.org/10.1109/TVCG.2013.191>
- [37] Johnny C. Lee, Jodi Forlizzi, and Scott E. Hudson. 2002. The Kinetic Typography Engine: An Extensible System for Animating Expressive Text. In *Proceedings of the 15th Annual ACM Symposium on User Interface Software and Technology* (Paris, France) (UIST '02). Association for Computing Machinery, New York, NY, USA, 81–90. <https://doi.org/10.1145/571985.571997>
- [38] Boyu Li, Linjie Qiu, Duotun Wang, Qianxi Liu, Ryo Suzuki, Mingming Fan, and Zeyu Wang. 2025. DesignMemo: Integrating Discussion Context into Online

- Collaboration with Enhanced Design Rationale Tracking. *Proc. ACM Hum.-Comput. Interact.* 9, 7, Article CSCW398 (Oct. 2025), 32 pages. <https://doi.org/10.1145/3757579>
- [39] Boyu Li, Linping Yuan, Zhe Yan, Qianxi Liu, Yulin Shen, and Zeyu Wang. 2024. AniCraft: Crafting Everyday Objects as Physical Proxies for Prototyping 3D Character Animation in Mixed Reality. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) (UIST '24). Association for Computing Machinery, New York, NY, USA, Article 99, 14 pages. <https://doi.org/10.1145/3654777.3676325>
- [40] Boyu Li, Lin-Ping Yuan, and Zeyu Wang. 2025. VideoCraft: A Mixed Reality-empowered Video Generation Workflow with Spatial Layer Editing for Concept Video Creation (UIST '25). Association for Computing Machinery, New York, NY, USA, Article 19, 16 pages. <https://doi.org/10.1145/3746059.3747606>
- [41] Haichuan Lin, Yilin Ye, Jiazhi Xia, and Wei Zeng. 2025. SketchFlex: Facilitating Spatial-Semantic Coherence in Text-to-Image Generation with Region-Based Sketches. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 546, 19 pages. <https://doi.org/10.1145/3706598.3713801>
- [42] Yuyu Lin, Jiahao Guo, Yang Chen, Cheng Yao, and Fangtian Ying. 2020. It Is Your Turn: Collaborative Ideation With a Co-Creative Robot through Sketch. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376258>
- [43] Feng-Lin Liu, Hongbo Fu, Xintao Wang, Weicai Ye, Pengfei Wan, Di Zhang, and Lin Gao. 2025. SketchVideo: Sketch-based Video Generation and Editing. In *Proceedings of the Computer Vision and Pattern Recognition Conference*. 23379–23390.
- [44] Feng-Lin Liu, Shi-Yang Li, Yan-Pei Cao, Hongbo Fu, and Lin Gao. 2025. Sketch3DVE: Sketch-based 3D-Aware Scene Video Editing. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers* (SIGGRAPH Conference Papers '25). Association for Computing Machinery, New York, NY, USA, Article 152, 12 pages. <https://doi.org/10.1145/3721238.3730623>
- [45] Shaoteng Liu, Tianyu Wang, Jui-Hsien Wang, Qing Liu, Zhifei Zhang, Joon-Young Lee, Yijun Li, Bei Yu, Zhe Lin, Soo Ye Kim, and Jiaya Jia. 2024. Generative Video Propagation. arXiv:2412.19761 [cs.CV] <https://arxiv.org/abs/2412.19761>
- [46] Vivian Liu, Rubaiat Habib Kazi, Li-Yi Wei, Matthew Fisher, Timothy Langlois, Seth Walker, and Lydia Chilton. 2025. LogoMotion: Visually-Grounded Code Synthesis for Creating and Editing Animation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 157, 16 pages. <https://doi.org/10.1145/3706598.3714155>
- [47] Zhicheng Liu, John Thompson, Alan Wilson, Mira Dontcheva, James Delorey, Sam Grigg, Bernard Kerr, and John Stasko. 2018. Data Illustrator: Augmenting Vector Design Tools with Lazy Data Binding for Expressive Visualization Authoring. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173697>
- [48] Zhaoliang Lun, Matheus Gadelha, Evangelos Kalogerakis, Subhansu Maji, and Rui Wang. 2017. 3D Shape Reconstruction from Sketches via Multi-view Convolutional Networks. In *2017 International Conference on 3D Vision (3DV)*. 67–77. <https://doi.org/10.1109/3DV.2017.00018>
- [49] Zhongjin Luo, Dong Du, Heming Zhu, Yizhou Yu, Hongbo Fu, and Xiaoguang Han. 2024. SketchMetaFace: A Learning-Based Sketching Interface for High-Fidelity 3D Character Face Modeling. *IEEE Transactions on Visualization & Computer Graphics* 30, 08 (Aug. 2024), 5260–5275. <https://doi.org/10.1109/TVCG.2023.3291703>
- [50] Jiaju Ma and Maneesh Agrawala. 2025. MoVer: Motion Verification for Motion Graphics Animations. *ACM Trans. Graph.* 44, 4, Article 33 (July 2025), 17 pages. <https://doi.org/10.1145/3731209>
- [51] Krystina Madej and Newton Lee. 2020. *Disney Stories*. Springer.
- [52] Timothy Matthews. 2012. *Sketch-based Digital Storyboards and Floor Plans for Authoring Computer-Generated Film Pre-visualisations*. Ph.D. Dissertation. Nelson Mandela Metropolitan University.
- [53] Kushin Mukherjee, Holly Huey, Xuanchen Lu, Yael Vinker, Rio Aguina-Kang, Ariel Shamir, and Judith E. Fan. 2023. SEVA: Leveraging Sketches to Evaluate Alignment Between Human and Machine Visual Abstraction. In *Proceedings of the 37th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (NIPS '23). Curran Associates Inc., Red Hook, NY, USA, Article 2934, 18 pages.
- [54] Tianrun Qiu and Yuxin Ma. 2025. AnyAni: An Interactive System with Generative AI for Animation Effect Creation and Code Understanding in Web Development. arXiv:2506.21962 [cs.HC] <https://arxiv.org/abs/2506.21962>
- [55] Dean Rubine. 1991. Specifying Gestures by Example. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques* (SIGGRAPH '91). Association for Computing Machinery, New York, NY, USA, 329–337. <https://doi.org/10.1145/122718.122753>
- [56] Sigurdur Gauti Samuelsson and Matthias Book. 2020. Eliciting Sketched Expressions of Command Intentions in an IDE. *Proc. ACM Hum.-Comput. Interact.* 4, ISS, Article 200 (Nov. 2020), 25 pages. <https://doi.org/10.1145/3427328>
- [57] Xinyu Shi, Yinghou Wang, Yun Wang, and Jian Zhao. 2024. Piet: Facilitating Color Authoring for Motion Graphics Video. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 148, 17 pages. <https://doi.org/10.1145/3613904.3642711>
- [58] Sarah Sonje, Harsh Dave, Jaswantsingh Pardeshi, and Sheetal Chaudhari. 2022. draw2code: AI based Auto Web Page Generation from Hand-drawn Page Mock-up. In *2022 IEEE 7th International conference for Convergence in Technology (I2CT)*. 1–7. <https://doi.org/10.1109/I2CT54291.2022.9824521>
- [59] Masaki Suwa and Barbara Tversky. 1997. What Do Architects and Students Perceive in Their Design Sketches? A Protocol Analysis. *Design Studies* 18, 4 (1997), 385–403. [https://doi.org/10.1016/S0142-694X\(97\)00008-2](https://doi.org/10.1016/S0142-694X(97)00008-2) Descriptive models of design.
- [60] John R Thompson, Zhicheng Liu, and John Stasko. 2021. Data Animator: Authoring Expressive Animated Data Graphics. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 15, 18 pages. <https://doi.org/10.1145/3411764.3445747>
- [61] Bekzat Tilekbay, Saelyne Yang, Michal Adam Lewkowicz, Alex Suryapranata, and Juho Kim. 2024. ExpressEdit: Video Editing with Natural Language and Sketching. In *Proceedings of the 29th International Conference on Intelligent User Interfaces* (Greenville, SC, USA) (IUI '24). Association for Computing Machinery, New York, NY, USA, 515–536. <https://doi.org/10.1145/3640543.3645164>
- [62] Tiffany Tseng, Ruijia Cheng, and Jeffrey Nichols. 2024. Keyframer: Empowering Animation Design using Large Language Models. <https://arxiv.org/abs/2402.06071>
- [63] Winger S.W. Tseng and Linden J. Ball. 2011. How Uncertainty Helps Sketch Interpretation in a Design Task. In *Design Creativity 2010*, Toshiharu Taura and Yukari Nagai (Eds.). Springer London, London, 257–264.
- [64] Barbara Tversky and Masaki Suwa. 2009. *Thinking with Sketches*. <https://doi.org/10.1093/acprof:oso/9780195381634.003.0004>
- [65] Yael Vinker, Tamar Rott Shaham, Kristine Zheng, Alex Zhao, Judith E Fan, and Antonio Torralba. 2025. SketchAgent: Language-Driven Sequential Sketch Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 23355–23368.
- [66] Jacob O. Wobbrock, Andrew D. Wilson, and Yang Li. 2007. Gestures Without Libraries, Toolkits or Training: A \$1 Recognizer for User Interface Prototypes. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology* (Newport, Rhode Island, USA) (UIST '07). Association for Computing Machinery, New York, NY, USA, 159–168. <https://doi.org/10.1145/1294211.1294238>
- [67] Chufeng Xiao, Deng Yu, Xiaoguang Han, Youyi Zheng, and Hongbo Fu. 2021. Sketchhairsalon: Deep Sketch-Based Hair Image Synthesis. *ACM Trans. Graph.* 40, 6, Article 216 (Dec. 2021), 16 pages. <https://doi.org/10.1145/3478513.3480502>
- [68] Shiyuan Yang, Liang Hou, Haibin Huang, Chongyang Ma, Pengfei Wan, Di Zhang, Xiaodong Chen, and Jing Liao. 2024. Direct-a-Video: Customized Video Generation with User-Directed Camera Movement and Object Motion. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) (SIGGRAPH '24). Association for Computing Machinery, New York, NY, USA, Article 113, 12 pages. <https://doi.org/10.1145/3641519.3657481>
- [69] Ryan Yen, Jian Zhao, and Daniel Vogel. 2025. Code Shaping: Iterative Code Editing with Free-form AI-Interpreted Sketching. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 872, 17 pages. <https://doi.org/10.1145/3706598.3713822>
- [70] Seo young Lee, Matthew Law, and Guy Hoffman. 2025. When and How to Use AI in the Design Process? Implications for Human-AI Design Collaboration. *International Journal of Human-Computer Interaction* 41, 2 (2025), 1569–1584. <https://doi.org/10.1080/10447318.2024.2353451> arXiv:https://doi.org/10.1080/10447318.2024.2353451
- [71] Lin-Ping Yuan, Boyu Li, Jindong Wang, Huamin Qu, and Wei Zeng. 2024. Generating Virtual Reality Stroke Gesture Data From Out-Of-Distribution Desktop Stroke Gesture Data. In *IEEE Conference Virtual Reality and 3D User Interfaces*. IEEE, 732–742.
- [72] Nurcan Yildizoğlu. 2024. Sketching Versus Digital Design Tools in Architectural Design. *Journal of Computational Design* 5 (08 2024). <https://doi.org/10.53710/jcode.1504947>
- [73] Robert C. Zeleznik, Andrew Bragdon, Chu-Chi Liu, and Andrew Forsberg. 2008. Lineogrammer: Creating Diagrams by Drawing. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology* (Monterey, CA, USA) (UIST '08). Association for Computing Machinery, New York, NY, USA, 161–170. <https://doi.org/10.1145/1449715.1449741>
- [74] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*. 3813–3824. <https://doi.org/10.1109/ICCV51070.2023.00355>

- [75] Lei Zhong, Chuan Guo, Yiming Xie, Jiawei Wang, and Changjian Li. 2025. Sketch2Anim: Towards Transferring Sketch Storyboards into 3D Animation. *arXiv preprint arXiv:2504.19189* (2025).

A User Rating in Three Stages

Fig. 12 illustrates the cumulative benefit of our design interventions across the three stages. We observe a substantial upward trend in Alignment and Control, shifting from predominantly negative ratings in Stage 1 to highly positive ones in Stage 3. This validates that while the VLM alone (Stage 1) often struggles with interpretation, the integration of clarification and refinement effectively bridges the gulf between user intent and system output. Notably, despite the additional interaction steps introduced in the final system, Effort ratings improved compared to Stage 1. This suggests that users perceived the targeted refinement process as significantly less laborious and frustrating than the repetitive trial-and-error redrawing required in the baseline. Furthermore, Flow and Ease remained consistently high, indicating that the added mechanisms were seamlessly integrated without disrupting the creative momentum.

B Visual Understanding Capability Experiment

Before the main study, we conducted pilot tests to check feasibility, focusing on prompt design and the VLM's ability to interpret free-form sketches. The results were encouraging: the prompts guided the model effectively, and the VLM showed strong potential in handling abstract input.

Here we show how the VLM with prompt engineering interpreted several Stage 1 sketches, illustrating how rough or symbolic drawings were mapped into animation concepts. This is not part of our core contribution but serves as an important foundation for the paper.

C User Study Survey and Interview Questions

This appendix provides the detailed instruments used in our three-stage user study to evaluate *SketchDynamics*. These include the participant background survey, the post-task Likert scale questionnaire, and the semi-structured interview guide.

C.1 Pre-study Questionnaire: Background and Experience

Before the experiment, participants provided background information to help us understand their prior expertise in animation and sketching:

- **Demographics:** Age, Gender, and Handedness.
- **Motion Graphics Experience:**
 - How often do you watch motion graphics or explainer videos? (1: Never, 7: Daily).
 - Have you ever created motion graphics using tools such as After Effects, CapCut, or others?
 - Rate your proficiency in video editing and animation software (1: Novice, 7: Expert).
- **Sketching Experience:**
 - How often do you use sketching for ideation or communication? (1: Never, 7: Daily).
 - Rate your comfort level with drawing/sketching (1: Very Uncomfortable, 7: Very Comfortable).

C.2 Post-task Questionnaire

After completing the tasks in each stage, participants rated their experience based on the following statements using a 7-point Likert scale (1: Strongly Disagree, 7: Strongly Agree)

- **Alignment:** The generated results reflected what I had in mind.
- **Ease:** The interface was easy to use and understand.
- **Flow:** The workflow felt smooth without disrupting my creative process.
- **Control:** I felt in control of how my input was interpreted.
- **Effort:** The effort required to use the system was reasonable.
- **Explore:** The system encouraged me to try new ways of expressing my ideas.

C.3 Semi-structured Interview Guide

The following questions guided our qualitative analysis across the three stages of the study:

C.3.1 General Questions (All Stages).

- What was your primary strategy for translating animation ideas into free-form sketches?
- Were there specific animation intents that you found particularly difficult to express through sketching?
- Did the system's interpretation of your sketches surprise you? In what way?

C.3.2 Stage 2: Clarification Cues.

- How did you perceive the clarification requests (e.g., multiple-choice, fill-value)? Did they feel like interruptions or helpful checks?
- Did the clarification process help you better define animation details (like timing or speed) that you hadn't initially considered?

C.3.3 Stage 3: Refinement Cues.

- How does the experience of refining a generated video compare to redrawing a storyboard from scratch?
- Did the ability to edit specific video frames increase your sense of control over the AI collaborator?
- Which did you find more useful for refinement: drawing on the frame or providing text prompts? Why?




Figure 12: Self-defined Likert scale results across the three stages. Alignment: The generated results reflected what I had in mind. Ease: The interface was easy to use and understand. Flow: The workflow felt smooth without disrupting my creative process. Control: I felt in control of how my input was interpreted. Effort: The effort required to use the system was reasonable. Explore: The system encouraged me to try new ways of expressing my ideas.

Table 1: Examples of sketch types and the animation intentions they convey

Sketch Type	Sketch Example	Text Script	Animation Description by VLM
Camera Motion		Virtual Camera Movement	Plan for a virtual camera sweep across six regions in a Z-shaped path: Play → Garden → Lake (top row), then down to Sports → Lawn → Cafe (bottom row), followed by an exit movement to the left.
Fade Out		Null	The word 'CHI' appears, then gradually fades while moving downward (represented with dashed lines).
Shape Changing		Null	A square morphs into a triangle and then into a line, shown sequentially within the canvas.
Moving Direction		Null	A pendulum suspended from a pivot with an arrow pointing left at the bob, indicating oscillatory motion.
Ghost State		Blue indicates future	The central circle enlarges to the blue dashed size while the five surrounding circles shrink to their dashed sizes.