

# VideoDiff: Human-AI Video Co-Creation with Alternatives

Mina Huh\*  
The University of Texas at Austin  
minahuh@cs.utexas.edu

Dingzeyu Li  
Adobe Research  
dinli@adobe.com

Kim Pimmel  
Adobe Research  
kipimmel@adobe.com

Hijung Valentina Shin  
Adobe Research  
vshin@adobe.com

Amy Pavel  
The University of Texas at Austin  
apavel@cs.utexas.edu

Mira Dontcheva  
Adobe Research  
mirad@adobe.com

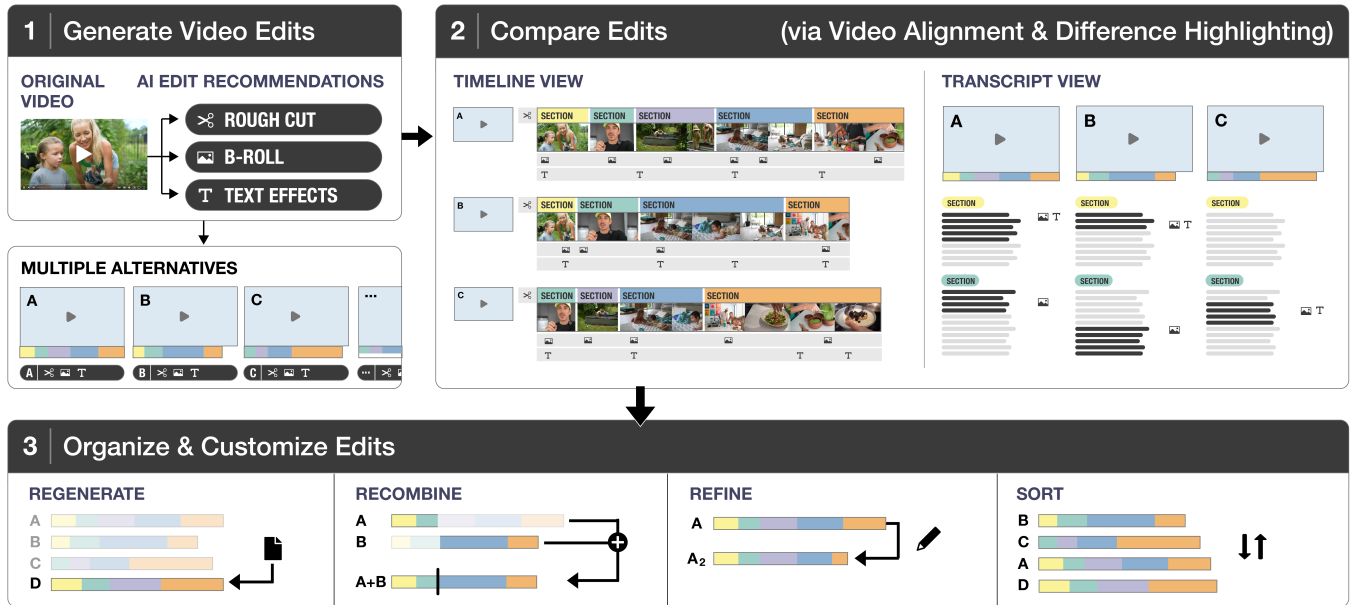


Figure 1: VideoDiff is a Human-AI co-creative system that supports video creators to explore multiple variations. 1) VideoDiff generates diverse AI recommendations for making rough cuts, inserting B-rolls, and adding text effects. 2) VideoDiff supports easy comparison by aligning videos and highlighting differences using timeline and transcript views. 3) Users can narrow down options by sorting and further customize them by refining and regenerating AI suggestions.

## Abstract

To make an engaging video, people sequence interesting moments and add visuals such as B-rolls or text. While video editing requires time and effort, AI has recently shown strong potential to make editing easier through suggestions and automation. A key strength of generative models is their ability to quickly generate multiple variations, but when provided with many alternatives, creators struggle to compare them to find the best fit. We propose VideoDiff, an AI video editing tool designed for editing with alternatives. With VideoDiff, creators can generate and review multiple AI recommendations for each editing process: creating a rough cut, inserting B-rolls, and adding text effects. VideoDiff simplifies comparisons

by aligning videos and highlighting differences through timelines, transcripts, and video previews. Creators have the flexibility to regenerate and refine AI suggestions as they compare alternatives. Our study participants (N=12) could easily compare and customize alternatives, creating more satisfying results.

## CCS Concepts

• Human-centered computing → Interactive systems and tools.

## Keywords

Video Editing, Authoring Tools, Generative AI, Human-AI Co-Creation

\*This work was done during the author's Adobe research internship

## ACM Reference Format:

Mina Huh, Dingzeyu Li, Kim Pimmel, Hijung Valentina Shin, Amy Pavel, and Mira Dontcheva. 2025. VideoDiff: Human-AI Video Co-Creation with Alternatives. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26-May 1, 2025, Yokohama, Japan. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3706598.3713417>

## 1 Introduction

As video has become a mainstream form of communication and storytelling, an increasing number of end users create and share videos. YouTube, which is considered the most popular video-focused social network, now has approximately 64 million creators worldwide [26]. Yet creating compelling videos is a complex and time-consuming task. Creators need to find key moments [79] and remove irrelevant and repetitive content [29, 37]. They also spend time making the video more engaging with visual effects, B-roll [35], text [83], and music [61]. Recent advances in video understanding and generative models have shown great potential for assisting video editing. Prior research has shown that AI tools can speed up multiple stages of video authoring including script writing [53], storyboarding [78], assembling clips into rough cuts [20, 69], identifying low-quality footage [29, 37], and adding B-roll [35]. Recent AI video products such as OpusClip [22], CapCut [13], and Vizard [72] further streamline video editing by automatically making cuts, and adding transition effects and subtitles.

One powerful new capability enabled by generative AI models is the quick generation of multiple variations. This allows creators to explore many alternative stories or B-roll placements simultaneously, potentially resulting in better final videos [24, 62]. While most existing video editing tools are designed to handle only one video version at a time, recent AI tools such as OpusClip [22] and Capcut [13] generate multiple variations of edited videos to offer diverse options to users. Despite the benefits of exploring alternatives in creative tasks, there are potential new burdens to users: 1) comparing the variations [31, 36] and 2) managing them over time [58, 62]. While prior work explored sensemaking and comparison of multiple AI generations in text [31, 58, 63], images [3, 36], and designs [50, 64], comparing multiple videos presents unique challenges due to the temporal nature of video. In this work, we explore this emerging approach to video editing centered on working with multiple variations.

To understand the opportunities and challenges of authoring videos with multiple alternatives, we conducted a formative study where 8 professional video creators were tasked with comparing multiple edited videos of the same source content. Creators in our study mentioned that comparison is a common practice throughout their current video editing process as they consider alternative narratives, visual assets, or lengths of videos. They also highlighted that having alternatives helps them reflect on their preferences and further plan editing directions, but it was time-consuming to manually create multiple versions. As AI speeds up the video creation process, we envision that future video editing tools will more commonly provide users with multiple variations.

However, reviewing alternatives can require a lot of time and effort because creators must watch all the videos to compare them. Unlike static images or text, which can be quickly scanned for differences, videos unfold over time, requiring creators to watch, pause, and replay to detect variations in stories, transitions, and visual effects. Based on these findings, we derive 6 design goals for video editing tools centered on working with alternatives. A tool to support creators in exploring video alternatives should minimize redundant watching (D1) and enable quick skimming of differences (D2) across different stages of the editing process (D3) in the most

appropriate modality (D4). Such a tool should also help creators easily organize and customize variations to support them in finding a version with fewer errors (D5) and best suits their needs (D6).

We introduce *VideoDiff*, an AI-powered video editing tool designed to enhance the creative potential of working with alternatives while minimizing the burden of comparison. VideoDiff generates diverse AI recommendations for each video editing task such as making rough cuts, inserting B-rolls, and adding text effects (Figure 1.1). For each task, creators can quickly skim the differences between variations using VideoDiff which aligns the variations and highlights the differences using timelines, transcripts, and video previews (Figure 1.2). Additionally, creators can sort variations to narrow down the options and further customize them by refining and regenerating AI suggestions (Figure 1.3). In this work, we focus specifically on working with narrative videos such as lifestyle vlogs, podcasts, how-to videos, and travel videos. Their descriptive dialogue makes them well suited to the current strengths of text-based large language models, which are good at extractive and abstractive summarization. Additionally, the source content for these types of videos is either one long video (e.g., presentation or an event) or multiple sources that can be sequenced in time because they capture chronological activities, such as a how-to demonstration or travel videos. This simplifies our visualization design to show how different versions relate to one another.

We evaluated VideoDiff in a within-subjects study with 12 video creators who compared VideoDiff with a baseline interface designed to encompass features of existing AI video editing tools (e.g., CapCut, OpusClip). Participants rated VideoDiff as more useful for quickly understanding the differences of multiple videos which helped them create and consider more diverse variations. Participants also expressed higher satisfaction with the final videos created with VideoDiff. The contributions of our paper are as follows:

- We identify the challenges of comparing multiple variations in video editing and derive design goals for systems to support video creation with alternatives.
- We propose *VideoDiff*, an AI-powered video editing tool that supports video creation with alternatives by enabling efficient generation, comparison, organization, and customization of variations.
- Our user study findings showcase the advantages and challenges of using VideoDiff.

## 2 Related work

### 2.1 Human-AI video co-creation

Traditional video editing tools such as Premiere Pro [2] or iMovie [4] require users to manually inspect the video, plan for what edits to make, and execute them by interacting with the complex UI in video editing tools. With the power of AI to understand the content and quality of videos, new forms of video editing systems have been introduced that speed up and automate the video creation process [13, 22, 23, 38]. Earlier systems supported users to structure narrated videos into multiple segments to help them quickly skim and select appropriate parts [37, 44, 68, 69, 77]. These systems label clips with metadata such as speaker and topic to let users navigate and search [37, 68], or recommend video segments that are semantically relevant to the script [45, 69, 77, 82]. Beyond organization of

clips, researchers have also proposed tools that recommend edits based on the video content and quality. B-Script [35] analyzes the video transcript and recommends possible B-roll positions and clips. AVscript [37] detects low-quality footage (e.g., blur, bad lighting) to cut out, and commercial tools such as Describe [23] and Vimeo [71] identify filler words and pauses for quick removal. Beyond post-editing, prior research has also proposed automatic video creation from various media including text documents [17], webpage [19], markdown [18], news articles [78], and music [48].

More recent tools proposed LLM-powered video editing tools that are capable of understanding users' editing intent in natural language and automatically executing them. ExpressEdit [67] allows users to edit videos by typing edit instructions and sketching on the video frame. LAVE [76] provides an LLM-agent that assists users in planning and performing edit actions while also enabling users to manually refine agent actions. To support the task of repurposing long videos into short and compelling videos, PodReels [79] and ROPE [75] explore the use of LLMs to create podcast summaries.

Existing Human-AI video co-creation tools already create multiple variations for users to better explore the design space. For instance, ReelFramer [78] suggests multiple narrative framing to let users explore and choose video topics, and ROPE [75] generates summaries of variable lengths. However, the new task of how users review, refine, and select among these variations of videos remains unexplored. In this era of AI-driven video creation, we investigate user challenges in reviewing and comparing AI-edited videos.

## 2.2 Video summarization, skimming, and browsing

Understanding and browsing video content is an important first step in video editing. Traditional video players commonly use a *seek bar* which indicates the progress of the video and allows viewers to navigate to different parts of the video. To address the difficulty of previewing upcoming scenes, many video players on the web including YouTube provide picture-in-picture previews of thumbnails when viewers hover over the seek bar and video chapters that segment video into multiple sections and add context to each portion for quick skimming.

There is a large body of work that summarizes a video as images by selecting the important keyframes [8, 10, 14, 39, 69]. Video editing tools such as Premiere Pro and iMovie utilize *filmstrip*, a visual representation of a video's frames displayed in a linear sequence. This format allows editors to get a quick overview of the entire video at a glance and identify parts for cutting or inserting effects. To enable seamless zoom over filmstrip, Video Tapestries eliminated borders between frames for spatial continuity [8]. Other works used grid-view of keyframes to summarize the videos [10, 39]. Boreczkly et al. proposed a comic book presentation for videos that displayed keyframes with different sizes based on their importance [10] and ElasticPlay visualizes an interactive summary based on dynamic time budget using a grid of thumbnails [39].

Researchers have also explored video browsing interfaces that leverage video transcripts. To support the navigation of informational videos, these systems highlight keywords in the transcript [14], or provide scaffolding steps [28, 84] and concepts [46]. Building upon prior research on video summarization and browsing

techniques, our work explores visual representations to support efficient review of multiple edited videos.

## 2.3 Comparison and sensemaking

Comparison is crucial in creative tasks, helping users understand the design space and refine their decisions. Variation theory [49] suggests that learning occurs through experiencing contrasts. Thus, exposing creators to alternatives can help them understand model affordances and make more informed decisions. [30]

To visualize comparison, 3 main techniques are used: juxtaposition, superposition, and explicit encoding [32]. *Juxtaposition* places distinct visual elements side-by-side to highlight similarities and differences. *Superposition* overlays visuals within a common frame and *Explicit encoding* reveals the predefined relationship between visualizations. Prior work utilized these techniques for video comparison by building sequential and parallel video players [6, 66], and highlighting the subtle differences with color overlay [6, 7]. To support the search of educational videos, VSedu supports video comparison based on topic relevance [9]. Video Lens facilitated the search in a large collection of sports videos via rapid playback and event-based search [51]. This line of work extends information visualization techniques to support the comparison of videos from multiple sources. Yet, they do not support video comparison in the editing context, where the comparison involves multiple variations of edited videos from the same source footage.

As AI continues to speed up and automate content generation, it is becoming more common to work with multiple variations in parallel in creative tasks including writing [12, 58, 62], image generation [36, 52, 55], video generation [56], and video editing [13, 22]. While facilitating divergent thinking and decision-making, this new workflow of working with variations demands an additional stage of comparison. This has led researchers to explore ways to compare and understand AI-generated content. Tools have provided multiple levels of abstraction to support users in managing complex information generated by AI [47, 63]. To facilitate collaboration with AI in creative tasks, researchers explored sensemaking and comparison of writing variations [42, 58, 62], images [3, 11, 36, 41], slides [25], storyboards [9], and designs [50]. However, video presents unique challenges due to the temporal nature and multi-modal complexity. We contribute an interface that supports sensemaking of multiple variations of videos and provide empirical insights into how creators utilize variations when co-creating videos with AI.

## 3 Formative study

To better understand how video creators work with alternatives and the support they need for comparing and managing them, we conducted a formative study with 8 professional video editors. The formative study consisted of semi-structured interviews and a comparison task with 3 videos.

### 3.1 Method

**Participants.** We recruited 8 professional video creators using Upwork<sup>1</sup> (P1-P8, §A Table 4). We compensated the participants with their self-set hourly rate (\$28, SD=5.3) for a 1-hour remote

<sup>1</sup><https://www.upwork.com/>

study conducted via Zoom. Participants had an average of 10.5 years of experience ( $SD=3.0$ ) and created a wide variety of videos including commercials, interviews, documentaries, vlogs, and short-form videos.

**Procedure.** We started the study with demographic and background questions about the participants' video editing experiences and current approaches for reviewing and comparing videos. In the comparison task, participants reviewed and compared 3 videos edited from the same source footage. We pre-uploaded videos to YouTube so that users could use a familiar interface that offered speed controls, hover thumbnails, and the video transcript. To become familiar with the original version, participants first skimmed the source footage. Then, the participants reviewed the 3 edited versions for 30 minutes. They were allowed to take notes to help them remember any details. At the end of the task, participants were asked to summarize the differences between the three videos and choose one they preferred, explaining their choice. We conducted a post-task interview to understand their strategies for comparing videos and any challenges they encountered.

**Materials.** For the comparison task, we used 2 different source videos (F1: cooking tutorial [70], F2: TED talk [65]). We chose these videos aiming for source footage that was largely unedited and varied in duration (F1: 1hr 11min, F2: 12 min) and visual information (F1: featuring many objects, F2: primarily talking head). We used 2 approaches to create multiple edited versions of the source video – 1) with AI and 2) with experts. This allowed us to see how users work with variations created today with the current capabilities of AI video tools and in the future when AI capabilities improve. For AI edits, we created 3 versions of the summary for each source video using an LLM to extract subsets of the video. Using the video transcript, we instructed GPT-4 to generate 3 summaries on the video topic (F1: How to make Eggs Benedict, F2: Talk highlights). We also hired 6 professional video editors to manually create 3 edited versions of F1 and F2<sup>2</sup>. The video editors made a variety of editing decisions including trimming content to enhance focus, adding subtitles, B-roll footage, and motion graphics, incorporating text overlays and background music, and making audio enhancements. Each participant was given 3 edited versions from the same source video created with the same approach either with AI or manually. For F1, LLM edited videos were 152 seconds ( $SD=24$ ) and expert edited videos were 451 seconds ( $SD=152$ ). For F2, LLM edited videos were 283 seconds ( $SD=110$ ) and expert edited videos were 323 seconds ( $SD=21$ ).

**Analysis.** To understand how participants explore alternatives, we analyzed the study recordings and participants' notes. We transcribed the interviews and participants' spontaneous comments during the task and grouped the transcript according to 1) current practices, 2) comparison strategies, and 3) opportunities and challenges of video editing via comparison.

## 3.2 Findings

**Current Practices.** Participants reported that decision-making via comparison is a common process throughout the video editing process. They mentioned having to review and select which clips to

include from the footage provided by their clients (P1, P3-4, P7-P8), decide on the appropriate B-roll images (P2, P4), and experiment with different logo styles or placements (P5-P6).

Participants often had to create multiple versions of videos upon clients' requests (P1, P3-P8). This was because clients often found it challenging to imagine the final look of the videos from the editors' descriptions alone. P5 and P6 who created ad videos for marketing companies created two versions for companies' A/B testing to see which version is more effective. Clients often requested videos that have different lengths, storylines, intro versions, caption styles, or music styles. P3 and P4 mentioned that clients often liked different aspects of multiple versions so they had to combine the versions in further editing. While creating multiple versions helped clients narrow down and make decisions, it required a lot of time and effort from video editors. P5 said *"I know they (clients) want to see three completely different videos, but I try not to because it takes forever. I only tweak small things that involve minor changes and can be done quickly."* As a workaround, participants created variations and got intermediate feedback in each stage (e.g., sharing multiple variations of scripts before moving on to the storyboard) so that they could minimize the amount of rework required later in the process (P1, P5, P8). Alternatively, P4 and P6 created multiple versions of the introductory part of a video instead of the full videos as introductions can capture the essence of what the rest of the videos will look like.

**Comparison Strategies - Watching Videos.** In the comparison task, participants watched the 3 versions of the edited videos to understand their differences. 5 participants watched each video from beginning to end before moving on to the next video (P1-P4, P6) while 3 participants switched to play other videos (P5, P7-P8). P7 first watched the introductions of all three videos because it was easier to remember and compare shorter segments and the introductions often captured the holistic style of the video. P5 and P8 switched to other versions of edited videos when they were not satisfied with the current version's edits (e.g., text effects) to check if other videos offered better alternatives. While reviewing edited videos, P5 revisited the original unedited video to check which part of the source video was missing. 4 participants re-watched the same video multiple times to identify parts they missed in the first review attempt (P3, P5-P6, P8).

**Comparison Strategies - Taking Notes.** 7 participants took notes while reviewing edited videos to better recall the differences (P2-P8). They took notes on the content covered or omitted (e.g., *yolk separation step is not included*), errors (e.g., *mid-sentence jumps, flash frames*), what they liked or disliked about the edits applied (e.g., *Nice intro music, Too many text effects*), how they would have edited differently (e.g., *I would have emphasized the end results by showing it longer*), and how the video compared to other versions (e.g., *This is better at showing the sequence of events compared to the first one*). Participants reviewing videos edited with LLM noted errors in the edits, such as sentence cut-offs and jump cuts. P1 did not take any notes because pausing and taking notes would make the reviewing process much longer.

**Challenges.** All participants mentioned the comparison process to be very time-consuming as they had to watch multiple videos to identify the differences. As a workaround, 5 participants sped

<sup>2</sup>See Supplementary Materials for examples.

up the video and skipped around the videos to find the edit points (P3-4, P6-P8). Although the transcripts of the edited videos were provided, participants noted that simply reviewing the transcript or thumbnail previews on the timeline cannot fully replace the watching experience. P1 explained “*Simply reading the transcript isn’t enough because the transcript might seem fine, but the visuals might have been messed up.*” P3 who checked for edit errors in the video noted that errors such as sentence cut-offs or jump cuts cannot be spotted using the transcript or by scrubbing through the thumbnails; she had to watch the entire video to identify them.

Comparing edited videos was also mentally demanding as participants had to compare multiple aspects of videos at the same time (e.g., storyline, visual effects, B-rolls, subtitles, music). As a result, participants often reviewed each version multiple times to focus on reviewing a single aspect at a time (P3, P5-P6, P8). P5 described “*I cannot focus on multiple aspects at once while reviewing. When I’m checking the colors, I only look at colors.*” When comparing the content, it was challenging to switch between multiple videos as they extracted different parts of the original video. For instance, P3 wanted to check if all videos included a specific cooking step but found it difficult as the step appeared in different timestamps across the videos.

Participants mentioned that it requires a lot of effort to manually identify edits made throughout the video, take notes to remember details, and understand both high-level and low-level differences among videos from their notes. In the post-task interview, P5 mentioned that he wanted AI to describe the high-level theme or goal of each edited version to help him better understand what he was reviewing and what to expect. P6 wanted to have a list of all the edits made in each video so that he would not miss subtle edits (e.g., color correction, small text effects) and could quickly jump to those specific edit points for preview.

### 3.3 Design Guidelines

Video editors in our formative study currently make decisions by comparison throughout the video editing process and mentioned the benefits of having multiple variations of edited videos. We envision that in near future, Human-AI video co-creation tools will more commonly provide users with multiple variations. Our formative study reveals design guidelines (D1-D6) to better support video creation with alternatives. Our work aims to address D1-D4 and D6 while leaving D5 as a future work.

#### D1. Minimize redundant watching by aligning variations

Participants found it difficult to repeatedly watch overlapping sections of edited videos to identify edit points and differences. They wanted to compare edited variations with the original video and across alternatives but found it challenging with mismatched timelines. Aligning variations can simplify comparisons and reduce redundant viewing.

#### D2. Support quick skimming by highlighting differences

Prior work explored comparison in text [31, 58] or images [3, 36], but the temporal aspect of videos makes it challenging to skim them for comparison. This can be time-consuming and tedious as video lengths increase or as more videos are compared. A system that describes or visualizes the differences among edited videos can help users quickly make decisions.

#### D3. Enable independent comparison for different editing stages

Compared to traditional video editing where editors go through an interactive process of making sequential edits and reviewing, recent AI products such as CapCut or OpusClip speed up this process and generate fully edited videos. The resulting videos can be mentally demanding to compare, as users have to consider multiple aspects (e.g., rough cuts, subtitle styles) together. Enabling users to compare aspects independently can reduce the cognitive load.

#### D4. Support comparison via multiple modalities

Relying on a single medium, like transcripts or thumbnails, is insufficient for comparing video variations. Edits in visuals or timing can be unnoticed when only text is reviewed. Supporting comparison through multiple modalities—such as transcripts, visual thumbnails, and audio—enables users to identify various types of differences in the most effective way and ensures a more thorough comparison of edits.

#### D5. Support verification of edit suggestions

As AI-powered tools automate suggestions, prior work has explored verifying AI’s recommendations [27, 36]. In video creation, automated edit suggestions can often lead to additional errors (e.g., trimming a clip leading to jump cut) and thus needs careful validation. Future work can explore automatically identifying and communicating such erroneous suggestions to users.

#### D6. Support management and customization of variations

Formative study participants who reviewed a small number of alternatives all mentioned that they wanted to further refine the edits as none of the versions fully aligned with their preferences. Scaling up the number of alternatives increases the chance of finding a preferred version but it also makes it more difficult to manage and select from a larger pool of variations.

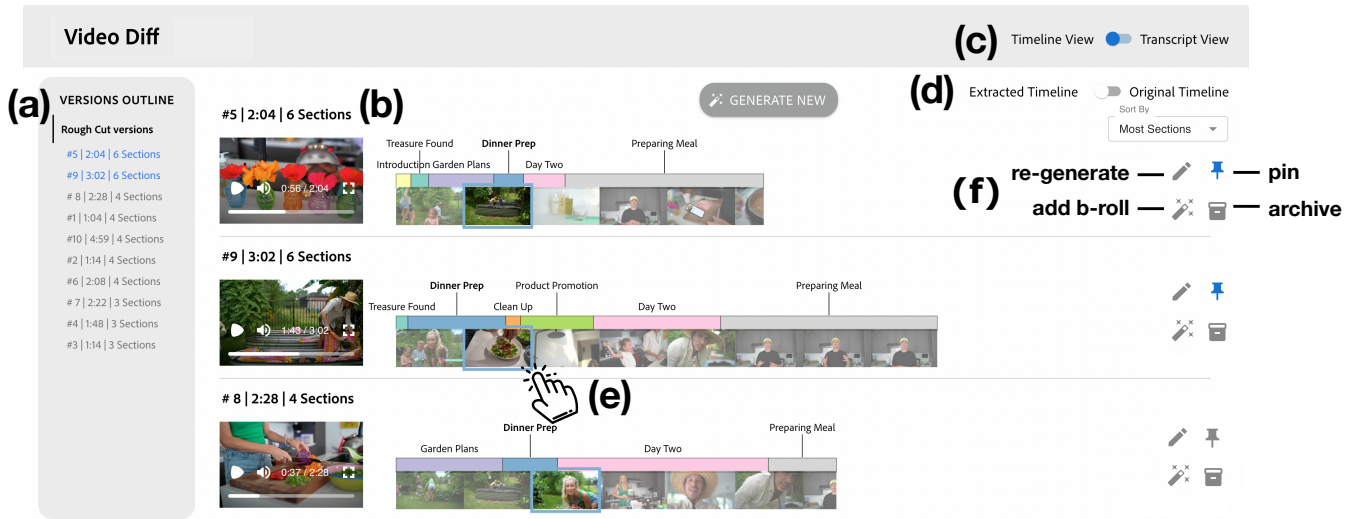
## 4 VideoDiff

We present VideoDiff (Figure 1), a human-AI video co-creation tool designed to support efficient video editing with alternatives. With VideoDiff, users can generate and review diverse AI recommendations for 3 different video editing tasks: making rough cuts, inserting B-roll images, and adding text effects (D3). VideoDiff supports easy comparison between alternatives by aligning videos (D1) and highlighting differences using timeline and transcript views (D2, D4). Users can organize and customize edits by sorting, refining, and regenerating AI suggestions (D6).

### 4.1 Scope

While VideoDiff provides visualizations that allow users to quickly identify and skim differences between video versions (D2), it is designed to go beyond the functionality of a *diff tool*, which typically focuses solely on highlighting changes or differences for review [6, 66]. Instead, we designed VideoDiff as a human-AI co-creation tool that facilitates iterative *generate-compare-refine* interactions, revealing how comparing alternatives can influence and enhance users’ creative processes.

To help set context for VideoDiff, we define the different stages of video editing. Rough cut creation is selecting good moments (or



**Figure 2: Overview of VideoDiff:** Users can view an outline of the variations in the current editing stage (a). In this figure, we see 10 rough cut variations. The user can play videos of these different versions (b) and compare them in the transcript or timeline view (c). Users can toggle between the edited and source timelines (d) to align videos to the source or edited context or click on each section to navigate directly to that part of the video (e). Users can also sort variations by duration and the number of sections included, as well as pin, archive, or edit variations according to their preferences (f).

clips) from source footage to create a compelling story. Inserting B-roll images or video helps to make the video more interesting and dynamic. To insert good B-roll, creators should find effective images or videos that illustrate what is being said and place them appropriately. Inserting text effects helps emphasize parts of the narration through animation and stylized text. The AI algorithms for each of these editing tasks are in themselves hard technical problems, and we do not focus on them in this work. Our focus is on supporting users to work with the generated variations as part of the editing process. However, because we must support some editing in order to test our ideas, we have implemented basic versions of editing algorithms that leverage LLMs to process the video transcript and recommend rough cuts, B-roll, and text effects. A holistic solution with multimodal analysis of video, audio, and narration can offer better editing suggestions in the future.

Most video editing softwares offer more than the three stages that VideoDiff supports, such as applying color correction or cleaning up the audio. Our goal is not to create a fully functional video editor, but rather to explore how video editing might change as AI technologies enter more and more of the editing stages. As AI easily generates multiple editing recommendations, the video editing task shifts to involve more curation beyond just editing, and we must consider how to best support this transition.

## 4.2 Interface

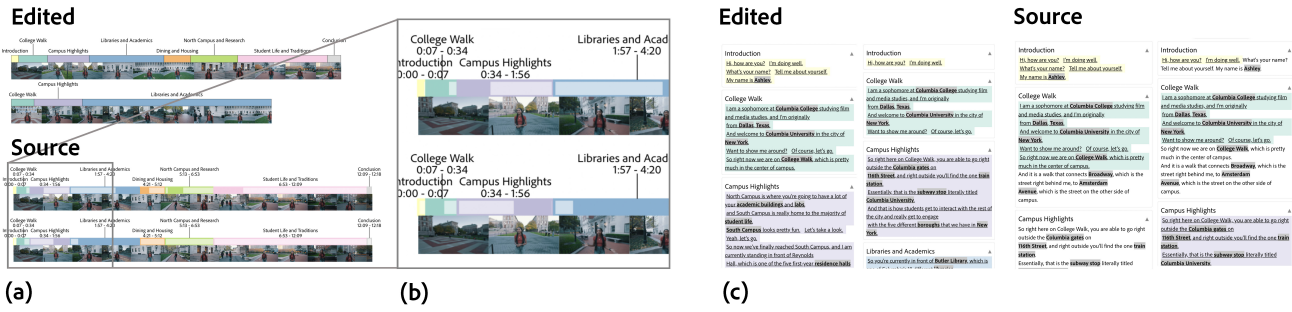
VideoDiff is a web-based video editing tool (Figure 2) where users can generate, review, and customize video alternatives. When the user uploads a video, VideoDiff first generates 10 rough cut recommendations. In the *versions outline* on the left, users can see the list of all variations in the current editing stage (Figure 2.a). Users can play the video of each version (Figure 2.b) or skim the differences



**Figure 3: At each editing stage, VideoDiff provides glanceable timelines for users to easily compare different variations. Users can click on any section, B-roll image, or text effect to jump to that part of the video and preview the effects.**

using a timeline or transcript view (Figure 2.c). Users can toggle between showing only the edited or all of the source content for additional context (Figure 2.d). They can also click on any section to navigate directly to that part (Figure 2.e). Users can sort, re-order, pin and archive the variations to organize them (Figure 2.f). Or they can refine or recombine existing variations or regenerate a new variation using text prompts.

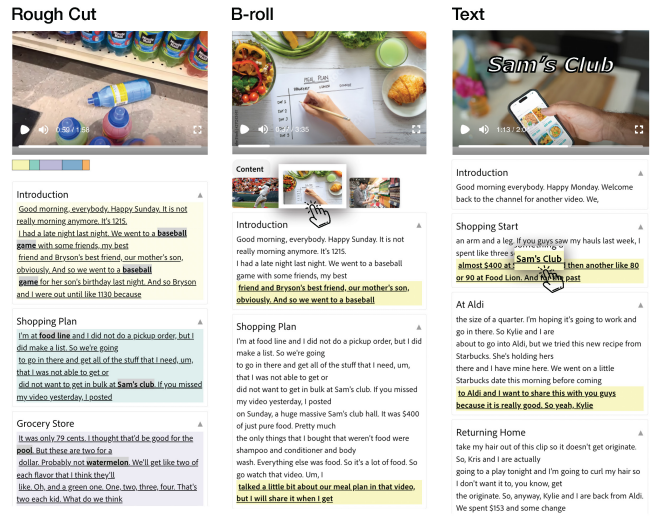
**Comparing Variations with Timeline View.** VideoDiff provides different visualizations of timelines at each video editing stage (Figure 3) so that users can quickly review and compare multiple alternatives (D3). In the rough cut stage, VideoDiff uses *sections* to visualize the timing and coverage of content in each variation. Drawing upon prior work that has shown that grouping footage



**Figure 4: Users can switch between the edited and original source timeline in the transcript (a) and timeline (b) views. This helps users see the edits in the context of the source content and compare which sections are included at a glance. The source view (c) shows users the location of the edited content in the context of the source view.**

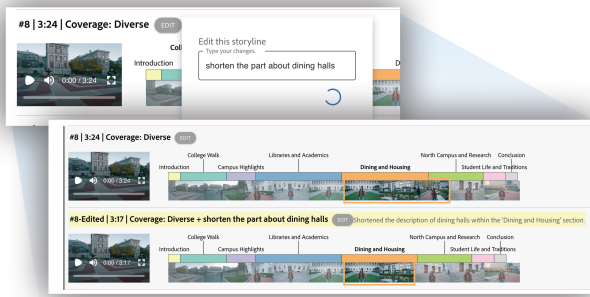
into thematically coherent chunks can help creators make video editing decisions [37, 45], we explore how chunking into sections can aid comparison of edited videos. Instead of segmenting each rough cut variation to identify sections, VideoDiff extracts sections from the source footage and applies them consistently across all variations. This allows users to easily align the variations and see how different versions include or exclude certain sections and cover varying parts or lengths of each (D1-D2). Users can toggle between edited and source timeline views (Figure 4.a) to align the videos based on the edited or source video timeline. The edited timeline allows users to quickly skim the video’s overall duration, along with the placement and proportion of each section. By hovering over a section, users can view relevant video thumbnails and click to play the video from that point. With the source timeline, users can easily understand which parts of the source video are extracted. In the formative study, video creators often compared edited videos to the source videos to verify any missing key information. Drawing from this observation and aligning with D1, we use the source video as an anchor for aligning rough cut variations for comparison. For example in Figure 4.b, while the edited timeline shows that both versions have a similar duration for “Campus Highlights” section, the source timeline reveals that each version covers different parts of the source footage on “Campus Highlights”. By hovering over the lighter-colored boxes indicating excluded parts, users can see which video thumbnails and topic keywords are not covered.

Once the user chooses a rough-cut version, they can generate 10 videos with different B-roll recommendations. In the B-roll stage, VideoDiff’s timeline view shows the B-roll thumbnails on top of the rough cut timeline bar, along with the keyword used to search for the B-rolls (Figure 3). Users can hover over the B-roll thumbnails to see which video scene each B-roll covers and click thumbnails to play the video and preview how the B-roll is inserted into the footage. After selecting a B-roll version, users can generate 10 new videos with different text effect suggestions. Instead of showing video thumbnails with text effects, we display the keywords where the text effects are applied, as text effects on thumbnails are too small to skim. Users can hover over a keyword to view the narration sentence for context, and click on the keyword to play the video and preview how the text effects are integrated into the footage.



**Figure 5: At each editing stage, VideoDiff provides glanceable transcripts so that users can easily review and compare different variations. In rough cut transcripts, visually concrete keywords [45] are emphasized in bold, allowing users to easily skim through the content of each variation. Users can click on section headings, B-roll images, or text effects to jump to that part of the video and preview the effects.**

**Comparing Variations with Transcript View.** VideoDiff’s transcript view allows users to quickly skim the transcripts to understand the differences of variations in each editing stage (Figure 5). The transcript is divided into sections and users can click on the section headings to play the video from the starting point of each section. During the rough cut phase, VideoDiff aids users in maintaining context while toggling between transcript and timeline views by showing a mini timeline within the transcript view. It also uses consistent color coding to denote sections across both views. In the transcript, visually concrete keywords [44] are emphasized



**Figure 6: Using VideoDiff, users can edit a variation, recombine multiple variations, and generate a new variation using text prompts. For each new generation, VideoDiff summarizes the changes so that users can easily verify the result.**

bold, allowing users to skim through the content of each variation<sup>3</sup>. Similar to the timeline view, users can also switch to the source transcript view (Figure 4.c) and see the complete transcript and identify which parts of the source text are excluded in each section. We synchronize the scrolling across multiple transcripts, allowing users to easily compare variations side-by-side (D1).

For quick skimming of B-roll and text effects options, VideoDiff highlights sentences in the transcript where these effects are applied (D2, Figure 5). Users can click on the B-roll thumbnails or text effects to preview the video from the moment the effect appears.

**Organizing Variations.** To help users to easily explore and narrow down the search space of alternatives, VideoDiff supports sorting of variations in each stage. In the rough cut stage, users can sort based on the edited video duration and the number of sections, and in the B-roll and text effects stages, users can sort based on the number of effects and the media type of B-rolls (image, illustration, video) or style of text effects (title, subtitle, lower thirds). Additionally, users can manually pin preferred versions to the top or archive unwanted versions to the bottom, which is also reflected in the outline view.

**Customizing Variations.** When users are not satisfied with the initial recommendations of VideoDiff, they can further edit existing versions or generate new alternatives (D6, Figure 6). Users can provide a text prompt to guide the generation of new alternatives (e.g., *Show many text effects when I talk about grocery items.*), or click “*Surprise me*” to have VideoDiff suggest a new alternative that is different from existing variations. Users can also recombine existing versions by specifying the version IDs and how to merge them in the text prompt (e.g., *Use first two B-roll images from #3 and last B-roll image from #7.*). By default, generated new results are pinned to the top for easy discovery. Using VideoDiff, users can also edit existing versions with a prompt (e.g., *Shorten the part when I’m talking about the meal plans.*). The new generation from the edit prompt is displayed right below the original version for quick comparison and VideoDiff also describes specific changes made for quick verification of edits (D2). For example, in Figure 6 the user tells VideoDiff to “shorten the part about dining halls” and

<sup>3</sup>We used GPT-o to identify visually concrete keywords with few-shot examples.

the system responds: “Shortened the description of dining halls within the Dining and Housing section”.

### 4.3 Implementation & Prompt Engineering

We implemented VideoDiff using React.js and d3.js. For embedding a video player, we used Remotion [57] to render the edited video and overlay B-rolls and text effects. When users upload a video, VideoDiff uses OpenAI’s Whisper API to transcribe the video. VideoDiff is powered by OpenAI’s GPT-4o and uses prompt engineering for 1) segmenting video into sections and identifying visually concrete keywords, 2) generating multiple alternatives of edit recommendations for rough cuts, B-rolls, and text effects, and 3) parsing and executing users’ new generation prompts and summarizing changes. To ensure VideoDiff provides diverse edit recommendations in each stage, we use *augmentation prompts* (§B) to control the generation of suggestion (e.g., by specifying duration and section coverage for each rough cut recommendation).

## 5 User Evaluation

To understand how well VideoDiff assists users in comparing multiple AI-edited videos, we conducted a within-subjects study with 12 video editors comparing VideoDiff to a baseline. Our study aims to investigate the following research questions:

### RQ1. How well does VideoDiff support video comparison?

Compared to the baseline, VideoDiff significantly decreases the time in video comparison (H1), improves comprehension and accuracy in video comparison (H2), lowers cognitive load (H3), and is more useful for video comparison (H4).

### RQ2. How well does VideoDiff support video authoring?

Compared to the baseline, VideoDiff significantly increases satisfaction in the final video (H5), provides better creativity support (H6), and is more useful for video authoring (H7).

### RQ3. How does VideoDiff impact video editing workflows?

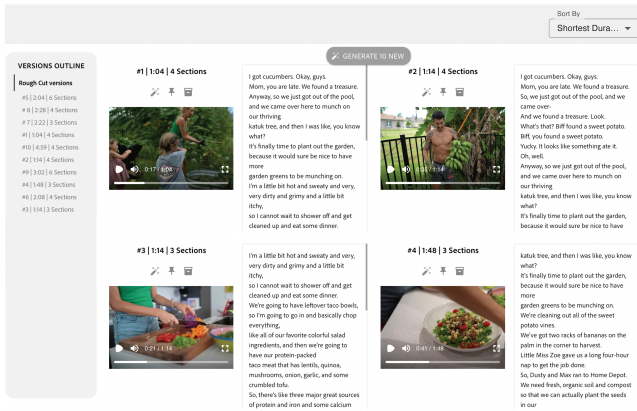
## 5.1 Method

**Participants.** We recruited 12 participants with diverse video editing experiences using mailing lists (P9-P20, Table 4). 6 of them (P9-P14) described themselves as proficient, having 8.33 years of video editing experience (SD=2.94). The other 6 participants (P15-P20) identified as beginners with 3.83 years (SD=1.17) of experience. We compensated \$75 to professional video editors and \$30 to amateur editors for the 1.5-hour remote study conducted via Zoom.

**Baseline.** The baseline (Figure 7) shares a similar UI design and features with existing AI video editing tools that support generating multiple videos (e.g., OpusClip, CapCut). When the user uploads source footage, the baseline system generates 10 different videos with different rough cuts and B-rolls applied. Users can sort the clips based on the duration and the number of effects.

**Materials.** We selected 3 videos (V0-V2, Table 1) from YouTube that feature primarily raw footage with few edits, real-world camera footage rather than screen recordings, and have narration. V0 was used in the tutorial session for both VideoDiff and baseline conditions. For the main study, we selected two videos (V1-V2) from the same YouTube channel which are similar in terms of length,





**Figure 7: The baseline interface shares a similar UI design and features with existing AI video editing tools that support generating multiple videos (e.g., OpusClip, CapCut).**

amount of narration, and scene changes. We asked participants to watch and familiarize with the original videos prior to the study.

Video ID	Video Type	Duration	Task	URL
V0	Campus Tour	12:28	Tutorial	[1]
V1	Grocery Haul	12:41	Comparison	[60]
V2	Grocery Haul	13:10	Comparison	[59]

**Table 1: Videos used in the evaluation study.**

**Procedure.** We started the study with demographic and background questions about the participants’ current video editing practices. Next, we gave a 10-minute tutorial on both VideoDiff and the baseline interface using V0. The participants then completed two tasks: the *comparison task (RQ.1)* and the *authoring task (RQ.2-RQ.3)*. The comparison task studies how users use VideoDiff when the comparison is an end goal in itself and the authoring task explores when the comparison is the means to accomplish the higher-level goal of creating the final video. In the comparison task, participants reviewed 10 different videos using VideoDiff or baseline to answer comparison questions (Table 2). We derive these questions from the formative study to consider what aspects video editors compare to select the video they preferred: the content of the videos (e.g., which key visual scenes are included or which key concepts are mentioned) and the editing techniques applied (e.g., the style, placement, and frequency of visual effects). We aimed to cover diverse types of tasks, including both single-select and multi-select formats, as well as questions that required participants to assess narration, visuals, or a combination of both. Also, our questions cover both comparison modes identified by Gleicher et al. [32]: *explicit comparison* where targets are known and available (e.g., Q4. *Which video mention her salad recipe?*) and *implicit comparison* where targets are hidden, requiring more exploratory search (e.g., Q3. *Which video features the highest proportion of talking heads?*). In each interface condition, users were assigned to one of V1-V2. For each question type, we measured the time taken, answer accuracy, and interaction logs to understand which system components participants utilized

to perform each comparison. For accuracy, we assigned a rating of 1 if the answer was completely correct, 0.5 if the participant got half or more of the answers right for multi-select questions, and 0 if less than half or none were correct. After each interface, we conducted a post-stimulus survey that included the following ratings: mental demand, performance, effort, frustration, and usefulness of the system in understanding differences between videos. All ratings were on a 7-point Likert scale.

In the authoring task, participants reviewed the video suggestions and further edited videos using VideoDiff or baseline. In each interface condition, users were assigned to one of V1-V2. We guided participants with an editing scenario: “*You have created this initial video about grocery haul. To edit this video you can review AI’s 10 different edit suggestions to select the final video or regenerate, refine, or recombine edit suggestions.*” After each interface, we conducted a post-stimulus survey to measure users’ level of satisfaction with the final video, creativity support indexes (exploration, engagement, effort/reward tradeoff, expressiveness) [16], and usefulness of the system in video authoring. In both the comparison and authoring tasks, the order of the interfaces and videos was counterbalanced and randomly assigned to participants.

## 5.2 Comparison Task Results

Overall participants rated VideoDiff to be significantly more useful in understanding differences (Figure 9,  $\mu=4.92$ ,  $\sigma=1.16$  vs.  $\mu=2.25$ ,  $\sigma=1.06$ ;  $Z=-2.77$ ;  $p<0.05$ ). In this section, we share 1) participants’ strategies for reviewing and comparing videos to answer comparison task questions, 2) task results (RQ1), and 3) perceived workload (RQ1).

**Comparison Strategies.** To narrow down the search space, 5 participants in the baseline condition (P9, P11, P14, P17, P19) and 6 participants in VideoDiff condition (P11, P14, P17-P20) first sorted variations based on the comparison question (e.g., sort by longest duration for Q1). With VideoDiff, 2 participants first identified which section is relevant to the question (e.g., “Meal Plan” section for Q9) then only checked within the variations that covered the section (P9, P18). All participants using VideoDiff switched between timeline and transcript views depending on the question type. They used the edited view to find what each variation covers, and 2 participants additionally utilized the source view to align videos side by side (P10, P19). 3 participants were initially unsure which view would be more helpful for a specific comparison question and tried switching views multiple times (P12, P14-P15).

To answer the questions that require visual comparison between videos, all 12 participants in the baseline condition frequently interacted with the video player of each variation directly. Participants scrubbed the video player timeline to skim through the visual content of the video or click-jumped to navigate to different parts of the video. After finding the visual of interest in one video, 4 participants tried to navigate to the similar timestamp in other videos but realized they needed to re-search as all videos extract different parts of the video. As searching in each variation took time, P12 took notes to remember the details in each variation. She noted “*10 is a lot to compare at once. I need to take notes.*” Because participants could not preview thumbnails of multiple videos in parallel, 2 participants tried playing multiple videos synchronously but found it

Edit Stage	Category	Example Questions
Rough-cut	Single-select	(audio) Q1. Which video talks about most grocery items? (audio) Q2. Which video mentions the calories of strawberry lemonade? (visual) Q3. Which video features the highest proportion of talking heads?
	Multi-select	(audio) Q4. Which videos mention her salad recipe? (visual) Q5. Which videos show the scissors?
B-roll & Text effects	Single-select	(visual) Q6. Which video has the most B-rolls showing food? (visual) Q7. Which video shows the B-roll image when talking about the protein bar? (both) Q8. Which video shows a B-roll image when talking about Z-bar?
	Multi-select	(visual) Q9. Which videos show B-rolls with hamburgers? (both) Q10. Which B-rolls are shown when talking about the restaurant?

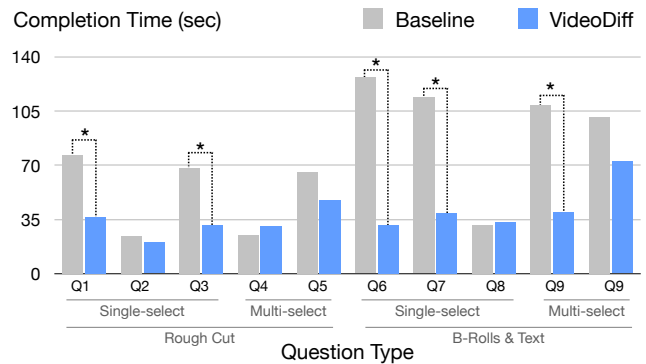
**Table 2: Questions used in the user study’s comparison task included both single-select (choosing one answer that meets the criteria) and multi-select (selecting all answers that meet the criteria) formats. These questions required participants to check either the narration, visuals, or both, with the specific questions varying depending on the assigned video.**

challenging to compare them as they were not aligned (P13-P14). As it was time-consuming to review the visuals of each variation, 4 participants used a workaround by searching for relevant keywords in the transcript (P14, P17-P19). P19 searched “I got ” to check what grocery items are mentioned in the narration.

Participants using VideoDiff mainly checked the timeline view to identify the shots included in each variation for rough cut comparison. If a timeline did not display a thumbnail of the visual they were searching for, 3 participants further ensured that the missing thumbnail was not simply due to the periodic nature of frame extraction while the visual is actually present in the video. P9 and P15 located sections that have thumbnails with similar background and shot type, navigated to the part and scrubbed the video player to skim the visuals. Similarly, P19 reviewed the *source timeline view* to check whether the frame of interest was present outside of selected sections. When looking for specific B-roll images, P18 mentioned the benefit of visual keyword search using the B-roll keywords in the timeline view. To check the video context when a B-roll or text effect was applied, participants either reviewed the thumbnails below the effects or clicked on the effects to preview them.

For questions that involve checking the audio (*i.e.* narration) of the video, participants in both conditions used keyword search (ctrl-F) in the transcript. In the baseline condition, 5 participants read the transcript of each variation sequentially which was time-consuming (P9, P11, P14, P18-P19). P19, a low-vision participant noted “In this version [Baseline], it is mostly reading than watching so very tiring for my eyes!” With VideoDiff, 2 participants first searched where the relevant part appears in the transcript in one of the variations, then scrolled horizontally to quickly compare whether other variations cover the part using *source view* (P9, P14).

**Time and Accuracy for Comparison Task.** We report the results of the comparison task in Figure 8 (task completion time) and Table 3 (task accuracy). On average, participants spent 38 seconds (SD=14) for each comparison question with VideoDiff, roughly half the time spent with the baseline (74 seconds, SD=38). In 5 questions that required checking multiple parts of the video (*e.g.*, Which video has the most B-rolls showing food?), participants using VideoDiff were significantly faster than in Baseline. In the baseline, 8 participants failed to find the correct answer for one or more questions within the 3-minute limit (P9, P12-13, P15-P16, P18-P20).



**Figure 8: Average task completion time to answer comparison questions in Table 2. ( $p < 0.05$  is marked with \*)**

5 participants provided a rough guess answer instead of carefully comparing variations after watching the first few videos (P9, P12, P16-P18). P16 notes “In only three minutes? I’ll just have to guess as I cannot watch all these videos.”

Baseline participants often provided incorrect or incomplete answers when the question involved visual comparison. 3 participants who scrubbed the timeline to search for an object missed it when it only appeared for a short amount of time (P11, P17-P18). While 2 participants searched keywords related to the visual in the transcript to speed up the search, some videos only mentioned the object and did not show it (P12, P17). While VideoDiff’s filmstrip in the timeline enabled participants to more quickly compare the visuals of multiple variations, 5 participants missed the visual search question (Q5) when the object asked only appeared briefly in the video and was not captured in the filmstrip.

**Cognitive Load in Video Comparison.** Figure 9 shows the distribution of survey ratings. Our NASA-TLX [34] results indicate that VideoDiff required significantly less mental demand ( $\mu=4.33$ ,  $\sigma=1.4$  vs.  $\mu=2.75$ ,  $\sigma=0.97$ ;  $Z=-2.08$ ;  $p<0.05$ ), temporal demand ( $\mu=4.33$ ,  $\sigma=1.72$  vs.  $\mu=2.33$ ,  $\sigma=1.15$ ;  $Z=-2.36$ ;  $p<0.01$ ), effort ( $\mu=4.42$ ,  $\sigma=1.08$  vs.  $\mu=2.92$ ,  $\sigma=1.08$ ;  $Z=-2.62$ ;  $p<0.05$ ), and frustration ( $\mu=3.75$ ,  $\sigma=2.01$  vs.  $\mu=2.33$ ,  $\sigma=1.37$ ;  $Z=-1.78$ ;  $p<0.05$ ). We did not see any significant

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10
<b>Baseline</b>	0.67 (0.49)	1.00 (0.00)	0.42 (0.51)	0.96 (0.14)	0.63 (0.31)	0.50 (0.52)	0.36 (0.50)	0.92 (0.29)	0.46 (0.40)	0.67 (0.25)
<b>System</b>	0.92 (0.29)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	0.75 (0.34)	1.00 (0.00)	0.90 (0.32)	0.91 (0.30)	0.96 (0.14)	0.86 (0.23)

**Table 3: Task Accuracy for answering comparison questions in Table 2. The value represents the average accuracy (SD).**

difference in performance ( $\mu=3.08$ ,  $\sigma=1$  vs.  $\mu=2.17$ ,  $\sigma=0.94$ ;  $Z=-1.66$ ;  $p>0.05$ ). Participants using baseline mentioned the challenge of finding relevant parts with a video using a video player and transcript and having to repeat the process for multiple variations. P15 noted *“It is more tedious than demanding. All videos are not visually distinct so I have to do a lot of manual digging.”*

After using VideoDiff, participants appreciated having separate views to compare videos across timelines and transcripts. P11 mentioned the benefit of having both in separate views *“If I’m shown both views at once, it is too overwhelming and will not fit multiple versions in a page.”* P10, P13, and P19 highlighted that the same color coding of sections across timelines and transcripts helped them pick up information easily. P13 said *“While this one [VideoDiff] has much more information with many different view options, they are not overwhelming and all very helpful and complementary.”* P10 expressed that she wants to use VideoDiff in the future to show her clients multiple versions of videos she has created.

### 5.3 Authoring Task Results

Overall participants rated VideoDiff to be significantly more useful for creating videos (Figure 9,  $\mu=3.25$ ,  $\sigma=0.87$  vs.  $\mu=5.42$ ,  $\sigma=0.51$ ;  $Z=2.97$ ;  $p<0.05$ ). In the followings, we share 1) participants’ strategies for reviewing and customizing variations for creating a video (RQ3) 2) how well VideoDiff supports creativity (RQ2) and 3) future improvements suggested by participants.

**Reviewing Variations and Converging.** In the baseline condition, participants mentioned the difficulty of understanding the differences and quickly identifying which one aligns best with their preference. Thus, 3 participants quickly selected the longest video as they were likely to cover most of the original video (P12-P13, P16). P16 mentioned *“I cannot understand what each version is about, so I’ll choose the longest one and import it into my own video editor to cut down further.”* Similarly, when reviewing the baseline’s B-roll recommendations, P18 chose a version with the most B-rolls as it was easier to remove than add them manually. 4 participants who first edited videos with the baseline mentioned feeling overwhelmed by starting with 10 alternatives and wanted fewer variations (P10-P12, P14). However, when they used VideoDiff to create videos in the following sessions, P10 and P14 changed their opinion and stated that they liked having 10 variations as a start. P14 said *“Because I can quickly skim the differences with visuals and zoom out, the same amount of suggestions don’t feel overwhelming here.”*

When creating videos with VideoDiff, all participants started in the timeline view and narrowed down the search. To closely compare final candidates or verify edits, 10 participants switched to the transcript view for a side-by-side comparison. P16 said *“You cannot zoom out and view all text in the transcript at once. It’s possible with the timeline view.”* Similarly, P17 compared two views: *“Timeline view is helpful for creating bones of the videos, deciding what to show*


*for how long. Transcript helps with individual details.”* 2 participants never used the transcript view and finished editing in the timeline view (P17, P19). P19 said *“All video editing software that I use shows timeline, so I’m not used to editing with the transcript view.”*

In the rough cut stage, participants sorted and shortlisted variations based on the coverage of sections they were interested in. 3 professional video editors highlighted the benefit of segmenting the video into sections (P9, P11, P14). P14 mentioned *“This is exactly what I do manually every time I start video editing, using colors to differentiate my shots in Premier Pro. This system [VideoDiff] will save me so much time when I’m working on 2-3 hour-long videos.”* P15 who narrowed down the rough cut versions to 2 candidates first added B-rolls to both versions to inform decision-making in the rough cut stage. She explained *“If one version has a part that I can add nice B-rolls, I’ll go for that [Rough cut version].”* When reviewing B-roll and text effect recommendations, participants checked whether the effects were evenly distributed throughout the video using timeline view (P12, P17), whether the B-roll aligns with the narration in the transcript view (P9, P11), and whether the B-rolls and text effects are visually aesthetic and match the style of the video by previewing in the video player (P13).

**Generating and Editing Variations.** In the baseline condition, only 2 participants generated a new set of variations (P11, P13) while the other 10 participants selected the final video from the initial recommendation without further change. P19 said *“I wouldn’t generate new videos because that would require another comparison!”* After generating new recommendations, P11 noted *“It is difficult to know how these videos are different from the original variations.”*

Participants who created videos using VideoDiff made 4.33 edits to existing variations (SD=2.42). 2 participants combined existing versions (P13, P19) and 3 participants generated new variations (P11, P16-P17). During the rough cut stage, many participants made edits related to sections, such as adding or removing sections, adjusting the coverage of sections, or reordering them. Participants made more edits focused on visual changes in the timeline view (P20 edited *“Show less talking heads, show more grocery items”*), while they made more edits related to narration in the transcript view (P17 edited *“Do not mention the store name”*). While most participants iteratively made edits on the first rough cut variation they selected, P18 switched to another version after making 3 edits on a rough cut version. He described *“After making a few edits on this one, I realized that I didn’t really like it. It’s really useful to explore multiple paths at once, and be able to switch between them easily.”*

In the B-roll and text effects stage, participants made edits to add/remove an effect, move the effect to a different part of the video by specifying a timestamp or a transcript sentence, and change the B-roll images or text styles. Participants noted that VideoDiff’s description of specific changes makes it easy to verify new changes without having to play the videos before and after the changes.



**Figure 9: Distribution of the rating scores for the Baseline ● and VideoDiff ● (1 = negative, 7 = positive) in the two tasks. Note that a lower value indicates positive feedback and vice versa. The asterisks indicate the statistical significance as a result of Wilcoxon text ( $p < 0.05$  is marked with \* and  $p < 0.01$  is marked with \*\*).**


**Video Diff as a Creativity Support Tool.** We assessed VideoDiff and the baseline using creativity support index [16] (Figure 9). VideoDiff supported efficient exploration and comparison between variations ( $\mu=2.5, \sigma=0.9$  vs.  $\mu=5.33, \sigma=0.78$ ;  $Z=3.07$ ;  $p<0.05$ ), which enabled participants to quickly narrow down the search and spend more time customizing them. This led to higher effort–reward trade-off ( $\mu=3.42, \sigma=1.24$  vs.  $\mu=5.67, \sigma=0.89$ ;  $Z=2.89$ ;  $p<0.05$ ) and higher satisfaction in the final result in VideoDiff than in the baseline ( $\mu=3.67, \sigma=0.98$  vs.  $\mu=5.25, \sigma=1.06$ ;  $Z=2.63$ ;  $p<0.05$ ). After creating the final video with VideoDiff, P13 said “I can see that this version covers all sections in a balanced way and has many aesthetic B-rolls, so I’m confident that I chose the right one.”

Participants were also more engaged in the creative process with VideoDiff than the baseline ( $\mu=2.83, \sigma=0.83$  vs.  $\mu=5.67, \sigma=0.89$ ;  $Z=3.05$ ;  $p<0.05$ ). P15 compared the experiences in both conditions: “In the first session [Baseline], I had to read the redundant transcript over and over, and having to read most of the time when I’m trying to create a video is not so fun. With this one [VideoDiff], I didn’t need to do repetitive work and could experiment what AI can do, so really enjoyed using it!” For expressiveness, we did not see any significant difference between two conditions ( $\mu=2.83, \sigma=1.27$  vs.  $\mu=4.75, \sigma=1.14$ ;  $Z=2.46$ ;  $p>0.05$ ). All participants using the baseline interface felt disengaged from the creative process due to limited control over the generation and a lack of understanding of the variations. P19 stated “There was no creation in this process, I was just selecting like a judge.” Similarly, P11 said “While being able to generate 10 recommendations is impressive, I’m not providing any input so I don’t feel like I’m the one creating something.” In the VideoDiff condition, participants showed mixed perspectives towards VideoDiff as a creativity support tool. P11 described “Instead of spending most of the time reviewing, it [VideoDiff] helped me to quickly move on to the creative stage and spend more time and effort there.” P14 elaborated “This tool [VideoDiff] helped me with tedious choices so that I can spend more time on meaningful choices.” On the other hand, P10 mentioned that VideoDiff does not make her feel as expressive as using her existing tools. “While this is helpful for making things faster, it feels like I’m supervising AI instead of me being creative. I might feel less attached to what I create with this tool.”

**Future Improvements for VideoDiff.** Participants provided suggestions on how to improve VideoDiff in the future. First, P13 and P20 suggested that VideoDiff can provide alternatives whenever participants request edits to existing variations. Current VideoDiff provides only one edited variation based on the user’s edit prompt. Thus, when the new variation does not align with the user’s intent, they have to refine the edit prompt multiple times. P20 described “I cannot always be specific because I forget to specify all the details in the prompt or I often do not know what I want. When I ask it to “shorten grocery haul”, it can generate 3-4 versions with different parts dropped.” Second, 5 participants (P9-P11, P14, P16) wanted more direct control with VideoDiff. P11 described “While I can change the duration or change B-rolls using prompt-based editing, some of these are easier with directly stretching the timeline or right clicking on B-rolls.” P9 also suggested having a slider to control the number of B-rolls and text effects easily. P16, an amateur video editor mentioned “Instead of exporting this and further refining using Premier Pro, I just want to finish everything here and directly publish without having to work twice.” Third, P14 mentioned that future VideoDiff can improve its generation pipeline to consider visuals into account when generating rough cuts and inserting B-rolls. P14 explained “I want B-rolls to be placed where the jump cuts are or when there are no interesting visuals. Now it covers the original video’s bananas with a B-roll bananas.” Finally, P19 a low vision participant provided insightful feedback to make VideoDiff more accessible to blind and low vision video creators. For low vision users who may find reading transcript challenging, VideoDiff can include more detailed subheadings summarizing the transcript’s content, in addition to the current section titles, to reduce the amount of reading required. For screen reader users, providing more descriptive text for visuals (e.g., alt-text that describes differences between 2 similar B-roll images [36]) can support accessible visual comparison of videos.

## 6 Exploratory Case Studies

The controlled user study (Section 5) showed that video creators found it easier to comprehend differences and create videos using VideoDiff compared to the baseline. To learn how creators would use VideoDiff to edit their own footage, we conducted an exploratory study with 3 video creators (P9, P21-P22) who brought their footage.



**Figure 10: In exploratory case studies, creators edited their own footage (V3-V5) using VideoDiff.**

## 6.1 Method

We recruited 3 video editors (2 professionals, 1 amateur) using mailing lists and social media (P9, P21-P22). P9 also participated in the controlled user evaluation study. Before the study, we collected footage from each participant (Figure 10). P9 and P22 who wanted to use multiple clips in the study concatenated them into a single source video. During a 1hr remote study session, we asked participants background questions, provided a tutorial of VideoDiff, invited participants to edit their own footage with VideoDiff, and asked participants semi-structured interview questions about their experience. We compensated participants \$75 via Amazon Gift Card for preparing the footage and participating in the study.

## 6.2 Three Vignettes: VideoDiff in Context

**V3: Science Video Podcast.** P21 is a college student who creates videos for a personal YouTube channel and school projects. For his physics class, he created a video podcast that parodies golf announcers to explain Newton’s law in a fun way. His goal for editing was to make an informational and entertaining video while describing all three of Newton’s laws. With VideoDiff’s initial rough cut suggestions, he first shortlisted versions with many sections which are likely to cover all three laws. P21 noted *“It’d be nice if I could ask it to filter out versions that do not cover all three laws”*. While he chose a longer rough cut version to further edit, he also pinned a shorter version, which primarily consisted of jokes, to use as a teaser for his main video. When reviewing VideoDiff’s B-roll recommendations, he looked for a version with a more consistent style of B-rolls as he was making an educational video. He liked having abundant options for B-rolls and recombined suggestions to create a version that has many infographic illustrations. For text effects, he wanted to insert a chapter text for each law but noticed that most versions focused on formulas or jokes. He generated a new version by specifying *“Show a lower-third text of what each chapter is about.”* In the future, P21 wished that VideoDiff would gradually learn from his selection of alternatives and offer more personalized alternatives. *“Like reinforcement learning, I want it to learn my preferences over time so I can spend less time customizing.”*

**V4: Artist Intro Video.** P22 is a professional video creator who often edits music/art videos and documentaries for clients. He is making an artist intro video with the footage he got from a client and used VideoDiff to make edits. His goal for editing was to fit the video into 2 minutes while talking about the major art pieces. When reviewing rough cut variations, P22 first noticed that #7 and #8 are both close to 2 minutes, and compare them side-by-side in the transcript to choose a version that mentions more art pieces of the artist. While he tried to edit the video with a prompt *“Remove shots with flashing lights.”*, he noticed that they still remained in the video.

In the B-roll stage, P22 did not apply any B-roll and explained *“I want to show the artist’s own art images rather than stock images.”* He suggested that in the future, VideoDiff can recommend zoom and pan animations for B-roll, as well as movements for text effects, and display them side by side for comparison. He wanted to continue using VideoDiff as part of his workflow as the alternatives gave him new ideas. *“My brain starts to lock once I start manually editing. Even when I try to tweak to something different, it gets fixated easily. It will be always useful to have alternatives.”*

**V5: IoT Setup Tutorial.** P9 is a freelancer video creator who edits videos for YouTuber’s vlogs and product reviews. One of the YouTubers asked him to edit their raw footage of a IoT setup tutorial to remove mistakes and make a polished video. P9 first chose one rough cut version and tried to make a section longer by editing with a prompt *“Add more footage for hotspot installation.”*. After noticing that VideoDiff generated a too long version, he switched to the transcript view to check what specific sentences he want to add and refined the edit prompt. After reviewing the footage, P9 noticed a lot of repetition in the sentences as the YouTuber corrects themselves. He edited the video to remove the repeated sentences and also prompted VideoDiff to cut when the YouTuber is out of the frame. However, he noticed that VideoDiff hallucinated that it had cut the parts out when they were still present. He mentioned that while VideoDiff is great at making a rough version of a video, he still needed to rely on other tools to manually remove filler words and polish the video. P9 described *“Usually making a rough cut for a video this long takes about 2 hours, but today I did it in 5 minutes. Although I still have to polish it up, this is so useful!”*

## 7 Discussion

In this section, we reflect on our findings from the design and evaluation of VideoDiff. We also discuss future opportunities to improve and extend VideoDiff.

**Creating Videos with Alternatives.** VideoDiff is a tool designed to support human-AI co-creation of video with alternatives. Our focus on alternatives is driven by prior work that highlights the value of divergent thinking in creative process [24, 58, 62]. Video creators in our user study highlighted 3 benefits of having alternatives. First, it helped to consider alternative edit styles that they are not familiar with. Second, when users give the AI an ambiguous edit request, having multiple alternatives allows them to choose the one that best matches their intent, rather than relying on a single suggestion from the AI that might be incorrect. Third, with alternatives, creators can quickly prototype multiple versions of end-to-end videos adapted to different audiences, settings, and platforms. While editing videos with alternatives is not a one-size-fits-all solution, it represents existing practices in video editing and other generative AI tasks. In this work, we focus on minimizing the burden of working with alternatives in video creation and demonstrate how VideoDiff achieves it. Still, some users may prefer incremental editing approaches, such as refining a single video [76], or workflows that prioritize error correction over creative decision-making [37]. In the future, we will consider diverse workflows by letting users select the number of suggestions (one vs many) and further explore how the scale of alternatives or the timing of providing alternatives can impact users’ creation process.

**Scope and Use Cases of VideoDiff.** We reflect on the scope of VideoDiff based on user groups, creation stage, and control granularity based on the taxonomy of roles that creativity support tools play [21]. First, VideoDiff can support a diverse range of users. Novices or creators with accessibility needs (e.g., screen reader users) who find existing video editing software complicated to use can benefit from VideoDiff’s suggestions. Novices with limited experience with the authoring tool often do not know where to start [5] and find it easier to start with examples rather than from scratch [40]. For experienced editors, VideoDiff reduces the time spent on routine tasks such as planning and making rough cuts, allowing them to focus on fine-grained adjustments that improve the overall quality of their work.

In terms of the creation stage, VideoDiff primarily supports idea generation and execution but not evaluation, requiring users to review and evaluate the quality. While automating execution, VideoDiff still emphasizes user control, allowing for post-editing and customization through natural-language prompts. In terms of control, VideoDiff provides indirect control (i.e., text prompt) to prioritize a low learning curve over precise editing. This design choice contrasts with traditional video editing software, which offers precise but complex editing capabilities. Thus, it will be more useful in cases when users want to create videos under time constraints (e.g., social media content, event recaps) or budget limits (e.g., personal projects or prototyping for feedback before detailed editing). While users can export the EDL to further edit in other video editing tools, future VideoDiff can incorporate more direct control (e.g., timeline manipulation, dragging B-roll to adjust timing).

**Support for Flexible Comparison.** VideoDiff enables quick comparison of video alternatives by supporting users to align multiple videos and highlighting differences using timeline and transcript views. While creators using VideoDiff were successful with various comparison tasks, 2 participants were confused about which views to utilize for a few comparison tasks. To further minimize creators’ mental load in search and comparison, a future iteration of VideoDiff can directly take creators’ natural language comparison queries (e.g., *Which video has fewer narration mistakes in the intro of the video?*). It can extract relevant parts of each video and determine the best views to communicate the differences and present them to creators. This can further reduce creators’ effort to switch views and help scale up the number of variations for comparison. Also, VideoDiff’s pipeline can consider broader types of comparison beyond what is based on narration and visual content – shot types (e.g., *Which video features a higher proportion of a talking head?*), audio (e.g., *Which videos use calm music?*), errors (e.g., *Which video has the least sentence cut-offs?*), or abstract comparison queries (e.g., *Which video has a more hooking intro?*) Also, while D4 highlights comparison through multiple modalities, current VideoDiff focuses on visual and narration differences and does not consider sound or music differences. Prior work has explored audio tagging of videos [86] and generating sound effects to augment video skimming [54]. Similarly, future VideoDiff can consider audio-based comparison by surfacing differences in background noises, music inserted, or speech characteristics. As the pipeline evolves to identify and describe these various types of differences, finding effective ways to visualize them becomes another important challenge.

**Implications of VideoDiff on Creativity.** While AI-powered creativity support tools risk design fixation [73] and lack of control [62], we designed VideoDiff to mitigate these concerns. First, VideoDiff supports divergent thinking. Participants in our study mentioned that it is easy to get fixated on one video editing style with existing tools, whereas VideoDiff inspired them to explore alternative ideas, which they found valuable to incorporate. Second, VideoDiff supports customization. Instead of generating a single result for users to accept as-is, it offers multiple options, enabling users to select and refine suggestions. By not delegating all edit decisions to AI, VideoDiff ensures creators remain in control of the final result. Finally, VideoDiff promotes efficiency by automating tedious edits, allowing creators to spend more time in the creative stages.

To further support a deeper sense of agency and ownership, future iterations of VideoDiff can involve users in the planning and ideation stage [15] and take user inputs before generating suggestions. For instance, users could specify requirements such as duration or style, provide text-based instructions, or upload reference videos to guide the editing process. They could also upload a partially edited video—for example, one with the first half completed—and request the system to finalize the remaining edits. To support more flexible user control, future VideoDiff can allow users to directly accept or reject individual suggestions as in prior works [33, 43]. To improve the transparency of AI suggestions, we can provide explanations for the generated suggestions [74] or offer access to intermediate steps of LLM chaining used in generating them [81]. This can enable users to better understand, utilize, and experiment with the tool [85]. Furthermore, future work can consider a personalized recommendation pipeline within VideoDiff, where the tool adapts to a user’s unique style and preferences over time. This would not only speed up the editing process but also give creators a more tailored experience.

As AI continues to accelerate the video creation process, there is also a risk that users may not thoroughly review or verify all the generated content. VideoDiff can help mitigate this by encouraging creators to do more sensemaking before making final selections. By offering clear visualizations of differences and allowing users to explore how each version fits their specific goals or audience, VideoDiff can assist creators in making more informed decisions.

**Environmental and Societal Impact.** We discuss the long-term impact of AI-powered creativity tools such as VideoDiff on the environment and society. Like other AI tools, the language model that powers VideoDiff requires computational resources for training and inference which involves carbon emissions [80]. Yet, our study demonstrates that VideoDiff is highly valuable to users by streamlining complex editing tasks and enabling creative exploration. Also, VideoDiff may reduce the time users spend on traditional editing workflows, which often rely on resource-intensive video editing software and involve repeatedly rendering, storing, and transferring multiple large video files via cloud services. Instead, VideoDiff represents edit variations as structured text, eliminating the need to store multiple edited versions as full video files. We expect the financial and environmental costs of the models to change over time, and future work can explore how well smaller models can support tasks of suggesting video edits.

Generative AI tools like VideoDiff also raise concerns about job displacement in fields that rely heavily on manual video editing and production. However, VideoDiff does not automate the entire video editing process and retains key editing decisions with users who select and refine the suggestions. For novices, VideoDiff can lower the entry barrier to video editing and open doors for new roles. For experienced users, it reduces the time spent on tedious, repetitive tasks, allowing them to focus more on higher-order creative endeavors. By fostering both accessibility and efficiency, VideoDiff can promote collaboration between humans and AI in a way that creates opportunities.

**Limitations.** VideoDiff demonstrates the usefulness of video creation with alternatives in 3 representative video editing tasks: making rough cuts, inserting B-rolls, and adding text effects. Our study participants expressed excitement to use VideoDiff for a wider range of edits (e.g., motion graphics, and transitions). Future versions of VideoDiff could incorporate more diverse editing types.

While VideoDiff currently takes a single source video as input, study participants mentioned that VideoDiff will be useful for generating and comparing alternative curations of multiple clips when provided with a huge library of footage. To support this, we need to explore ways for users to efficiently skim through the coverage and sequence of different clips, while also accounting for the possibility that a single clip may appear multiple times in the final edit.

VideoDiff is powered by an LLM (GPT-4o) to generate edit recommendations and to follow users' commands to generate new variations. Our pipeline is prone to transcription errors and LLM hallucinations as few study participants have pointed out. GPT-4o also has a limit for input tokens (128k) and cannot process long videos with extensive length of transcripts. Also, the current pipeline does not consider visual input when generating edit recommendations and thus cannot correctly follow edit requests related to the visual content. In the future, large multimodal models that take visual input (e.g., GPT-4V) can be integrated to improve the performance of the pipeline.

## 8 Conclusion

We presented VideoDiff, a human-AI video co-creation tool designed to support efficient video editing with alternatives. Through a functional prototype and quantitative and qualitative user feedback, we showed that users want to work with variations when they are editing videos and would like flexible tools that offer easy ways to minimize watching and maximize creative potential. VideoDiff inspired users to explore alternative ideas while retaining key edit decisions to users who can select and refine AI suggestions. All study participants wanted to integrate it into their workflow to create more satisfying videos. Generative AI tools offer new promising capabilities that will only improve over time. It is up to the HCI community to create new experiences that harness that power into better and better creative tools.

## Acknowledgments

We thank Columbia College for giving permission to use the video *Columbia University Campus Tour*. Other videos are CC licensed (*What We Ate Today + Planting Our Garden*). Part of this research is supported by Mina's Google Ph.D. fellowship.

## References

- [1] Columbia Undergraduate Admissions. 2023. Columbia University Campus Tour. <https://www.youtube.com/watch?v=OkxRkNjPHIk>
- [2] Adobe. [n. d.]. Adobe Premiere Pro. <https://www.adobe.com/products/premiere.html>
- [3] Shm Garangano Almeda, JD Zamfirescu-Pereira, Kyu Won Kim, Pradeep Mani Rathnam, and Bjoern Hartmann. 2024. Prompting for Discovery: Flexible Sense-Making for AI Art-Making with Dreamsheets. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–17.
- [4] Apple. [n. d.]. Apple iMovie. <https://support.apple.com/imovie>
- [5] Narges Ashtari, Andrea Bunt, Joanna McGrenere, Michael Nebeling, and Parmit K Chilana. 2020. Creating augmented and virtual reality applications: Current practices, challenges, and opportunities. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13.
- [6] Adam Baker, Carl Gutwin, Justin Matejka, and Ian Stavness. 2024. Interaction Techniques for Comparing Video. In *Graphics Interface 2024 Second Deadline*.
- [7] Guha Balakrishnan, Frédo Durand, and John Guttag. 2015. Video diff: Highlighting differences between similar actions in videos. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 1–10.
- [8] Connelly Barnes, Dan B Goldman, Eli Shechtman, and Adam Finkelstein. 2010. Video tapestries with continuous temporal zoom. In *ACM SIGGRAPH 2010 papers*. 1–9.
- [9] Elena Benedetto, Gabriele Romano, Ilaria Torre, Mario Vallarino, and Gianni Viardo Vercelli. 2024. A Visual Comparison interface for educational videos. In *Proceedings of the 2024 International Conference on Advanced Visual Interfaces*. 1–5.
- [10] John Boreczky, Andreas Girgensohn, Gene Golovchinsky, and Shingo Uchihashi. 2000. An interactive comic book presentation for exploring video. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 185–192.
- [11] Stephen Brade, Bryan Wang, Mauricio Sousa, Sageev Oore, and Tovi Grossman. 2023. Promptify: Text-to-image generation through interactive prompt exploration with large language models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–14.
- [12] Daniel Buschek, Martin Zürn, and Malin Eiband. 2021. The impact of multiple parallel phrase suggestions on email input and composition behaviour of native and non-native english writers. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [13] CapCut. [n. d.]. CapCut. <https://www.capcut.com/>
- [14] Minsuk Chang, Mina Huh, and Juho Kim. 2021. Rubyslippers: Supporting content-based voice navigation for how-to videos. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–14.
- [15] Shizhe Chen, Bei Liu, Jianlong Fu, Ruihua Song, Qin Jin, Pingping Lin, Xiaoyu Qi, Chunting Wang, and Jin Zhou. 2019. Neural storyboard artist: Visualizing stories with coherent image sequences. In *Proceedings of the 27th ACM International Conference on Multimedia*. 2236–2244.
- [16] Erin Cherry and Celine Latulipe. 2014. Quantifying the creativity support of digital tools through the creativity support index. *ACM Transactions on Computer-Human Interaction (TOCHI)* 21, 4 (2014), 1–25.
- [17] Peggy Chi, Tao Dong, Christian Frueh, Brian Colonna, Vivek Kwatra, and Irfan Essa. 2022. Synthesis-Assisted Video Prototyping From a Document. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–10.
- [18] Peggy Chi, Nathan Frey, Katrina Panovich, and Irfan Essa. 2021. Automatic instructional video creation from a markdown-formatted tutorial. In *The 34th Annual ACM Symposium on User Interface Software and Technology*. 677–690.
- [19] Peggy Chi, Zheng Sun, Katrina Panovich, and Irfan Essa. 2020. Automatic video creation from a web page. In *Proceedings of the 33rd annual ACM symposium on user interface software and technology*. 279–292.
- [20] Pei-Yu Chi, Joyce Liu, Jason Linder, Mira Dontcheva, Wilmot Li, and Bjoern Hartmann. 2013. Democut: generating concise instructional videos for physical demonstrations. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*. 141–150.
- [21] John Joon Young Chung, Shiqing He, and Eytan Adar. 2021. The intersection of users, roles, interactions, and technologies in creativity support tools. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*. 1817–1833.
- [22] Opus Clip. [n. d.]. Opus Clip. <https://www.opus.pro/>
- [23] Descript. [n. d.]. Descript: Edit Videos & Podcasts Like a Doc | AI Video Editor. <https://www.descript.com/>
- [24] Steven P Dow, Alana Glassco, Jonathan Kass, Melissa Schwarz, Daniel L Schwartz, and Scott R Klemmer. 2010. Parallel prototyping leads to better design results, more divergence, and increased self-efficacy. *ACM Transactions on Computer-Human Interaction (TOCHI)* 17, 4 (2010), 1–24.
- [25] Steven M Drucker, Georg Petschnigg, and Maneesh Agrawala. 2006. Comparing and managing multiple versions of slide presentations. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*. 47–56.
- [26] Fabio Duarte. [n. d.]. YouTube Content Creator Statistics (2024). <https://explodingtopics.com/blog/youtube-creator-stats>

- [27] Kasra Ferdowsi, Ruanqianqian Huang, Michael B James, Nadia Polikarpova, and Sorin Lerner. 2024. Validating AI-Generated Code with Live Programming. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–8.
- [28] C Ailie Fraser, Joy O Kim, Hijung Valentina Shin, Joel Brandt, and Mira Dontcheva. 2020. Temporal segmentation of creative live streams. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [29] Ohad Fried, Ayush Tewari, Michael Zollhöfer, Adam Finkelstein, Eli Shechtman, Dan B Goldman, Kyle Genova, Zeyu Jin, Christian Theobalt, and Maneesh Agrawala. 2019. Text-based editing of talking-head video. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–14.
- [30] Simret Araya Gebreegziabher, Yukun Yang, Elena L Glassman, and Toby Jia-Jun Li. 2024. Supporting Co-Adaptive Machine Teaching through Human Concept Learning and Cognitive Theories. *arXiv preprint arXiv:2409.16561* (2024).
- [31] Katy Ilonka Gero, Chelse Swoopes, Ziwei Gu, Jonathan K Kummerfeld, and Elena L Glassman. 2024. Supporting Sensemaking of Large Language Model Outputs at Scale. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–21.
- [32] Michael Gleicher, Danielle Albers, Rick Walker, Ilir Jusufi, Charles D Hansen, and Jonathan C Roberts. 2011. Visual comparison for information visualization. *Information Visualization* 10, 4 (2011), 289–309.
- [33] Han L Han, Miguel A Renom, Wendy E Mackay, and Michel Beaudouin-Lafon. 2020. Textlets: Supporting constraints and consistency in text documents. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [34] Sandra G Hart. 2006. NASA-task load index (NASA-TLX); 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting*, Vol. 50. Sage publications Sage CA: Los Angeles, CA, 904–908.
- [35] Bernd Huber, Hijung Valentina Shin, Bryan Russell, Oliver Wang, and Gautham J Mysore. 2019. B-script: Transcript-based b-roll video editing with recommendations. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [36] Mina Huh, Yi-Hao Peng, and Amy Pavel. 2023. GenAssist: Making image generation accessible. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–17.
- [37] Mina Huh, Saelyne Yang, Yi-Hao Peng, Xiang'Anthony' Chen, Young-Ho Kim, and Amy Pavel. 2023. AVscript: Accessible Video Editing with Audio-Visual Scripts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [38] invideo AI. [n. d.]. invideo AI Script Generator. <https://invideo.io/tools/script-generator/>
- [39] Haojian Jin, Yale Song, and Koji Yatani. 2017. Elasticplay: Interactive video summarization with dynamic time budgets. In *Proceedings of the 25th ACM international conference on Multimedia*. 1164–1172.
- [40] Joy Kim, Mira Dontcheva, Wilmot Li, Michael S Bernstein, and Daniela Steinsapir. 2015. Motif: Supporting novice creativity through expert patterns. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 1211–1220.
- [41] Janin Koch, Nicolas Taffin, Michel Beaudouin-Lafon, Markku Laine, Andrés Lucero, and Wendy E Mackay. 2020. Imagesense: An intelligent collaborative ideation tool to support diverse human-computer partnerships. *Proceedings of the ACM on human-computer interaction* 4, CSCW1 (2020), 1–27.
- [42] Steffen Koch, Markus John, Michael Wörner, Andreas Müller, and Thomas Ertl. 2014. VarifocalReader—in-depth visual analysis of large text documents. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 1723–1732.
- [43] Philippe Laban, Jesse Vig, Marti Hearst, Caiming Xiong, and Chien-Sheng Wu. 2024. Beyond the chat: Executable and verifiable text-editing with llms. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–23.
- [44] Mackenzie Leake, Abe Davis, Anh Truong, and Maneesh Agrawala. 2017. Computational video editing for dialogue-driven scenes. *ACM Trans. Graph.* 36, 4 (2017), 130–1.
- [45] Mackenzie Leake, Hijung Valentina Shin, Joy O Kim, and Maneesh Agrawala. 2020. Generating audio-visual slideshows from text articles using word concreteness. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [46] Ching Liu, Juho Kim, and Hao-Chuan Wang. 2018. ConceptScape: Collaborative concept mapping for video learning. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–12.
- [47] Michael Xieyang Liu, Tongshuang Wu, Tianying Chen, Franklin Mingzhe Li, Aniket Kittur, and Brad A Myers. 2024. Selenite: Scaffolding Online Sensemaking with Comprehensive Overviews Elicited from Large Language Models. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–26.
- [48] Vivian Liu, Tao Long, Nathan Raw, and Lydia Chilton. 2023. Generative disco: Text-to-video generation for music visualization. *arXiv preprint arXiv:2304.08551* (2023).
- [49] Ferenc Marton. 2014. *Necessary conditions of learning*. Routledge.
- [50] Justin Matejka, Michael Glueck, Erin Bradner, Ali Hashemi, Tovi Grossman, and George Fitzmaurice. 2018. Dream lens: Exploration and visualization of large-scale generative design datasets. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–12.
- [51] Justin Matejka, Tovi Grossman, and George Fitzmaurice. 2014. Video lens: rapid playback and exploration of large video collections and associated metadata. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. 541–550.
- [52] Midjourney. [n. d.]. Midjourney. <https://www.midjourney.com/>
- [53] Piotr Mirowski, Kory W Mathewson, Jaylen Pittman, and Richard Evans. 2023. Co-writing screenplays and theatre scripts with language models: Evaluation by industry professionals. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–34.
- [54] Zheng Ning, Brianna L Wimer, Kaiwen Jiang, Keyi Chen, Jerrick Ban, Yapeng Tian, Yuhang Zhao, and Toby Jia-Jun Li. 2024. SPICA: Interactive Video Content Exploration through Augmented Audio Descriptions for Blind or Low-Vision Viewers. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–18.
- [55] OpenAI. [n. d.]. DALL-E 3. <https://openai.com/index/dall-e-3/>
- [56] OpenAI. 2024. Sora OpenAI video generation model. <https://openai.com/sora>
- [57] Remotion.de. 2024 (accessed Sept 11, 2024). Remotion: Make videos programmatically in react. <https://www.remotion.dev/>
- [58] Mohi Reza, Nathan Laundry, Ilya Musabirov, Peter Dushniku, Zhi Yuan Yu, Kashish Mittal, Tovi Grossman, Michael Liut, Anastasia Kuzminykh, Joseph Jay Williams, et al. 2023. ABScribe: Rapid Exploration of Multiple Writing Variations in Human-AI Co-Writing Tasks using Large Language Models. *arXiv preprint arXiv:2310.00117* (2023).
- [59] Brittany Rose. 2024. Huge ALDI Grocery Haul | Shop with me and haul!! <https://www.youtube.com/watch?v=ti43mLvsug>
- [60] Brittany Rose. 2024. WEEKLY GROCERY HAUL AND MEAL PLAN | Shop with me!! <https://www.youtube.com/watch?v=IxPofzdXxM>
- [61] Steve Rubin and Maneesh Agrawala. 2014. Generating emotionally relevant musical scores for audio stories. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. 439–448.
- [62] Sangho Suh, Meng Chen, Bryan Min, Toby Jia-Jun Li, and Haijun Xia. 2024. Luminate: Structured Generation and Exploration of Design Space with Large Language Models for Human-AI Co-Creation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–26.
- [63] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. 2023. Sensecape: Enabling multilevel exploration and sensemaking with large language models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–18.
- [64] Amanda Swearngin, Chenglong Wang, Alannah Oleson, James Fogarty, and Amy J Ko. 2020. Scout: Rapid exploration of interface layout alternatives through high-level design constraints. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13.
- [65] TED. [n. d.]. Does Working Hard Really Make You a Good Person? | Azim Shariff | TED. [https://www.youtube.com/watch?v=z\\_ca983qJcI](https://www.youtube.com/watch?v=z_ca983qJcI)
- [66] Atima Tharatipyakul and Hyowon Lee. 2018. Towards a better video comparison: comparison as a way of browsing the video contents. In *Proceedings of the 30th Australian Conference on Computer-Human Interaction*. 349–353.
- [67] 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*. 515–536.
- [68] Anh Truong and Maneesh Agrawala. 2019. A Tool for Navigating and Editing 360 Video of Social Conversations into Shareable Highlights.. In *Graphics Interface*. 14–1.
- [69] Anh Truong, Floraine Berthouzou, Wilmot Li, and Maneesh Agrawala. 2016. Quickcut: An interactive tool for editing narrated video. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 497–507.
- [70] Babish Culinary Universe. [n. d.]. Eggs Benedict | Basics with Babish Live. <https://www.youtube.com/watch?v=ROAmMJeTO6Y>
- [71] Vimeo. [n. d.]. Vimeo AI-Powered Video Platform. <https://vimeo.com/>
- [72] Vizard. [n. d.]. Vizard. <https://vizard.ai/>
- [73] Samangi Wadinambarachchi, Ryan M Kelly, Saumya Pareek, Qiushi Zhou, and Eduardo Velloso. 2024. The Effects of Generative AI on Design Fixation and Divergent Thinking. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–18.
- [74] Qian Wan, Siying Hu, Yu Zhang, Piaohong Wang, Bo Wen, and Zhicong Lu. 2024. "It Felt Like Having a Second Mind": Investigating Human-AI Co-creativity in Prewriting with Large Language Models. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–26.
- [75] Bryan Wang, Zeyu Jin, and Gautham Mysore. 2022. Record Once, Post Everywhere: Automatic Shortening of Audio Stories for Social Media. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–11.



- [76] Bryan Wang, Yuliang Li, Zhaoyang Lv, Haijun Xia, Yan Xu, and Raj Sodhi. 2024. LAVE: LLM-Powered Agent Assistance and Language Augmentation for Video Editing. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*. 699–714.
- [77] Miao Wang, Guo-Wei Yang, Shi-Min Hu, Shing-Tung Yau, Ariel Shamir, et al. 2019. Write-a-video: computational video montage from themed text. *ACM Trans. Graph.* 38, 6 (2019), 177–1.
- [78] Sitong Wang, Samia Menon, Tao Long, Keren Henderson, Dingzeyu Li, Kevin Crowston, Mark Hansen, Jeffrey V Nickerson, and Lydia B Chilton. 2024. Reel-Framer: Human-AI co-creation for news-to-video translation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–20.
- [79] Sitong Wang, Zheng Ning, Anh Truong, Mira Dontcheva, Dingzeyu Li, and Lydia B Chilton. 2024. PodReels: Human-AI Co-Creation of Video Podcast Teasers. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 958–974.
- [80] Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, et al. 2022. Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 214–229.
- [81] Tongshuang Wu, Michael Terry, and Carrie Jun Cai. 2022. Ai chains: Transparent and controllable human-ai interaction by chaining large language model prompts. In *Proceedings of the 2022 CHI conference on human factors in computing systems*. 1–22.
- [82] Haijun Xia, Jennifer Jacobs, and Maneesh Agrawala. 2020. Crosscast: adding visuals to audio travel podcasts. In *Proceedings of the 33rd annual ACM symposium on user interface software and technology*. 735–746.
- [83] Liwenhan Xie, Zhaoyu Zhou, Kerun Yu, Yun Wang, Huamin Qu, and Siming Chen. 2023. Wakey-Wakey: Animate Text by Mimicking Characters in a GIF. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–14.
- [84] Saelyne Yang, Jisu Yim, Aitolkyn Baigutanova, Seoyoung Kim, Minsuk Chang, and Juho Kim. 2022. SoftVideo: Improving the Learning Experience of Software Tutorial Videos with Collective Interaction Data. In *Proceedings of the 27th International Conference on Intelligent User Interfaces*. 646–660.
- [85] Catherine Yeh, Gonzalo Ramos, Rachel Ng, Andy Huntington, and Richard Banks. 2024. GhostWriter: Augmenting Collaborative Human-AI Writing Experiences Through Personalization and Agency. *arXiv preprint arXiv:2402.08855* (2024).
- [86] Zheng Zhang, Zheng Ning, Chenliang Xu, Yapeng Tian, and Toby Jia-Jun Li. 2023. PEANUT: A Human-AI Collaborative Tool for Annotating Audio-Visual Data. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–18.

## A PARTICIPANTS

PID	Gender	Age	Job	Experience	Videos Created
1	Male	37	Video editor	Expert (15 yrs)	Commercials, TV shows, Targetted Ads
2	Female	35	Video editor	Expert (14 yrs)	Commercials, Trailers, Wedding/Birthday videos
3	Female	29	Video editor	Expert (10 yrs)	Vlogs, Documentaries, Wedding videos
4	Male	26	Studio director	Expert (6 yrs)	Interviews, Training videos, Political Ads
5	Male	29	Video editor	Expert (10 yrs)	Real estate videos, Company Brand/Training videos
6	Male	20	Video editor	Expert (8 yrs)	Documentaries, Social media clips, Sports highlight
7	Male	28	Videographer	Expert (13 yrs)	Music videos, Documentaries, Ads, Short-forms
8	Male	32	Video editor, 2D Animator	Expert (8 yrs)	Promotion videos, Explainer/Tutorial videos, Celebration videos
9	Male	22	Video editor	Expert (4yrs)	Product reviews, Vlogs
10	Female	29	Video editor	Expert (8yrs)	Commercials, Educational videos
11	Male	37	Motion graphics designer	Expert (10yrs)	Corporate videos
12	Male	33	Videographer	Expert (10yrs)	Documentaries, Social media reels
13	Male	25	Motion graphics designer	Expert (6yrs)	Promotional videos, Ads
14	Male	35	Motion graphics designer	Expert (12yrs)	Corporate training videos, Presentations
15	Male	18	Film school student	Amateur (3yrs)	Documentaries
16	Female	25	3D animator	Amateur (3yrs)	Music videos
17	Female	28	Product manager	Amateur (3yrs)	Personal videos
18	Male	24	Software engineer	Amateur (4yrs)	Personal videos
19	Male	33	Graduate Student	Amateur (6yrs)	Product demo videos
20	Male	27	Software engineer	Amateur (3yrs)	School projects
21	Male	18	Student	Amateur (4yrs)	Personal videos, Explainers
22	Male	43	Video Editor	Expert (5yrs)	Music/art videos, documentaries

**Table 4: Demographics of study participants (Formative study: P1-P8, Controlled user study: P9-P20, Exploratory case study: P9, P21-P22)**

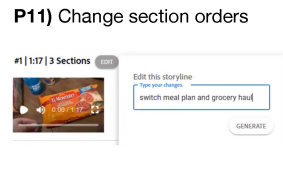
## B Augmentation Prompts

Edit Stage	Attribute	Options
Rough Cut	Length	- Make the total duration to be less than 2 minutes. - Make the total duration to be 2-4 minutes. - Make the total duration to be 4-5 minutes.
	Focus	- Focus on 1-2 sections for a coherent summary, minimize additional sections. - Focus on 3-4 sections, minimize additional sections. - Extract short parts from many sections for diversity.
B-Roll	Media	- Illustration - Photo - Video
	Frequency	- Limit B-Rolls to 2-3. - Limit B-Rolls to 4-5. - Limit B-Rolls to 7-10.
Text Effects	Style	- Lower - Title - Subtitle
	Frequency	- Limit pull quote sentences to 2-3. - Limit pull quote sentences to 4-5. - Limit pull quote sentences to 7-10.

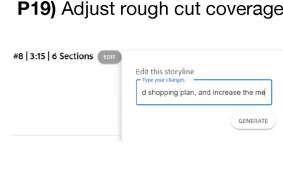
Table 5: Augmentation prompts to generate diverse prompts for rough cuts, B-rolls, and text effects. For B-rolls and text effects, the media type is used for searching stock API and embedding text styles using Remotion.

## C User Study Examples

**P11) Change section orders**



**P19) Adjust rough cut coverage**



**P20) Remove B-roll**

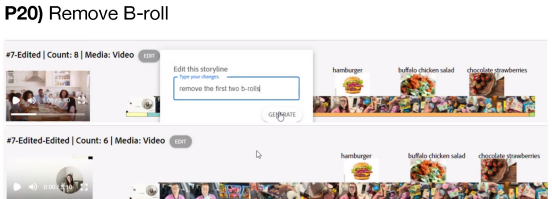


Figure 11: Examples of user edits with VideoDiff