SlideSpecs: Automatic and Interactive Presentation Feedback Collation

JEREMY WARNER, University of California, Berkeley, USA
AMY PAVEL, University of Texas, Austin, USA
TONYA NGUYEN, University of California, Berkeley, USA
MANEESH AGRAWALA, Stanford University, USA
BJÖRN HARTMANN, University of California, Berkeley, USA

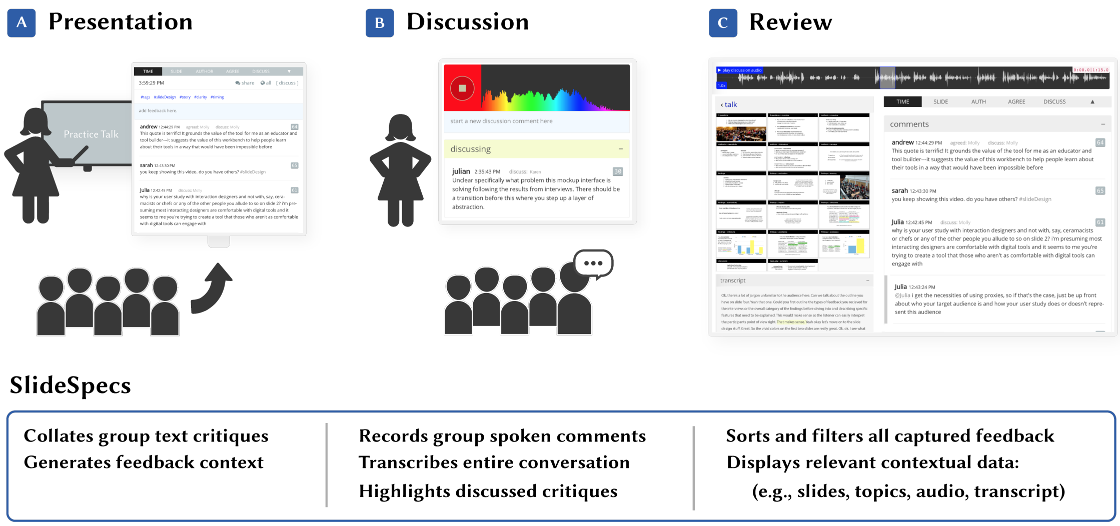

Fig. 1. SlideSpecs supports presenters in three phases: **Presentation**, **Discussion**, and **Review**. During the Presentation, SlideSpecs collates audience text critiques referencing talk slides, feedback tags, or other critiques. During the Discussion, the interface records and transcribes the discussion audio. Participants can easily reference previous critiques, and a facilitator annotates discussion topics. During Review, SlideSpecs provides presenters with rich contextual data that links the feedback gathered from both previous phases.

Presenters often collect audience feedback through practice talks to refine their presentations. In formative interviews, we find that although text feedback and verbal discussions allow presenters to receive feedback, organizing that feedback into actionable presentation revisions remains challenging. Feedback may lack context, be redundant, and be spread across various emails, notes, and conversations. To collate and contextualize both text and verbal feedback, we present *SlideSpecs*. SlideSpecs lets audience members provide text feedback (e.g., 'font too small') while attaching an automatically detected context, including relevant slides (e.g., 'Slide 7') or content tags (e.g., 'slide design'). SlideSpecs also records and transcribes spoken group discussions that commonly occur after practice talks and facilitates linking text critiques to relevant discussion segments. Finally, presenters can use SlideSpecs to review all text and spoken feedback in a single contextually rich interface (e.g., relevant slides, topics, and follow-up discussions). We demonstrate the effectiveness of SlideSpecs by deploying it in eight practice talks with a range of topics and purposes and reporting our findings.

Manuscript submitted to ACM

CCS Concepts: • **Human-centered computing → Interactive systems and tools**; **Collaborative interaction**.

Additional Key Words and Phrases: interactive systems; feedback tools; image matching; context generation; slide presentations

ACM Reference Format:
Jeremy Warner, Amy Pavel, Tonya Nguyen, Maneesh Agrawala, and Björn Hartmann. 2023. SlideSpecs: Automatic and Interactive Presentation Feedback Collation. In *IUI 2023: 28th ACM Conference on Intelligent User Interfaces, March 27–31, 2023, Sydney, Australia.* ACM, New York, NY, USA, 22 pages.

1 INTRODUCTION

Presentations are a foundational way of sharing information with others across education, business, science, and government. Unfortunately, there are also abundant examples of ineffective or even misleading talks. Assuming presenters want to give the best version of their message and communicate as clearly as possible – what goes wrong?

Preparing an effective presentation takes time and work. The more a skilled presenter makes it look easy, the more work and talk revisions they've likely done beforehand. When you consider the potential broader impact of a talk, this work is justified. A 20-minute talk for a 60-person audience (e.g., a reasonable conference talk) consumes 20 hours of audience time. Talk videos can reach an even broader – a 600-view video of that same talk takes 200 hours of viewer time [4]. When framed this way, the time and work spent refining presentations are more clearly justified.

To improve the effectiveness of slide presentations, presenters often give practice talks to live audiences and receive presentation critiques. Presenters may receive feedback from audience members through discussion after the presentation and written text critiques (e.g., shared notebook, email). To leverage audience feedback, presenters must record and distill critiques from an open-ended discussion, recall critiques and their corresponding contexts (e.g., relevant slides, audience-suggested solutions), and organize comments across multiple authors and mediums. Each task is challenging and potentially error-prone, especially given a limited amount of time for discussion and clarification. Presenters can quickly lose track of relevant feedback, leaving potentially valuable audience critiques unaddressed.

To understand the presenters' challenges when refining talks, we conducted formative interviews with 14 participants. Received feedback ranged in *scope* (e.g., specific to a single slide vs. general presenting tips) and *subject* (e.g., slide design, narrative). Spoken feedback during and after the practice talk provided opportunities for discussion and clarification; however, bandwidth was limited (e.g., how many issues can be raised), individual discussions may not make the best use of the entire audience's time, and public discussion could bias or inhibit other audience feedback. Alternatively, audience members sometimes shared written text feedback with the presenter either privately (e.g., an email, index cards) or publicly (e.g., a shared online document). Text feedback provides a written record of critiques and offers more space to share concerns. Both verbal and written feedback benefit these presenters, but capturing critique context and organizing feedback across sources into a more valuable and accessible form remains challenging.

To support the automatic and interactive collation of audience feedback, we present SlideSpecs (Figure 1). SlideSpecs supports three common phases in the presentation revision process: (A) *Presentation*, (B) *Discussion*, and (C) *Review*.

In (A), the *Presentation*, the presenter uses their preferred slide software while the audience uses the SlideSpecs *feedback-providing interface* (Figure 2). This interface shows the talk slides, the speaker's current position within the slide deck, and the other audience critiques (optionally). Audience members can write critiques, include relevant *scope*, provide their critique's *subject*, reply to other comments, and flag other critiques for agreement or further discussion.

In (B), the post-talk *Discussion*, the presenter and audience can use the *discussion interface* to guide and capture conversations (Figure 4). SlideSpecs records and transcribes the conversation in this phase, including elaborating on existing comments and new critiques the audience may raise. A talk facilitator can dynamically associate each spoken audience comment with an existing text comment or designate the comment as a new point of critique (Figure 5).

In (C), the feedback *Review*, presenters access a contextualized record of all critiques from (A) and (B) in the *reviewing interface*. This interface shows transcribed discussion segments and displays these transcripts alongside relevant linked text critiques and discussion topics. Presenters can also sort, filter, and search all feedback with keywords, slide numbers, tags, and agree/discuss votes. The automatically generated context, transcribed discussion segments, and audience-contributed tags help presenters better understand their critique and revise their talk accordingly.

SlideSpecs is the first system to collate and contextualize synchronously generated text and spoken presentation critiques in a unified interface. SlideSpecs collects and contextualizes text and verbal critiques across both practice presentations and group discussions. SlideSpecs uses a live group-chat style interface with lightweight tagging for feedback to support elaboration on other audience member critiques (Figure 2). SlideSpecs also provides a flexible set of filtering and sorting tools that presenters can apply to feedback across the entire process of practicing their talks.

To assess how using SlideSpecs affects presenters' ability to leverage critiques for talk revision, we evaluate SlideSpecs across the presentation, discussion, and review phases (Figure 3). We demonstrate the effectiveness of SlideSpecs by deploying it in eight practice presentations (Table 1) and reporting our findings. We also reflect on several future benefits that further automation and text summarization may provide presenters (Sec. 9).

Our contributions include:

- Design implications for group slide-feedback interfaces derived from formative interviews
- SlideSpecs, a novel system for collating audience text and spoken presentation feedback
- A novel screenshot-based slide-detecting technique for automatic feedback contextualization
- An evaluation applying SlideSpecs to eight talks across different research groups and topics

2 RELATED WORK

Our work draws on prior research ranging from systems that support creating and improving presentations, generating context for meetings and discussions, tools for gathering and reviewing feedback, and text summarization techniques.

2.1 Supporting Better Presentations

Past works have addressed providing synchronous, textual feedback written during in-class presentations [50] through a forum-style interface and critique during design reviews [19, 23, 38, 39], though they do not capture the post-presentation discussion, nor does it contextualize feedback in the presentation. While existing online slideware (e.g., Google Slides) may allow multiple users to comment on talk slides simultaneously, these platforms do not facilitate references between presentation feedback and post-presentation discussion. This research diverges from prior work by focusing on the domain of practice presentations to small groups, where long in-person discussions after the talks are common.

Prior work considered tools to improve the production and presentation of slides primarily by reducing the rigidity of linear slides, providing additional context for slide material, or slide generation tools [2, 15, 26, 28, 29, 49, 53]. Such slideware lies on a spectrum from traditional slide presentation tools (e.g., Google Slides, PowerPoint, Keynote), which provide a linear sequence of rectangular slides, to IdeaMaché [29], which allows freeform presentations in the form of a navigable collage. In between lie tools like Fly [28], NextSlidePlease [53], and CounterPoint [15], which each allow the user to arrange slides in a 2D space and enable linear or branching paths through the space during presentations (to allow presentations to change dynamically based on time or audience feedback). Drucker et al. demonstrate an effective way

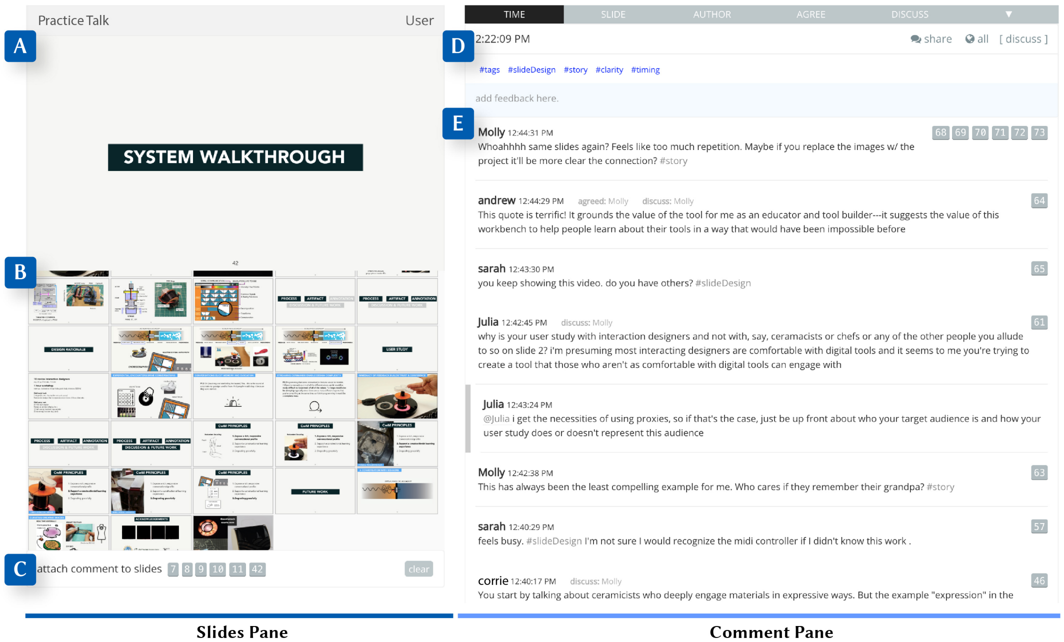

Fig. 2. An overview of the SlideSpecs *feedback-providing* interface. In the *Slides Pane*, a detailed slide image view is shown (A), and hovering over the smaller slide thumbnails (B) temporarily updates the large image, current slide (A), allowing users to scan over talk slides. Audience members can select these thumbnails to attach their comments to designated presentation slides (C). In the *Comments Pane*, the interface features different ways of sorting the comments (D), a list of tags contained by comments, and a text area entering feedback. Last, the interface displays a live-updating list of the audience's comments (E), which includes comment metadata (i.e., comment author, creation time, and slides referenced).

to track presentation changes and refinements over time [11]. Other tools focus specifically on the presenter's oration rather than general feedback [1, 37]. To increase feedback context, DynamicSlide [20] infers links across presentation videos and slides while others [42, 43] link slide text and images to presenter speech and promote accessibility. As SlideSpecs lets audiences provide feedback on uninterrupted rehearsal presentations, SlideSpecs supports any linear presentation. It also lets audiences peek forward or look back at previous slides to craft general presentation feedback.

2.2 Contextualizing Meetings and Discussions

As SlideSpecs helps presenters record and organize spoken critiques produced during a discussion phase, we build on prior work in supporting meeting discussions. Producing a reusable record of meeting events remains a challenging task. Early work found that meeting note-takers faced difficulties including dropping crucial facts, using uninterpretable abbreviations, finding too little time to write notes, and missing participation opportunities while note-taking [22, 58]. We also find it challenging for presenters to write down critiques during presentation feedback sessions while listening and responding to the conversation. To improve meeting recall despite spotty notes, prior work recorded the meeting and then linked handwritten [6, 58] or typed meeting notes [5] taken during the meeting to their media timestamps. Other work indexed meetings by the speaker-segmented transcript [57], by domain-specific artifacts [9, 13, 32] referenced during the meeting (e.g., slides [32]), or by automatically extracted important moments [27, 33, 35].

We investigate how to let the presenter focus on organizing text-based critiques, using the meeting transcript primarily to provide more fidelity for written critiques. To enable this, we investigate how to help the facilitator efficiently establish links between discussion segments and relevant existing text critiques during the meeting. Unlike following an agenda with few items [13] or writing notes directly [5, 6, 58], locating specific text critiques requires the facilitator to search for particular comments and the audience to promote specific comments for discussion.

Our work relates to prior work on using group discussion chats during a presentation and organizing chats for later review. Prior work on digital backchannels [31] during presentations involves either forum-style promotion of questions and feedback [17, 46, 50] or real-time chat [45, 47]. In contrast to prior work that investigated general real-time chat during a presentation, we explore the use of real-time chat to collectively generate slide critiques where threading can be used for critique elaboration rather than general chat. To use real-time comments for later review, we allow lightweight markup of text chat similar to Tilda [64, 65]. Instead of focusing on general meeting markup, we explore how to design lightweight tagging of comments specific to critiques (e.g., location, critique type).

2.3 Feedback Review Support

Many systems exist for recording and reviewing critique in asynchronous and synchronous scenarios for different media (e.g., single page graphic designs [30], PDFs [62]) and audiences (e.g., crowds [30], peers [24, 50], teams [39, 41]). For instance, prior work has addressed capturing and organizing critiques [30, 39, 41, 50, 60] primarily in asynchronous settings where all communication is mediated through the critique interface. However, one unique aspect of practice talks is the synchronous setting in which all participants see the talk simultaneously and then discuss feedback post-presentation. Past work has studied in-person collaboration on creative tasks [56], the use of Google Docs for providing feedback in classes [50], and the use of Google Docs compared to in-person meetings for design collaboration [21].

Using SlideSpecs, audience members write critiques using real-time messaging on a shared chat channel, promoting specific critiques and threaded conversations. Researchers have also suggested techniques for improving peer and novice feedback, including introducing rubrics [24, 30, 55], computing reviewer accuracy [7], and providing live suggestions on the feedback [12]. While these prior works focus on supporting critique from peers and providing guidance on how to give better critique, our research emphasizes supporting critique in a group setting where participants typically have more experience and motivation. The work that has focused on helping experts provide feedback in asynchronous contexts, both in small teams [41, 62, 63] and classrooms [14], has not addressed the larger question of how to organize feedback from multiple critiques across multiple modes of media. We aim to address this open question: how can we effectively support synchronous critique from an experienced audience in a group setting across multiple media modes?

2.4 Text Summarization

One core issue when receiving feedback is the sheer amount of feedback. Automatic text summarization performance continues to improve, enabling new possible applications [51, 52]. Text summarization techniques have been applied across many domains (outside of meetings) to help the reader, including news [25, 61], sports [36], food reviews [8], auctions [18] and email [10]. Efficient text summarization can also be applied to real-world physical documents for accessibility needs [3]. SlideSpecs provides the audience with shared context to help reduce redundant comments.

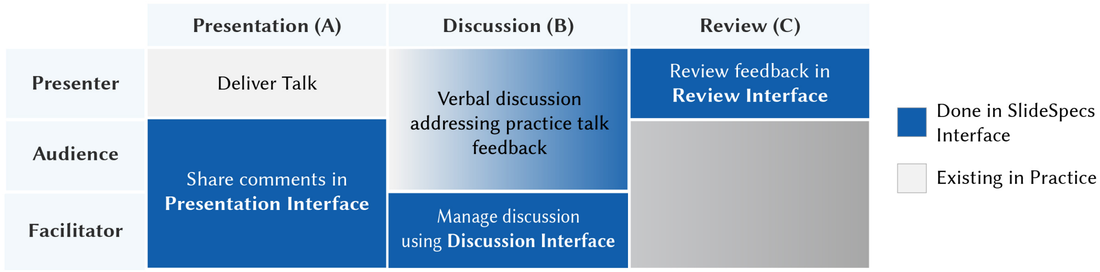

Fig. 3. SlideSpecs supports three phases (presentation, discussion, and review) and three distinct participant roles (presenter, audience, and facilitator). The presenter delivers their practice presentation while audience members and the facilitator share comments using the *Presentation Interface* (A). During the post-presentation discussion (B), the presenter and audience discuss practice talk feedback while the facilitator annotates the discussion. The presenter later reviews all feedback at their leisure in the *Review Interface* (C).

3 FORMATIVE STUDY: EXISTING PRACTICES

3.1 Method

To learn more about existing practices for giving and receiving feedback on presentations, we analyzed existing guides for creating presentations and interviewed 14 presenters (N1-14). We found these presenters via university mailing lists that connect a wide range of domains. Our participant presenters represented a range of fields (e.g., AI, HCI, Optimization and Testing, Design, Student Government) and eventual talk venues (e.g., company presentations, conferences, workshops, job talks). All the presenters we talked to also provided feedback to others on practice talks in the past. Specifically, we asked presenters a series of questions to find out: *What are the benefits and drawbacks of methods currently used to give and receive feedback on practice presentations?* We recorded the audio of these formative interviews. We then categorized and reviewed participant answers into *role* (e.g., audience member, presenter), *phase* (e.g., feedback, discussion, revision), and *type* (text vs. spoken) before summarizing the categorized notes in our findings.

3.2 Findings

3.2.1 Presentation Feedback Sources. Participants reported using methods for giving and receiving feedback that fell into one or more of the following categories: (1) text feedback on the slide deck alone, (2) personal or shared text feedback during the practice talk (e.g., personal notes, email, Google Docs, shared whiteboard), and (3) spoken feedback at the end of the practice talk (e.g., in-person or via video conferencing). Two participants mentioned pre-talk slide deck feedback (slide design and story critique only); all but one mentioned in-session text feedback (6 using shared text feedback); all participants received spoken feedback. Most participants reported gathering feedback across all possible scopes, though one only sought feedback focused on the work itself (rather than any presentation-specific feedback).

3.2.2 Text Feedback. Participants mentioned that more audience members contributed to a larger quantity and diversity of feedback when using a shared text medium (e.g., Google Docs or shared whiteboard) or personal text notes (corroborating Jung et al. [21]). Reflecting on the quantity and content of text feedback, N3 noted that such feedback is 'more critical, can be a bit meaner, but that's important and even better since it's being given to help'. Participants also mentioned preferring text feedback, particularly for specific lower-level critiques, as major feedback might require discussion: 'Higher-level structural feedback wasn't useful in comments because I couldn't ask them what they meant. (N2)', the discussion gives presenters a 'better sense of what the audience is confused about (N8)'.

Participants reported that when preparing personal notes to deliver critiques, they take rough notes during the presentation, refining the notes (e.g., rephrasing, ideating solutions, or removing comments brought up in discussion) before sending them to the presenter afterward. For instance, N7 mentioned that her initial text 'notes are harsh and blunt and have to be reframed' before sending them to the presenter for their review. On the other hand, when crafting shared notes, some participants focused all their time on the initial comment but did not revisit comments to revise before sharing them, reporting it 'feels like the responsibility of the presenter at that point (N14)'.

In addition, participants noted that in providing shared feedback, they still regulated their text feedback in the presence of other group members either to avoid the judgment of senior members (evaluation apprehension [56]) or to avoid duplicate work: 'I had the feeling that everyone was doing the same thing but better (N12)', 'the advisor knows better (N2)'. In some cases, participants found the shared feedback to be distracting or discouraging or spent time reading senior member feedback purposefully to learn how to provide feedback: 'First feedback session, didn't write a single thing, but I read a lot of the feedback because I don't know how to give feedback on this type of work'...'everyone is having all of these insights, but I'm not. (N11)'. Group dynamics influence the feedback process.

3.2.3 Verbal Feedback. Participants mentioned finding that verbal feedback provides the most opportunity for discussion: 'way more effective than Google Docs back and forth (N1)'. Still, it was very time-limited as only one speaker could voice ideas at a time (production blocking [56]). Participants emphasized that especially in time-limited scenarios, the most senior participants spoke most: 'the feedback process is almost like, oh the most senior person or the most opinionated person gives a high-level overview. Everyone else echoes and then adds on. (N5)' With limited time, and senior participants speaking, participants cited that most spoken feedback consisted of high-level or delivery feedback: 'More high-level feedback with structure (N3)', with the exception of lengthy in-person slide walk-through sessions or interrupting questions: 'in my research group, there's a big culture of asking questions mid-talk (N6)'. When in the role of giving feedback, participants mentioned that they considered several issues when considering whether to provide spoken feedback: how other participants may perceive their feedback: 'you want to be astute about it (N13)', the importance of their feedback to improving the talk: 'Structure, flow, there has to be a story that's going on (N9)', their familiarity with the subject, number of times speaking: 'I don't want to monopolize the conversation, once I say my two things (N13)', and the ease of communicating a critique (e.g., whether the relevant slide appeared on screen).

3.2.4 Revision Process. When recording notes, both in the case of providing and receiving feedback, participants mentioned that they sometimes forget the relevant context of notes (e.g., what slide, what content, which audience member). But, presenters found such context important when revising the presentation. When addressing feedback, participants suggested prioritizing based on several factors: (1) importance or trustworthiness of feedback author, (2) ease of fixing the problem (e.g., typos, font choice), (3) importance of the problem to overall argument, (4) Slide order (e.g., start with Slide 1), and (5) ease of remembering suggested problems (e.g., fix the problems that you remember first, and then go back to document to fill in the gaps). Most presenters formally organized their comments before revision: 'I make a to-do list after I synthesize all comments (N8)', 'a to-do list and cross out items that are finished (N7)', 'organized by a PM on google docs (N1)'. More experienced presenters expressed more selectivity in addressing comments: 'I just think about the notes I was given and try to catch the issues when I practice (N3)'. For those that organized their received feedback, the messiness (e.g., their own quickly handwritten or typed discussion notes or the group's shared notes) and source disparity (e.g., spread across emails) of existing notes was a common complaint.

3.3 Classifying Critiques

In addition to our formative interviews, we also analyzed several existing guides for creating and giving effective presentations [4, 34, 48, 54]. We use each guideline from these guides as a possible source of critique. Two defining dimensions for critiques emerged from analyzing these guides: *scope* and *subject*.

3.3.1 Scope. A critique typically references one of four types of locations in a slide presentation: a single slide (e.g., 'The results slide contains too much text'), a contiguous set of slides (e.g., 'the motivation section could be shortened'), a non-contiguous set of slides (e.g., 'switch the position of these two slides'), or the entire talk (e.g., 'pause more often'). We refer to critiques about any subset of slides as *local* and critiques about the presentation as a whole to be *global*.

3.3.2 Subject. Presentation feedback falls into four main subject categories: *content, slide design, story*, and *delivery*. *Content* comments address the core of the presented work (e.g., 'consider interviewing the people that used your system'). Such comments may express ideas or concerns about the work, but they also may reflect a misunderstanding or need for clarification within the presentation. *Slide design* addresses the visual layout of slide content, the amount of slide content, and consistency across slides. *Story* covers presentation structure, supporting examples, argument, and audience-presentation-fit. *Delivery* addresses aspects of giving the talk, including the script, the speed of narration, and gestures or body positioning. A comment in any of these categories may be positive (e.g., 'great use of color!') or critical (e.g., 'I don't understand this example'). These critique category labels are not mutually exclusive.

4 DESIGN IMPLICATIONS

We discuss design implications determined from prior work and our formative work, then discuss our approach to addressing these implications with SlideSpecs. Based on our analysis of presentation guides and formative interviews, we present five key implications for designs in the space of recording and organizing presentation feedback:

Support existing workflows for group feedback. Formative interviews reveal that groups give feedback at different times (e.g., before, during, or after the presentation), using different media (e.g., text, spoken, digital, written), and with varying levels of privacy (e.g., public to group/private). Preserving a group's social processes is advantageous in new technology adoption (Grudin's third principle [16]), though this can drastically limit the space of acceptable designs. A balance should be struck between adapting to existing structured processes and effectively collecting and organizing feedback from diverse inputs.

Support a variety of critique contexts. The context of feedback (e.g., authorship, critique type, slide) aids critique comprehension and prioritization during revision. Presenters and presentation guidelines reveal several important contextual features, and presenters will value these features differently.

Mitigate audience attention demands. Audience members already manage multiple challenging tasks when providing feedback, including comprehending talk content, identifying critical issues, composing feedback, and managing social expectations. While audience members may be best suited to provide additional information about their critiques, a challenge exists in balancing any extra effort with present cognitive demands, especially as audience members are not always the primary beneficiary of the work [16].

Organize presentation and discussion feedback into action items. Most presenters use a distilled list of action items to revise their presentation. However, they generally receive feedback via blocks of raw text and an open-ended, non-linear discussion of presentation issues. A design should support transforming these disparate inputs (e.g., text critiques and non-linear discussion) into usable, distinct action items.

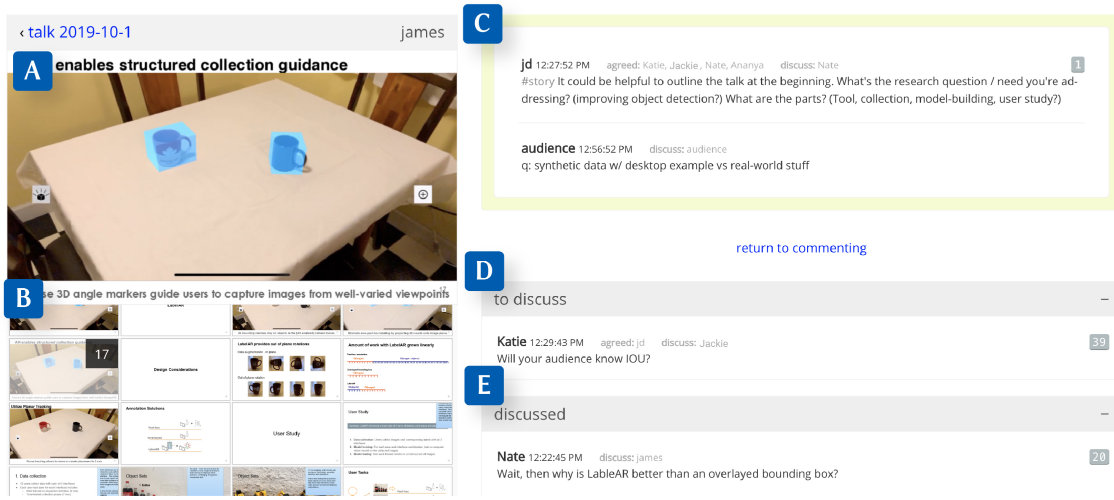

Fig. 4. The *Discussion Participant* interface. (A) The slide pane detail view. (B) The listing of slide thumbnails, which also allows discussion participants to tag and filter comments by a specific slide. (C) Comments currently under discussion. These are synced with the comments shown in Fig. 5B. (D) Comments queued for discussion, initially populated by comments that the audience marked for discussion. (E) Comments already discussed. Once the facilitator marks these comments as discussed, they move into this section.

Reduce note-taking demands. Currently, presenters mainly capture and organize their critiques from discussions by taking notes. However, note-taking is demanding and error-prone: a well-documented challenge due to the high cognitive load of note-taking [40, 44, 59]. As the presenter frequently participates in post-presentation discussions, they are suited particularly poorly to take high-quality, contextually vivid discussion notes.

5 COLLATING FEEDBACK WITH SLIDESPECS

With these design implications in mind, we developed **SlideSpecs**. SlideSpecs supports collating presentation feedback across three observed practice talk phases: the *presentation*, the post-presentation *discussion* and the feedback *review* (Figure 3). First, the SlideSpecs *presentation interface* records and contextualizes audience text critiques and facilitates audience collaboration with a lightweight tagging scheme and reply mechanism (Figure 2). The post-presentation *discussion interface* surfaces relevant text critiques to inform verbal discussions (Figure 4), and helps a facilitator to link discussion segments to related critiques (Figure 5). The *review interface* then lets presenters all collated feedback alongside the corresponding context (e.g., slides, discussion, topic) and a transcript of linked discussion segments (Figure 6).

5.1 Presentation Phase

SlideSpecs lets presenters upload a PDF of their presentation slides, and the system generates a thumbnail image for each slide in the presentation. To allow the audience to provide and organize text feedback (e.g., by author, by slide) during the presentation, SlideSpecs features a feedback-providing interface (Figure 2). It features two main panes: the *Slides Pane* that lets audience members view and select relevant slides for critique and the *Comments Pane* that lets audience members provide their comments and view and interact with other audience comments.

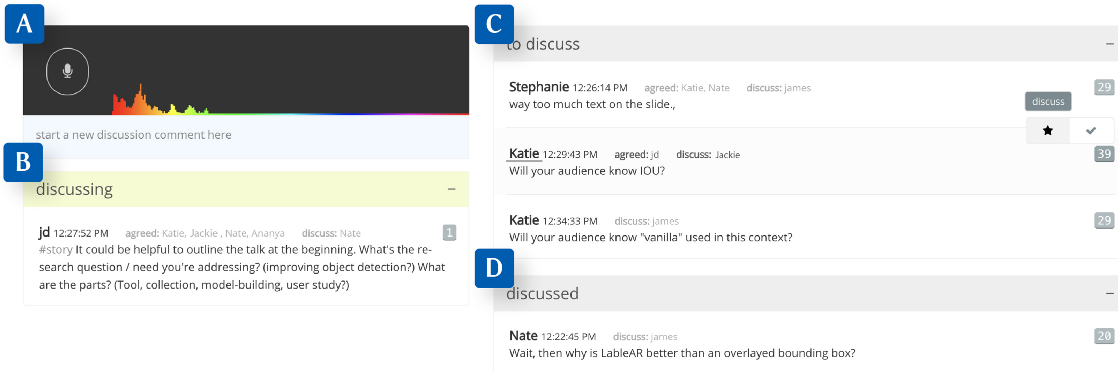

Fig. 5. The *Discussion Facilitator* interface. (A) A microphone button to toggle audio capture and a visualizer for the audio input to verify that the discussion audio is active. The text entry below the audio controls can submit new discussion comments and works as a search tool to find relevant comments. (B) The "discussing" pane shows comments actively being discussed. These comments sync with the discussion participant view. (C) Comments queued up for discussion. (D) Previously discussed comments.

5.1.1 Slides Pane. Audience members can view and attach feedback to presentation slides with the slide pane (Figure 2a, b, c). By hovering over a slide, audience members may view each slide in greater detail. Audience members can also select a set of slides to attach their comments to reference (Figure 2a, b, c). In this case, the audience member selects a range of slides (e.g., 'slides 7-11 & 42') for their feedback to reference (Figure 2b, c). If a comment references specific slides, the slide numbers will appear alongside the comment in the comment pane. By default, the slides pane displays the current slide, or the slide most recently matched to a presenter screenshot (Figure 2a)).

5.1.2 Comment Pane. The comment pane allows audience members to share their feedback within a comment field (Figure 2d), marked "add feedback here." SlideSpecs detects the currently presented slide (details in Section 6.1), allowing the audience to contextualize their feedback automatically. After sharing a comment, it appears in the comment list (Figure 2e) along with the author's name (Molly), the time the comment was shared (12:44 pm), and any referenced slides (e.g. 'slides 68-73'). To attribute an explicit category to the comment, the audience can include tags in the comment's body (e.g., '#slidedesign, #delivery, & #story'). Audience members can interact with shared comments by writing a reply, agreeing with the comment, and flagging a comment for discussion. The comments pane also features a *focus* mode, which only shows self-authored comments.

5.2 Discussion Phase

After the presenter has finished going through their talk, the discussion interface supports the process of recording and linking the group discussion through two views: the *participant view* and the *facilitator view*. The audience and the presenter use the *participant view*, which has a toggle for using the *feedback providing interface*. The discussion participant can be viewed by the presenter and audience and includes all comments marked for discussion (Figure 4). A single audience member acts as the *facilitator*, who handles marking which comments as they are discussed. The facilitator view allows recording and transcribing discussion audio and marking new for discussion (Figure 5).

5.2.1 Discussion Participant View. To help audience members and the presenter select comments for discussion, SlideSpecs features a *participant view* that can be projected onto the presentation screen. The participant view contains

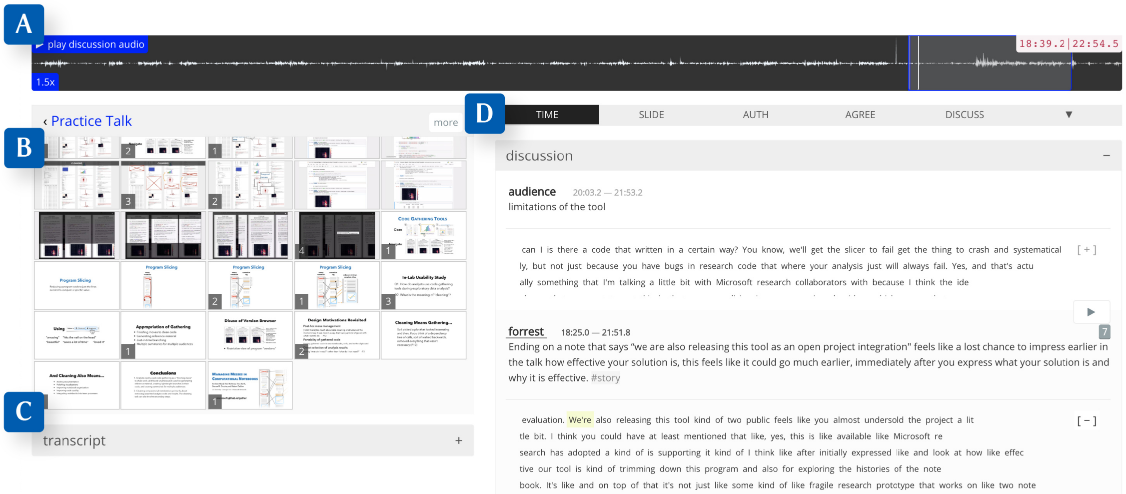

Fig. 6. The *Review* interface. (A) The audio waveform from the verbal discussion is shown along with playback and speed controls. Users can scrub and jump directly on the waveform to navigate through the audio file. (B) The slides pane. Clicking on a slide allows filtering the visible comments to only those which refer to the selected slide. (C) The full transcript of the discussion (shown collapsed). (D) The listing of all comments and discussion topics. Users can hover over a referenced comment to highlight the time range in which the comment was discussed in the waveform and view the directly related transcript segment. This connection is learned from the facilitator marking specific comments as being discussed during the discussion phase.

the slides pane (as in Figure 2) to ease slide referencing and a list of comments marked for discussion by the audience. By default, comments are shown in a *To Discuss* list until the group addresses them. Audience members can promote any comment for discussion by flagging the comment on their personal display.

5.2.2 Discussion Facilitator View. SlideSpecs uses a human facilitator efficiently handle searching and marking discussion topics in the specialized and technical domains we studied. A purely algorithmic solution could potentially automatically cross-correlate verbal and written comments, though supporting jargon-dense technical research domains remains hard. The facilitator uses a unique view to record the discussion audio and to mark when specific audience comments are being discussed. The timing information captured about when comments were discussed can be leveraged later by the presenter in the review interface. Each comment status is linked to the participant view so that when a comment is marked as being "discussed" by the facilitator, each participant's view also updates.

To quickly find relevant comments for discussion, the facilitator can leverage the list of comments flagged for discussion or create a new topic. The facilitator can also edit a topic's content to better match the discussion, given the facilitator will not often be sure what will be discussed ahead of time. The text input here doubles as a search box: when its text is updated, any comments with matching content or metadata are shown below the input as suggested topics for discussion. Multiple comments can be marked as being "discussed" at once, allowing for more leniency on the facilitator's timing as discussions will not always follow clear topic demarcations. These features serve to reduce the number of duplicate questions and topics, which can streamline the presenter's review process. Existing AI-based NLP techniques could help alleviate the task burden placed on the human facilitator during this phase. This includes finding comments that are related to the ongoing spoken discussion and summarizing redundant comments. These exciting directions are all further discussed in our future work section (Sec. 9).

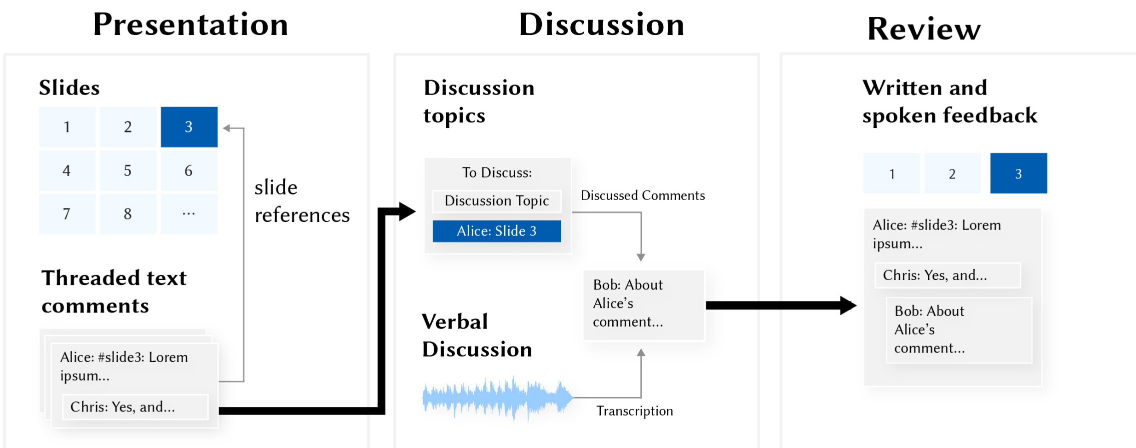

Fig. 7. During the presentation, audience members write critiques and respond to existing critiques. Optionally, audience members can include context with their text critiques (e.g., specific slides, tags for feedback type). The threaded text comments flagged for discussion appear in the Discussion interface. During the discussion, the facilitator marks which topics are being discussed, matching existing text critiques where possible. SlideSpecs records and transcribes the discussion audio, allowing the presenter to review the collated text and verbal feedback together.

5.3 Review Phase

SlideSpecs includes a *reviewing interface* to let the presenter efficiently review feedback. The reviewing interface features (A) an audio player/waveform, (B) the slides pane, (C) the discussion transcript, and (D) a sorting/filtering pane coupled with the complete comments list and the list of each discussed comment. The presenter can use the transcript to index into part of the discussion: clicking on a word will start playing the audio file at that point. Additionally, every discussed comment can start playing the audio for the time range that it was discussed and features a trimmed version of the transcript for that time range. The discussion audio can be played back at 1x, 1.5x, or 2x speed. To let the presenter view the slides and tags that received the most comments, we overlay the number of comments on each slide in the slide pane. The interface displays the number of comments with each tag in the sorting and filtering pane (Figure 6).

6 IMPLEMENTATION

An overview of the SlideSpecs system architecture is shown in Figure 7, which highlights the SideSpecs data processing pipelines. Here, we describe our slide-matching, web interface, facilitator search, and transcription implementations. The source code for this application is available online: https://github.com/BerkeleyHCI/SlideSpecs.

6.1 Slide Matching

To automatically provide relevant location hints for feedback, SlideSpecs predicts the currently active talk slide. We leverage color histogram information matching and optical character recognition for the presenter's screen. While presentation software keeps track of the active slide, presenters used a wide range of presentation tools (observed in Section 3: PowerPoint, Keynote, PDF, HTML/JS). This variety makes authoring a uniform plugin to monitor slide updates difficult. Instead, we use a custom technique for detecting the current slide based on presenter screenshots.

To enable slide-based references and automatic slide updates, the presenter uploads their presentation slides as a PDF to SlideSpecs. We automatically generate reference thumbnails for each slide from the presentation that serve as templates for matching. For each slide image, we precompute the included text using Tesseract's OCR Engine[1]. We also precompute a color histogram for each slide using OpenCV[2]. This is used to find the *currently active* presentation slide.

Before their talk, the presenter downloads author-provided Python software that captures the projected screen image at 1-second intervals. For each screenshot, we first compute the color histogram of the image (how much of each color the image contains). Each presentation generally starts at the first slide and progresses linearly forward. Since the ordering of slides is known, we can constrain and refine our predictions. We compare the histograms using the spatial distance correlation[3] along each color channel (RGB). If we have a confident histogram spatial match (e.g., less than a 0.01% difference) within a five-slide range, we immediately return that prediction. Still, some slides contain videos or have incremental builds/animations – visual complexity that will distort matches purely based on a color histogram.

If there is no confident histogram match in range, we recognize text within the screenshot (again with Tesseract[4]). We tokenize the extracted screenshot text with the spaCy NLP pipeline[5]. We compare the common tokens for each slide in the expected five-slide range against their precomputed slide text tokens. We attempt to match the screenshot to one of these slides based on the highest shared tokens and return that slide as a prediction. If there is still no confident match, we search over the entire slide range (e.g., the presenter advanced quickly, rewound, or a previous prediction was incorrect). We then return the highest match over all slides, allowing the system to adapt to fast transitions or backtracking. This live prediction is streamed to each audience member's interface, automatically updating the inferred active slide. This prediction helps automatically label audience-provided critiques with valuable context.

6.2 Web Interface

We used React.js to build each of the SlideSpecs interfaces. We deployed SlideSpecs with Meteor.js, enabling live updates across many clients in real-time. To access SlideSpecs, the presenter shares a unique SlideSpecs link with the audience. To guide the discussion, we enable users to flag comments they would like to discuss, which causes the comments to appear later in the audience *Discussion* interface. The facilitator can also view and control these flagged comments.

6.3 Transcription

To unify written feedback entered during the talk and spoken feedback during the discussion phase, we employ speech recognition using Google's Cloud[6]. SlideSpecs sends the recorded audio and retains returns individual speech tokens, timestamps, and confidence ratings. This transcript is available for context in the final presenter *Review* interface.

6.4 Discussion Context

To further link from the *Presentation* and *Discussion* phases, our facilitator *Discussion* interface enables both recording discussion audio and marking comments as the audience discusses them. This context is recorded and appears later in the presenter *Review* interface. The facilitator comment box doubles as a search interface that can display written

[1]https://github.com/tesseract-ocr/tesseract
[2]https://opencv-tutorial.readthedocs.io/en/latest/histogram/histogram.html
[3]https://docs.scipy.org/doc/scipy/reference/generated/scipy.spatial.distance.correlation.html
[4]https://github.com/tesseract-ocr/tesseract
[5]https://spacy.io/models/en#en_core_web_md
[6]https://cloud.google.com/speech-to-text

Table 1. An overview of feedback collated with SlideSpecs. Presentation lengths ranged from 32 to 84 total slides (mean: 48.6, s.d.: 31.6). Audience members contributed 86% of text comments during *Presentation* and 14% of comments during *Discussion*. *N* refers to the active group size (everyone who contributed at least one text or spoken comment), including the presenter, the facilitator, and the audience (mean: 9.5, s.d.: 3.8). Time (for Presentation/Discussion) refers to the minutes spent on each phase (mean: 17.0, s.d.: 3.5). Many comments display interaction with features beyond static comments (e.g., a slide reference, agreement, discussion marks, replies, or topic tags). We show the percentage of matching feedback for comments with different attributes (agreed, discussed, and slide-referencing comments). For example: on average, 68% of comments referenced at least one slide.

				Presentation		*Discussion*			comments with:				
#	topic	N	slides	time	comments	time	comments		agree	discuss	slide ref.	reply	tag
1	Haptic UIs	14	57	20	72	39	10		9	22	57	22	6
2	Notebooks	9	33	18	32	23	10		10	13	30	13	5
3	Tutorials	9	84	18	25	13	5		17	7	19	7	12
4	Electronics	8	40	15	37	7	4		21	9	26	9	11
5	Debugging	5	46	18	52	37	6		7	15	37	15	3
6	Virtual Reality	16	83	22	115	35	11		25	36	95	36	8
7	Computer Vision	10	41	14	38	15	12		11	16	33	16	6
8	Sensemaking	5	32	11	18	19	14		6	14	14	3	10
	average:	9.5	52.0	17.0	48.6	24.7	16.5	*% total:*	23%	29%	68%	6%	13%

presentation comments matching author, content, and slide number. Each new keystroke instantly updates the list of related comments below the comment box, allowing the facilitator to tag relevant comments for context rapidly.

7 EFFECTIVENESS STUDY: USING SLIDESPECS

To determine the effectiveness of SlideSpecs' feedback collation, we deployed SlideSpecs in eight practice presentations. These presentations took place at two large research universities and spread across four different research groups.

7.1 Method

We recruited presenters to give practice talks by sending emails via university mailing lists. For each talk, we first introduced SlideSpecs to the audience with a 5-minute tutorial. The list of presentation topics is shown in Table 1.

During the talk, the audience used personal laptops or mobile devices to contribute written critiques with SlideSpecs. After the talk, the group verbally discussed feedback while a talk facilitator recorded the audio and marked comments for discussion. Facilitators were selected based on their relative seniority and familiarity with the presentation topic. Four unique facilitators (all senior Ph.D. Students) worked to both note the discussion topic and mark any relevant existing comments from the presentation during the discussion phase over these eight talks. We encouraged discussion participants (audience members) to discuss their feedback as they usually would. After the presentation and discussion, the presenter used SlideSpecs to review critiques. We logged all participant feedback and interaction with SlideSpecs.

After the talk, we sent an optional survey to each audience member with open-ended and Likert scale questions addressing likes, dislikes, and usefulness of system features. Each presenter and facilitator completed a similar survey on their experience with SlideSpecs. We also interviewed each presenter after they updated their presentations to learn if and how they used SlideSpecs to revise their talk. The audio of each presenter interview was recorded. We gathered more open-ended feedback on their experience and learned how SlideSpecs compared to previously used methods.

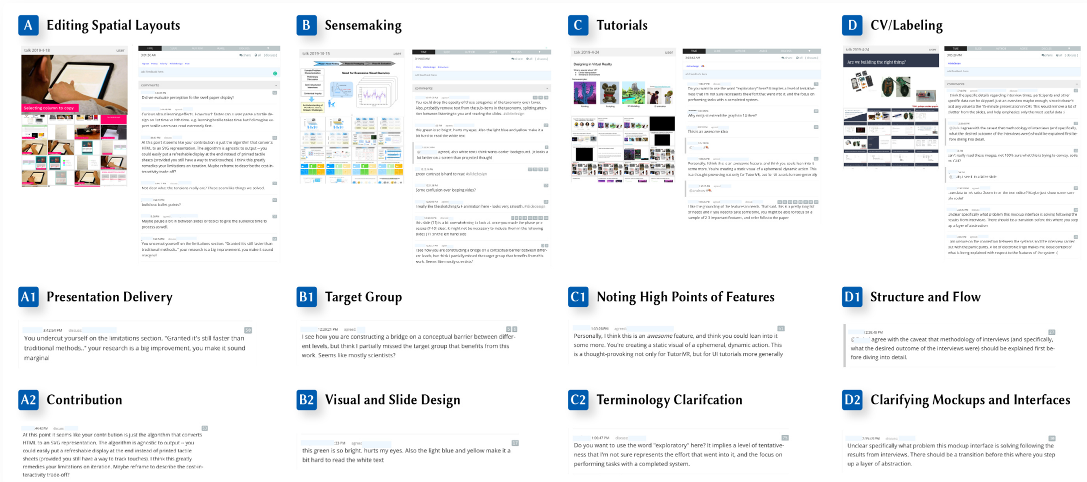

Fig. 8. Participants used SlideSpecs for a wide variety of practice talks (A-D). They received feedback concerning many different aspects of their presentation, from slide design, delivery, and requests for clarification. We highlight select feedback from four different practice presentations. These comments showcase the range of topics that the audience's feedback addressed; eight comments (A1-D2) are shown in detail alongside a rough characterization of what part of the presentation the comment addresses.

7.2 Findings

SlideSpecs effectively supported presenters by collating feedback into a single accessible and context-rich location. Presenters reported that using SlideSpecs improved feedback organization, provided valuable context, and reduced redundancy. When surveyed, 85% of presenters and responding audience members reported they'd use the tool again (Figure 9). In total, 52 unique audience members provided feedback over the eight presentations we deployed SlideSpecs. Of those, 14 audience members (A1-A14) reported on their experience in a voluntary survey (6/8 talks had at least one responding audience member). A1, A3, and A7 mentioned using collaborative tools for feedback (Google Sheets/Docs) but preferred SlideSpecs. We report further on the experience of both the audience members and the presenters.

SlideSpecs provided presenters valuable feedback organization. All presenters reported receiving useful feedback and that seeing the feedback organized by slides was useful (Figure 9). 5/8 presenters specifically rated organizing comments by authoring time that people provided critiques as useful for revision: it gave a "blow-by-blow account of how people reacted" during the presentation and discussion (P1). Beyond slides and comment time, one presenter sorted by slides tagged comment types. Another presenter combined sorting techniques to prioritize feedback with the most consensus (e.g., agreement), then by the most discussed, and then by slide to form a *to-do* checklist.

SlideSpecs provided presenters valuable feedback context. Audience members used SlideSpecs to link slide references to 68% of critiques (313/461). They used tags (e.g., *#design*, *#story*, *#content*) to categorize their comments into subject types much less frequently (13%, 60/461). This difference could partially be because of the manual nature that comments had to be tagged. 8/14 reporting audience members explicitly mentioned liking being able to link slides to their critiques using slide images. Linking slides to images "saves typing", lets audience members reference slides automatically when generating critiques, and frees audience members from recalling slide numbers. Presenters found the transcription and recording helpful when revisiting slides further from the original practice talk. 4/8 presenters

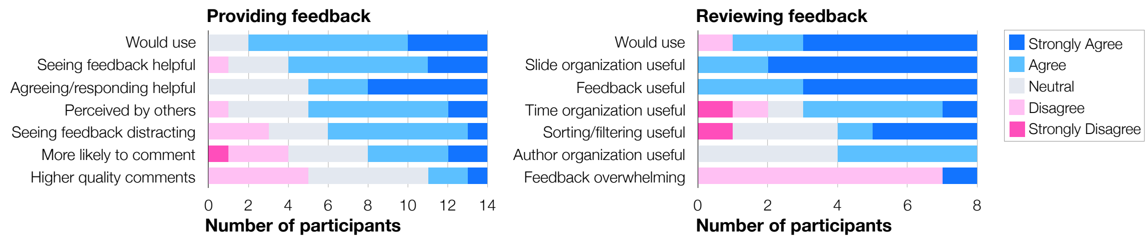

Fig. 9. Likert scale responses for 14 audience members (left) and eight presenters (right) in the group deployment. All except one presenter and two audience members agreed that they would use the interface in the future (85%). Likert scale responses reflect that audience members broadly find seeing the real-time of others to be very helpful (10/14) though somewhat distracting (6/14). Audience members noted seeing and agreeing/responding to feedback was helpful. All 8/8 presenters favored slide organization during the feedback review and found the feedback collation useful.

reported reading or listening to the entire discussion section recording during their review. While most participants used either the transcript or recording to listen to segments of interest, one participant reported that in the future, they would use the transcript for efficiency but "If there was ambiguity, even if just tone, I would play the audio to see what people actually meant" (P7). P3 described their process of supplementing listening with written comments: 'Took me about 20 minutes to listen through the whole thing. I went back and back-filled with additional comments onto the written comments. A lot of that stuff came up in the discussion that normally I think I just would've forgotten.'

Shared group awareness help reduce redundant presentation feedback. Audience members collaboratively provided, on average, 48 written comments per talk, providing both local and global critiques. Audience members also interacted with others during critique: 23% of comments had at least one agreement (106 agreements over eight talks). 7/14 audience members reported liking seeing other audience members' critiques in real-time as it: "helps reduce redundant feedback" (A2) by providing an "easy medium to share and agree" (A10). Each agreement was potentially an instance of feedback that the presenter would have received multiple copies of.

Shared group awareness help promote thematic feedback consensus. Several presenters reported noticing more thematic consistency in the verbal discussion that followed the practice talk after using SlideSpecs. P1 reported: 'During the verbal feedback session after my talk, people spoke not only about their individual comments but also about themes they saw emerging on SlideSpecs. So, I felt like I got more coherent verbal feedback that covered all the major issues in a pretty coherent way.' A8 reported: 'During the discussion, another audience member said, 'Here are the themes I saw from the comments' – and this feedback was so useful.' The visibility of the unified thread of discussed comments helped the audience merge multiple related comments into higher-level themes during the discussion.

The Facilitator's tasks are demanding. While the presenters benefited from the discussion organization, the facilitators generally found the role to be demanding. For example, F1 reported that it could be unclear which comments are being discussed: 'Since I could choose to make a new topic or mark existing ones for discussion, I sometimes stumbled between just marking a new topic quickly and trying to find the most relevant comment to bring up for discussion'. F3 also reported on this challenge: 'Facilitating was challenging, as I was trying to both follow the conversation and think of what would be relevant for the presenter later on. Even though I used the comment search and author filtering, deciding the true "most relevant" comment was ambiguous, along with deciding when a comment was no longer being discussed'. While the context that the facilitator provided to the discussion was valuable, the role's tasks were complex and would benefit from automation support (e.g., text summarization, identifying relevant comments, and marking new topics).

7.3 Study Limitations

Participant Usage Period. Most presenter participants used our tool only once; thus, our findings have limitations common to first-use studies. We have limited information about how tool use might change over time and how it may impact the nature and flow of group discussions and practices.

Domain and Talk Format. Our sampling of presentations was limited to only include STEM fields. SlideSpecs likely doesn't support processes or talks that avoid slides or using computers (e.g., only verbal discussions).

In-Person Presentations Only. We only studied in-person practice talks in this study. However, remote and hybrid talks are increasingly common, and SlideSpecs may serve presenters in that context differently.

Audience Survey Response Rate. Despite having 52 unique audience participants in our practice talk sessions, only 14 participants sourced from 6/8 talks responded to our optional survey (27%). While we likely captured the strongest and most polarized perspectives, we missed out on reporting a broader set of user experiences.

8 DISCUSSION

Through building and evaluating SlideSpecs, we gained a thorough understanding of what a presentation feedback tool should support. We reflect on our performance through several design implications for feedback collation and review.

8.1 Peer awareness helps reduce feedback redundancy.

Our results demonstrated that incorporating peer awareness mechanisms (e.g., seeing/interacting with other audience feedback) reduced redundancy in the final collected feedback. For instance, audience members interacted with others during the critique by expressing agreement (e.g., liking comments, commenting in a thread to elaborate on another member's feedback), providing an easy mechanism for prioritizing specific feedback and reducing repeated statements.

8.2 No single tool best accommodates every feedback process.

A key design implication from our formative study was *supporting existing workflows for group feedback*, so we developed SlideSpecs to support a commonly observed (yet fairly linear) feedback process. SlideSpecs is designed to support existing practices that center around receiving in-person feedback. Still, no single tool can accommodate the entire range of feedback processes that different groups use. For example, SlideSpecs requires a deck of slides for the audience to comment on, and some talks forgo slides. Another assumption is that presenters will follow a linear pattern of Presentation, Discussion, and Review (Fig. 3). However, especially in longer talks, presenters may want to pause between each section to gather feedback incrementally. Another more volatile process might involve stopping anytime an audience member has a comment or question. To support more dynamic feedback structures like this, a different approach could be used where audio is constantly recorded/transcribed, regardless of what feedback phase is happening.

The eight presentations in Table 1 cover both a broad range of subject matter and a range of structures, including practicing for conference talks, grant meeting progress updates, and academic job talks. Despite this variation, many other types of talk structures exist that we did not evaluate (e.g., educational lessons, startup pitch decks). However, no evidence suggests that varying the internal structure of the talk makes SlideSpecs less relevant or valuable. Despite only brief training, the groups we worked with adopted SlideSpecs into many existing practice processes. On top of this, the presented talks in our study had a mix of intended purposes (conference, research update, job talk).

Hybrid in-person/online presentations also present unique challenges. Inviting online participants offers compelling benefits – allowing geographically remote participants, built-in recording features (e.g., Zoom), and more easily scaling the audience size. SlideSpecs could be deployed in these contexts as-is, though it would not be leveraging this new contextual data (e.g., multiple audio feeds for discussion, video/screen recordings for presentations). Still, the mixed social dynamic of hybrid meetings presents new challenges that SlideSpecs is not explicitly designed to support.

8.3 Feedback context helps presenters revise effectively.

SlideSpecs captures many forms of context automatically, which helps presenters revise their presentations more effectively. This included *scope:* what parts or slides the feedback is referring to, and *subject:* what category or type of feedback is being given. For instance, presenters may leverage the flexibility of SlideSpecs to sort the slides to achieve different revision goals: e.g., prioritize slides that require the most changes, fix comments with the most consensus, or revise the comments participants discussed first. This range of contextual data supports how presenters may weigh feedback. Given the value of feedback context, comments could also be automatically clustered around a tag or theme.

8.4 Feedback quality consists of more than context.

Another implicit assumption is that the audience can already provide valuable and relevant feedback. SlideSpecs makes providing feedback with enriching context simple, though it cannot make up for being unfamiliar with the talk domain or not knowing how to give effective feedback. While the shared context can inform more novice audience members about how to give feedback, it doesn't inherently instruct the audience *what* feedback to provide the presenter. While our contributing audience members varied from undergrads to professors, we did not measure feedback quality against expertise. A follow-up study could compare these qualities; participants in our formative study also reported implicitly weighting feedback with the provider's role (e.g., the lab PI vs. an undergraduate research assistant).

9 FUTURE WORK

An exciting vision is to integrate AI and NLP techniques to boost the usefulness of tools like SlideSpecs. We review a few of the most promising directions for future work in supporting group feedback understanding and organization.

9.1 Reducing the facilitator and audience workload.

An integral part of feedback classification and distilling key themes within SlideSpecs relies on discussion facilitators. For example, facilitators markup post-presentation discussions with links to comments and new topics, distilling key discussion points into actionable feedback. In addition, audience members contribute to the context by tagging comments, using side references in feedback, and discussion facilitation. We suggest that future systems focus on reducing facilitator and audience work by automatically semantically grouping and labeling feedback. Discussions could be enhanced if a system could automatically recommend topics based on the magnitude of similar feedback. Overall, future work can enhance feedback systems by automating most of the facilitator's tasks and reducing audience members' labor by inferring more context.

9.2 Refining and distilling gathered feedback into actionable changes.

In our study, we worked with groups that had roughly ten audience members (Table 1). While there is no *right* number of audience members, this size worked well for these presentations. A larger audience size could yield a higher quantity of valuable feedback and also increase the relevance of the automatic summarization of the provided comments. Future work could automatically generate more context by allocating comments (both written and spoken) to slides as they are entered based on their content. This may pave the way to supporting more dynamic feedback processes that feature common audience interruptions. While increased group awareness can reduce redundancy when receiving feedback from a group, participants still sometimes enter similar redundant comments. Adding an automatic comment summarization and aggregation pipeline into the *Review* phase would further streamline the presenter experience. Finally, LLMs could transform these aggregated feedback points into a concise list of more refined actionable changes.

9.3 Updating presentation structure and content from feedback automatically.

The work is not quite over once the presenter has reviewed and synthesized a list of concrete changes to make to their presentation. They must still manually review their slides and update the content according to the gathered feedback. Depending on the format of the presentation, it is very feasible to envision some of these transformations being automated. While proprietary formats like PowerPoint and Keynote may be less accessible, some presenters use open text-based slide formats, including LaTeX (with Beamer[7]), HTML (with Reveal.js[8]), and Markdown (with Marp[9]). Existing LLMs could transform this existing text-based slide format based on the feedback gathered from the audience. Future work could transform existing slide decks, including synthesizing new slides, modifying existing slides, cutting out slides, and restructuring or reordering the presentation slides.

10 CONCLUSION

We present SlideSpecs, a novel system for automatically and interactively collating and contextualizing talk feedback. Grounded in our formative study findings, SlideSpecs allows presenters to organize and review their presentation feedback effectively. SlideSpecs improves the revision process by unifying text and spoken feedback into a centralized space for seamless review. We demonstrate the effectiveness of SlideSpecs by deploying it in eight unique presentations across computer vision, programming notebooks, sensemaking, and more. We find that audience members contributed a wide variety of feedback successfully on a range of slide presentations. Presenters reported that using SlideSpecs while refining their talk improved their feedback organization, provided valuable context, and reduced redundant comments. Future collaborative revision and feedback systems can benefit users by integrating more automatic contextualization.

ACKNOWLEDGMENTS

This work would not have been possible without the presenters and research groups that were open to trying out SlideSpecs – thanks to all who participated and provided feedback with SlideSpecs. We also extend thanks to Pancham Yadav who assisted with the initial literature review. Thanks to James Smith for creating the project video figure.

REFERENCES

[1] Reza Asadi, Ha Trinh, Harriet J. Fell, and Timothy W. Bickmore. 2017. IntelliPrompter: Speech-Based Dynamic Note Display Interface for Oral Presentations. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction* (Glasgow, UK) *(ICMI '17)*. Association for Computing Machinery, New York, NY, USA, 172–180. https://doi.org/10.1145/3136755.3136818

[2] Reza Asadi, Ha Trinh, Harriet J. Fell, and Timothy W. Bickmore. 2018. Quester: A Speech-Based Question Answering Support System for Oral Presentations. In *23rd International Conference on Intelligent User Interfaces* (Tokyo, Japan) *(IUI '18)*. Association for Computing Machinery, New York, NY, USA, 583–593. https://doi.org/10.1145/3172944.3172974

[3] Karim Benharrak, Florian Lehmann, Hai Dang, and Daniel Buschek. 2022. SummaryLens – A Smartphone App for Exploring Interactive Use of Automated Text Summarization in Everyday Life. In *27th International Conference on Intelligent User Interfaces* (Helsinki, Finland) *(IUI '22 Companion)*. Association for Computing Machinery, New York, NY, USA, 93–96. https://doi.org/10.1145/3490100.3516471

[4] Gary Bernhardt. 2018. *How to Prepare a Talk.* Retrieved Sep 21, 2022 from https://www.deconstructconf.com/blog/how-to-prepare-a-talk

[7] https://en.wikipedia.org/wiki/Beamer_(LaTeX)
[8] https://github.com/hakimel/reveal.js/
[9] https://github.com/marp-team/marp

[5] Patrick Chiu, John S Boreczky, Andreas Girgensohn, and Don Kimber. 2001. LiteMinutes: an Internet-based system for multimedia meeting minutes. *WWW* 1 (2001), 140–149.

[6] Patrick Chiu, Ashutosh Kapuskar, Sarah Reitmeier, and Lynn Wilcox. 1999. NoteLook: Taking notes in meetings with digital video and ink. In *Proceedings of the seventh ACM international conference on Multimedia (Part 1)*. ACM, 149–158.

[7] Kwangsu Cho and Christian D Schunn. 2007. Scaffolded writing and rewriting in the discipline: A web-based reciprocal peer review system. *Computers & Education* 48, 3 (2007), 409–426.

[8] Yaakov Danone, Tsvi Kuflik, and Osnat Mokryn. 2018. Visualizing Reviews Summaries as a Tool for Restaurants Recommendation. In *23rd International Conference on Intelligent User Interfaces* (Tokyo, Japan) *(IUI '18)*. Association for Computing Machinery, New York, NY, USA, 607–616. https://doi.org/10.1145/3172944.3172947

[9] Laurent Denoue, Scott Carter, and Matthew Cooper. 2015. Searching Live Meeting Documents "Show Me the Action". In *Proceedings of the 2015 ACM Symposium on Document Engineering* (Lausanne, Switzerland) *(DocEng '15)*. Association for Computing Machinery, New York, NY, USA, 195–198. https://doi.org/10.1145/2682571.2797082

[10] Mark Dredze, Hanna M. Wallach, Danny Puller, and Fernando Pereira. 2008. Generating Summary Keywords for Emails Using Topics. In *Proceedings of the 13th International Conference on Intelligent User Interfaces* (Gran Canaria, Spain) *(IUI '08)*. Association for Computing Machinery, New York, NY, USA, 199–206. https://doi.org/10.1145/1378773.1378800

[11] 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* (Montreux, Switzerland) *(UIST '06)*. Association for Computing Machinery, New York, NY, USA, 47–56. https://doi.org/10.1145/1166253.1166263

[12] C. Ailie Fraser, Tricia J. Ngoon, Ariel S. Weingarten, Mira Dontcheva, and Scott Klemmer. 2017. CritiqueKit: A Mixed-Initiative, Real-Time Interface For Improving Feedback. (2017), 7–9. https://doi.org/10.1145/3131785.3131791

[13] Werner Geyer, Heather Richter, and Gregory D Abowd. 2005. Towards a smarter meeting record—capture and access of meetings revisited. *Multimedia Tools and Applications* 27, 3 (2005), 393–410.

[14] Elena L Glassman, Juho Kim, Andrés Monroy-Hernández, and Meredith Ringel Morris. 2015. Mudslide: A spatially anchored census of student confusion for online lecture videos. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1555–1564.

[15] Lance Good and Benjamin B Bederson. 2001. *CounterPoint: Creating jazzy interactive presentations*. Technical Report.

[16] Jonathan Grudin. 1995. Groupware and social dynamics: Eight challenges for developers. In *Readings in Human–Computer Interaction*. Elsevier, 762–774.

[17] Drew Harry, Joshua Green, and Judith Donath. 2009. Backchan.nl: integrating backchannels in physical space. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 1361–1370.

[18] Yoshinori Hijikata, Hanako Ohno, Yukitaka Kusumura, and Shogo Nishida. 2006. Social Summarization of Text Feedback for Online Auctions and Interactive Presentation of the Summary. In *Proceedings of the 11th International Conference on Intelligent User Interfaces* (Sydney, Australia) *(IUI '06)*. Association for Computing Machinery, New York, NY, USA, 242–249. https://doi.org/10.1145/1111449.1111500

[19] Hiroshi Ishii and Minoru Kobayashi. 1992. ClearBoard: a seamless medium for shared drawing and conversation with eye contact. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 525–532.

[20] Hyeungshik Jung, Hijung Valentina Shin, and Juho Kim. 2018. DynamicSlide: Reference-Based Interaction Techniques for Slide-Based Lecture Videos. In *The 31st Annual ACM Symposium on User Interface Software and Technology Adjunct Proceedings* (Berlin, Germany) *(UIST '18 Adjunct)*. Association for Computing Machinery, New York, NY, USA, 23–25. https://doi.org/10.1145/3266037.3266089

[21] Young-Wook Jung, Youn-kyung Lim, and Myung-suk Kim. 2017. Possibilities and Limitations of Online Document Tools for Design Collaboration: The Case of Google Docs. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 1096–1108.

[22] Fawzia Khan. 1993. *A survey of note-taking practices*. Hewlett-Packard Laboratories.

[23] Scott R Klemmer, Michael Thomsen, Ethan Phelps-Goodman, Robert Lee, and James A Landay. 2002. Where do web sites come from?: capturing and interacting with design history. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 1–8.

[24] Chinmay E Kulkarni, Michael S Bernstein, and Scott R Klemmer. 2015. PeerStudio: rapid peer feedback emphasizes revision and improves performance. In *Proceedings of the Second (2015) ACM Conference on Learning@ Scale*. ACM, 75–84.

[25] Philippe Laban, Elicia Ye, Srujay Korlakunta, John Canny, and Marti Hearst. 2022. NewsPod: Automatic and Interactive News Podcasts. In *27th International Conference on Intelligent User Interfaces* (Helsinki, Finland) *(IUI '22)*. Association for Computing Machinery, New York, NY, USA, 691–706. https://doi.org/10.1145/3490099.3511147

[26] Sahiti Labhishetty, Bhavya, Kevin Pei, Assma Boughoula, and Chengxiang Zhai. 2019. Web of Slides: Automatic Linking of Lecture Slides to Facilitate Navigation. In *Proceedings of the Sixth (2019) ACM Conference on Learning @ Scale* (Chicago, IL, USA) *(L@S '19)*. Association for Computing Machinery, New York, NY, USA, Article 54, 4 pages. https://doi.org/10.1145/3330430.3333668

[27] Dar-Shyang Lee, Berna Erol, Jamey Graham, Jonathan J Hull, and Norihiko Murata. 2002. Portable meeting recorder. In *Proceedings of the tenth ACM international conference on Multimedia*. ACM, 493–502.

[28] Leonhard Lichtschlag, Thorsten Karrer, and Jan Borchers. 2009. Fly: a tool to author planar presentations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 547–556.

[29] Rhema Linder, Nic Lupfer, Andruid Kerne, Andrew M Webb, Cameron Hill, Yin Qu, Kade Keith, Matthew Carrasco, and Elizabeth Kellogg. 2015. Beyond slideware: How a free-form presentation medium stimulates free-form thinking in the classroom. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*. ACM, 285–294.

[30] Kurt Luther, Jari-Lee Tolentino, Wei Wu, Amy Pavel, Brian P Bailey, Maneesh Agrawala, Björn Hartmann, and Steven P Dow. 2015. Structuring, aggregating, and evaluating crowdsourced design critique. In *Proc. CSCW'15*. ACM, 473–485.

[31] Joseph F McCarthy, Elizabeth F Churchill, William G Griswold, Elizabeth Lawley, Melora Zaner, et al. 2004. Digital backchannels in shared physical spaces: attention, intention and contention. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*. ACM, 550–553.

[32] Scott Minneman, Steve Harrison, Bill Janssen, Gordon Kurtenbach, Thomas Moran, Ian Smith, and Bill van Melle. 1995. A confederation of tools for capturing and accessing collaborative activity. In *ACM Multimedia*, Vol. 95.

[33] Gabriel Murray, Steve Renals, and Jean Carletta. 2005. Extractive summarization of meeting recordings. (2005).

[34] NCSL 2017. Tips for making effective Powerpoint presentations. http://www.ncsl.org/legislators-staff/legislative-staff/legislative-staff-coordinating-committee/tips-for-making-effective-powerpoint-presentations.aspx.

[35] Ani Nenkova, Kathleen McKeown, et al. 2011. Automatic summarization. *Foundations and Trends® in Information Retrieval* 5, 2–3 (2011), 103–233.

[36] Jeffrey Nichols, Jalal Mahmud, and Clemens Drews. 2012. Summarizing Sporting Events Using Twitter. In *Proceedings of the 2012 ACM International Conference on Intelligent User Interfaces* (Lisbon, Portugal) *(IUI '12)*. Association for Computing Machinery, New York, NY, USA, 189–198. https://doi.org/10.1145/2166966.2166999

[37] Xavier Ochoa, Federico Domínguez, Bruno Guamán, Ricardo Maya, Gabriel Falcones, and Jaime Castells. 2018. The RAP System: Automatic Feedback of Oral Presentation Skills Using Multimodal Analysis and Low-Cost Sensors. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge* (Sydney, New South Wales, Australia) *(LAK '18)*. Association for Computing Machinery, New York, NY, USA, 360–364. https://doi.org/10.1145/3170358.3170406

[38] Lora Oehlberg, Kyu Simm, Jasmine Jones, Alice Agogino, and Björn Hartmann. 2012. Showing is sharing: building shared understanding in human-centered design teams with Dazzle. In *Proceedings of the Designing Interactive Systems Conference*. ACM, 669–678.

[39] Jasper O'Leary, Holger Winnemöller, Wilmot Li, Mira Dontcheva, and Morgan Dixon. 2018. Charrette: Supporting In-Person Discussions around Iterations in User Interface Design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 535.

[40] Thierry Olive. 2004. Working memory in writing: Empirical evidence from the dual-task technique. *European Psychologist* 9, 1 (2004), 32–42.

[41] Amy Pavel, Dan B Goldman, Björn Hartmann, and Maneesh Agrawala. 2016. VidCrit: Video-based Asynchronous Video Review. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, 517–528.

[42] Yi-Hao Peng, Jeffrey P Bigham, and Amy Pavel. 2021. Slidecho: Flexible Non-Visual Exploration of Presentation Videos. In *The 23rd International ACM SIGACCESS Conference on Computers and Accessibility* (Virtual Event, USA) *(ASSETS '21)*. Association for Computing Machinery, New York, NY, USA, Article 24, 12 pages. https://doi.org/10.1145/3441852.3471234

[43] Yi-Hao Peng, JiWoong Jang, Jeffrey P Bigham, and Amy Pavel. 2021. Say It All: Feedback for Improving Non-Visual Presentation Accessibility. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 276, 12 pages. https://doi.org/10.1145/3411764.3445572

[44] Annie Piolat, Thierry Olive, and Ronald T Kellogg. 2005. Cognitive effort during note taking. *Applied cognitive psychology* 19, 3 (2005), 291–312.

[45] Alexander Pohl, Vera Gehlen-Baum, and François Bry. 2011. Introducing Backstage–a digital backchannel for large class lectures. *Interactive Technology and Smart Education* 8, 3 (2011), 186–200.

[46] Matt Ratto, R Benjamin Shapiro, Tan Minh Truong, and William G Griswold. 2003. The activeclass project: Experiments in encouraging classroom participation. In *Designing for change in networked learning environments*. Springer, 477–486.

[47] Jun Rekimoto, Yuji Ayatsuka, Hitoraka Uoi, and Toshifumi Arai. 1998. Adding another communication channel to reality: an experience with a chat-augmented conference. In *Conference on Human Factors in Computing Systems: CHI 98 conference summary on Human factors in computing systems*, Vol. 18. 271–272.

[48] Garr Reynolds. 2011. *Presentation Zen: Simple ideas on presentation design and delivery*. New Riders.

[49] Athar Sefid, Prasenjit Mitra, and Lee Giles. 2021. SlideGen: An Abstractive Section-Based Slide Generator for Scholarly Documents. In *Proceedings of the 21st ACM Symposium on Document Engineering* (Limerick, Ireland) *(DocEng '21)*. Association for Computing Machinery, New York, NY, USA, Article 11, 4 pages. https://doi.org/10.1145/3469096.3474939

[50] Amy Shannon, Jessica Hammer, Hassler Thurston, Natalie Diehl, and Steven Dow. 2016. PeerPresents: A web-based system for in-class peer feedback during student presentations. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. ACM, ACM, 447–458.

[51] Tian Shi, Yaser Keneshloo, Naren Ramakrishnan, and Chandan K. Reddy. 2021. Neural Abstractive Text Summarization with Sequence-to-Sequence Models. *ACM/IMS Trans. Data Sci.* 2, 1, Article 1 (jan 2021), 37 pages. https://doi.org/10.1145/3419106

[52] Steven J. Simske and Rafael Lins. 2018. Automatic Text Summarization and Classification. In *Proceedings of the ACM Symposium on Document Engineering 2018* (Halifax, NS, Canada) *(DocEng '18)*. Association for Computing Machinery, New York, NY, USA, Article 1, 2 pages. https://doi.org/10.1145/3209280.3232791

[53] Ryan Spicer, Yu-Ru Lin, Aisling Kelliher, and Hari Sundaram. 2012. NextSlidePlease: Authoring and delivering agile multimedia presentations. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 8, 4 (2012), 53.

[54] TED Staff. 2014. 10 tips for better slide decks. https://blog.ted.com/10-tips-for-better-slide-decks/.

[55] David Tinapple, Loren Olson, and John Sadauskas. 2013. CritViz: Web-based software supporting peer critique in large creative classrooms. *Bulletin of the IEEE Technical Committee on Learning Technology* 15, 1 (2013), 29.

[56] Andy Warr and Eamonn O'Neill. 2005. Understanding design as a social creative process. In *Proceedings of the 5th conference on Creativity & cognition*. ACM, 118–127.

[57] Pierre Wellner, Mike Flynn, and Maël Guillemot. 2004. Browsing recorded meetings with Ferret. In *International Workshop on Machine Learning for Multimodal Interaction*. Springer, 12–21.

[58] Steve Whittaker, Patrick Hyland, and Myrtle Wiley. 1994. Filochat: Handwritten notes provide access to recorded conversations. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 271–277.

[59] Steve Whittaker, Rachel Laban, and Simon Tucker. 2005. Analysing meeting records: An ethnographic study and technological implications. In *International Workshop on Machine Learning for Multimodal Interaction*. Springer, 101–113.

[60] Wesley Willett, Jeffrey Heer, Joseph Hellerstein, and Maneesh Agrawala. 2011. CommentSpace: structured support for collaborative visual analysis. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM, 3131–3140.

[61] Zhenmin Yang, Yonghao Dong, Jiange Deng, Baocheng Sha, and Tao Xu. 2021. Research on Automatic News Text Summarization Technology Based on GPT2 Model. In *2021 3rd International Conference on Artificial Intelligence and Advanced Manufacture* (Manchester, United Kingdom) *(AIAM2021)*. Association for Computing Machinery, New York, NY, USA, 418–423. https://doi.org/10.1145/3495018.3495091

[62] Dongwook Yoon, Nicholas Chen, François Guimbretière, and Abigail Sellen. 2014. RichReview: blending ink, speech, and gesture to support collaborative document review. In *Proc. UIST'14*. ACM, ACM, 481–490.

[63] Dongwook Yoon, Nicholas Chen, Bernie Randles, Amy Cheatle, Corinna E Löckenhoff, Steven J Jackson, Abigail Sellen, and François Guimbretière. 2016. RichReview++: Deployment of a Collaborative Multi-modal Annotation System for Instructor Feedback and Peer Discussion. In *Proc. CSCW'16*. ACM, ACM, 195–205.

[64] Amy X Zhang and Justin Cranshaw. 2018. Making sense of group chat through collaborative tagging and summarization. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 196.

[65] Amy X. Zhang, Lea Verou, and David Karger. 2017. Wikum: Bridging Discussion Forums and Wikis Using Recursive Summarization. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (Portland, Oregon, USA) *(CSCW '17)*. ACM, New York, NY, USA, 2082–2096. https://doi.org/10.1145/2998181.2998235