

ReSlide: Towards Effective Presentation Authoring for Evolving Narratives and Contextual Constraints

by

Linxiu Zeng

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Computer Science

Waterloo, Ontario, Canada, 2025

© Linxiu Zeng 2025

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Authoring slides for a public presentation requires speakers to navigate multiple contextual constraints that directly shape how content is structured and delivered. We first conducted a formative study with ten experienced presenters to understand what constraints are prioritized and how they influence slide deck creation. We identified three key constraints—time, audience, and communicative intent—and challenges in integrating them into slide authoring for a single presentation and long-term needs with diverse narratives. We designed **ReSlide**, a presentation authoring tool that helps presenters create and reuse slides by bridging contextual constraints with evolving narratives. We evaluated ReSlide in a within-subjects study with 12 participants against a baseline tool, and in an exploratory study with eight professional presenters using only ReSlide. Results indicate that Reslide’s novel features of constraint awareness and multi-granular slide reuse helped presenters effectively craft presentations in both one and multiple authoring cycles.

Acknowledgements

First and foremost, special thanks to my Supervisor!

A big shoutout to my lab mates and collaborators for all the suggestions, feedback, brainstorming, and random hallway conversations that somehow turned into thesis ideas. You made this journey not only productive but also genuinely fun. :)

And, of course, endless gratitude to my family and friends. Thank you for listening to my rants, keeping me caffeinated, and reminding me that life exists outside of LaTeX. This thesis wouldn't be here without you. I'll cherish all moments that we've spent together.

My dear Heavenly Father, I don't know how to express my deepest thanks to you. You have always been with me, supporting me and continually offering the best to me. Without You, I would not have had the wisdom, persistence, and courage to finish this work. Hallelujah!

Dedication

To my Heavenly Father, the one and the only God. To my family, friends, and all the beloved ones.

Table of Contents

Author’s Declaration	ii
Abstract	iii
Acknowledgements	iv
Dedication	v
List of Figures	x
List of Tables	xii
1 Introduction	1
2 Related Work	5
2.1 Contextual Constraints and Their Management in Presentations	5
2.2 Systems for Supporting Narrative Creation	7
2.3 Techniques for Reusing Presentations	8
3 Formative Study	10
3.1 Study Setup	10
3.2 Challenges in Current Presentation Authoring Process	11
3.2.1 C1: Fragmented and ad-hoc constraints handling.	11

3.2.2	C2: Difficulty in maintaining narrative coherence.	12
3.2.3	C3: Inefficient content management and reuse.	12
3.3	Design Goals	13
4	ReSlide	15
4.1	System Overview	15
4.2	Define Unit of Authoring with Section	16
4.3	Support Narrative Construction with Timeline and Jargon Detection	18
4.4	Enable Multi-Layered Presentation Reuse and Synchronization with Central Repository	20
4.5	Example Workflow	22
5	Evaluation	24
5.1	General Study Setup	24
5.2	Study 1: Comparative Study with Single-Session Authoring	25
5.2.1	Participants	25
5.2.2	Task	26
5.2.3	Procedure	26
5.3	Study 2: Exploratory Study with Multi-Session Authoring	27
5.3.1	Participants	27
5.3.2	Task	27
5.3.3	Procedure	27
5.4	Data Collection and Analysis	28
6	Results	29
6.1	Questionnaire Results from Study 1	29
6.2	Externalizing Constraints Made Intent More Tangible and Actionable (RQ1 + RQ2)	30
6.3	Structuring with Sections Shifted Presenters' Authoring Focus from Slides to Ideas (RQ2)	31

6.3.1	The Timeline Overview Assisted Idea Reflection and Organization .	31
6.3.2	Section-first Workflow Encouraged Narrative-centered Authoring . .	32
6.4	Visualizing Time and Emphasis Effectively Guided Presenters to Resolve Constraints (RQ1 + RQ2)	33
6.4.1	Reflecting on Sections Enables Accurate and Flexible Time Estimation	33
6.4.2	Leveraging Visual Indicators in Timeline Overview to Prevent Overloading and Unbalanced Idea Development	33
6.5	Jargon-checking Prompted Reflection and Inspiration (RQ1 + RQ2)	34
6.6	Modular Reuse and Synchronization Supported Flexible Authoring Across Versions (RQ2+ RQ3)	35
6.6.1	Exploratory Reuse for Idea Generation and Inspiration	36
6.6.2	Intentional Reuse for Targeted Content Adaptation	36
6.6.3	Saving and Synchronizing Across Versions to Manage Evolving Content	37
7	Discussion	38
7.1	Design Implications	39
7.1.1	Future Tools Should Support Context-aware Authoring	39
7.1.2	Future Tools Should More Carefully Balance Between AI-automation and Author Controls	39
7.1.3	Future Tools Should Include Support for Longitudinal Workflows .	40
7.2	Limitations and Future Work	40
8	Conclusion	42
	References	43
	APPENDICES	65
.1	Jargon Checking prompts	65
.1.1	Prompt for Expanded Audience Profile	65
.1	Prompts for Audience Context Expansion	65

.1.1	Prompt for Jargon Detection	67
.2	Evaluation Study Details	69
.2.1	Screening Questions	69
.2.2	Demographic information for with-subject participants	71
.2.3	Demographic information for exploratory participants	71
.2.4	Custom Likert Scale Questions	71

List of Figures

1.1	ReSlide supports presenters manage three contextual constraints—time, audience, and communicative intent—across multiple authoring cycles. (1) <i>Sections</i> serve as intermediate units between a presentation and individual slides, grouping multiple slides around a single idea with user-defined attributes of time and emphasis. (2) <i>Jargon checking</i> provides on-demand suggestions to adapt content to audience knowledge. Sections can be reused from the <i>Slide Repository</i> (B), created and managed during authoring (A, C), and synced back for long-term use (D). Switching between the <i>Timeline Overview</i> (A) and <i>Slide Editor</i> (C) allows presenters to flexibly move across levels of detail through Sections while maintaining awareness of constraints. For long-term usage, ReSlide highlights changes in reused slides and synchronizes them with the central repository at the end of each authoring cycle.	2
4.1	Each Section is visualized as a vertical bar below the timeline, displaying its associated slide thumbnails.	16
4.2	Sections serve as a core unit of authoring in RESLIDE. Each Section represents one idea and can contain multiple slides, with time duration and emphasis level set as parameters.	17
4.3	User-defined constraints in RESLIDE. At the presentation level (A), presenters specify overall duration and audience characteristics. At the section level (B), presenters define section duration and emphasis level.	18
4.4	The Timeline Overview highlights constraint conflicts. Sections with mismatched emphasis levels and time durations are color-coded: blue indicates no conflicts, while conflicts ranging from yellow to red. If the total duration exceeds the time limit, the exceeding section(s) are marked in dark red. . .	19

4.5	At the start of an authoring cycle, the presenter configures the <i>target audience</i> (A) by selecting expertise level and providing a brief description, which is then displayed in the Editor View (a1-a2). When the jargon checker is triggered (b1), the system scans the current slide for terms that may exceed the audience’s modeled expertise. If jargon is detected, the system 1) highlights the term in its original location (b2) and 2) lists all flagged terms in a Side Panel (b3), showing each term’s expression with an on-demand definition and two suggested simpler alternatives tailored to the modeled audience.	21
4.6	RESLIDE organizes content in a hierarchical structure where a presentation consists of sections, and each section consists of multiple slides. This structure allows content in multiple granularities to be stored in and retrieved from the central Slide Repository.	22
4.7	The workflow involves switching between the Timeline Overview and the Slide Editor, enabling users to alternate between high-level structuring and detailed content editing.	23
6.1	Participants’ responses in Study 1 for comparing presentation authoring experiences with Baseline and RESLIDE.	30
6.2	Four representative user-interaction sequences, selected from both the comparative and exploratory studies. These illustrate (1) different patterns of revisiting Timeline Overview, (2) different timing of jargon checking, and (3) different strategies of content reuse (exploratory vs. intentional). The Other category denotes residual actions such as text editing, image editing, or window switching.	32

List of Tables

1	Demographic information for within-subject participants	71
2	Demographic information for exploratory participants	72

Chapter 1

Introduction

Slides are a ubiquitous medium for people to present information effectively. Recognizing the need to support effective communication between presenters and audiences, many presentation textbooks decompose the entire presentation into planning, building, and delivering phases [130, 13, 58]. Planning a presentation is a fundamental step in creating a successful presentation. It involves considering many factors that shape the presentation, such as time limits [154, 92, 153], audience expectations [154, 153], presentation settings [58], or institutional requirements [130]. We refer to these factors as *contextual constraints*, which largely influence how presenters structure and deliver their presentation content, even on the same topic. For instance, presenters must adapt their narratives and slide decks when explaining a concept to domain experts versus a general audience with limited prior knowledge.

While prior research-based presentation tools recognize the importance of certain constraints, such as time management [154, 92, 153] and audience engagement [114], they typically address these at specific checkpoints or during delivery, rather than integrating them continuously throughout presentation authoring. Some other work focuses on introducing diverse representations of data and richer forms of interactions during delivery, such as gesture-based interaction [54, 51, 30] and interaction with physical objects [163, 162]. However, existing studies offer little insight into how different contextual constraints actually shape the slide-creation process. Similarly, presentation textbooks often assume those constraints can be fully accounted for during the planning stage [130]. This gap highlights the need to explore how contextual constraints influence the entire presentation authoring process, leaving open questions on how tools might better support presenters in managing contextual constraints to craft different narratives for various scenarios.




Figure 1.1: ReSlide supports presenters manage three contextual constraints—time, audience, and communicative intent—across multiple authoring cycles. (1) *Sections* serve as intermediate units between a presentation and individual slides, grouping multiple slides around a single idea with user-defined attributes of time and emphasis. (2) *Jargon checking* provides on-demand suggestions to adapt content to audience knowledge. Sections can be reused from the *Slide Repository* (B), created and managed during authoring (A, C), and synced back for long-term use (D). Switching between the *Timeline Overview* (A) and *Slide Editor* (C) allows presenters to flexibly move across levels of detail through Sections while maintaining awareness of constraints. For long-term usage, ReSlide highlights changes in reused slides and synchronizes them with the central repository at the end of each authoring cycle.

We first conducted a formative study with ten participants to better understand the common contextual constraints and how presenters integrate them into their slide-creation workflow. Our findings identified three common contextual constraints across experienced presenters: *time limitation*, *target audience’s characteristics*, and *presenter’s communicative intent* (i.e., the main goal of the presentation). While constraints were typically recognized at the beginning of the authoring cycle, participants struggled to remain mindful of these constraints across later stages. In addition, we identified three challenges in integrating these constraints into long-term presentation authoring, specifically how presenters adapted core narratives across multiple cycles in response to different constraints. Due to a lack of visibility on constraints, it is difficult for presenters to continuously account for them while crafting the narrative and slides, resulting in redundant iteration in later stages or suboptimal delivery (C1). Moreover, presenters complained about the inefficiency in managing content distribution and ensuring smooth transitions that accurately reflected their intended story flow (C2). Finally, current tools provide limited support for managing slides across long-term authoring, leaving presenters with time-consuming, inconsistent manual management and inefficient slide reuse (C3).

To address these challenges, we propose **ReSlide**, a presentation authoring tool designed to increase presenters’ attention on contextual constraints throughout a working cycle, as well as support efficient and flexible content construction in a single presentation and long-term authoring. The system introduces a unit of authoring called “*Section*” (Figure 1.1-1), a modular layer between the entire presentation and individual slides that gives presenters flexibility in organizing and refining ideas. At the start of an authoring cycle, RESLIDE allows presenters to set constraints (e.g., presentation duration, emphasis level of each section) and provides visualizations of these constraints (Figure 1.1-C), supporting reflection on progress toward meeting them throughout the cycle. The *Timeline Overview* (Figure 1.1-A) and *jargon checking* (Figure 1.1-2) further support narrative construction and adaptation to audience characteristics. Finally, RESLIDE enables presenters to save slides and content into a *central repository* (Figure 1.1-B) and reuse them at multiple granularities, reducing redundant work and allowing flexible adaptation of existing materials. Working together, these features enable presenters to maintain contextual constraints visible during authoring (C1), create a structured overview for improved narrative coherence (C2), and facilitate flexible reuse and synchronization of content across multiple authoring cycles (C3).

We evaluated RESLIDE through two studies: an in-lab user study with 12 experienced presenters focusing on creating a single presentation, as well as an exploratory study to examine the multiple-presentation authoring with eight professional presenters. In the first study, compared to a baseline tool with basic slide authoring functionalities, RESLIDE

helped presenters turn contextual constraints from peripheral checkpoints to active design factors that guided and shaped narrative construction with greater awareness. Sections with multi-granular, multi-stage reuse encouraged presenters to shift their focus from slides to ideas, enabling iterative content development for inspiration and adaptation within a single authoring cycle. The second study further showed that the synchronization process extended flexible content reuse across multiple authoring cycles. Together, these findings suggest that presentation tools which embed contextual constraints visibility, narrative structuring, and reusable modular content management directly into the authoring flow can effectively support presentation authoring.

In summary, the main contribution of this work is threefold:

- We presented *a formative study* that revealed how experienced presenters tackled the common contextual constraints and identified three major challenges in narrative construction and content management.
- We developed RESLIDE, *a presentation authoring tool* that incorporates constraint awareness, narrative structuring, and multi-granularity reuse and synchronization into long-term authoring workflows.
- We conducted *a within-subjects study* with 12 experienced presenters and *an exploratory study* with eight professional presenters to examine the usability and effectiveness of RESLIDE through both quantitative and qualitative measures.

Chapter 2

Related Work

In this section, we review prior theories, guidelines, and systems related to handling contextual constraints, creating narratives, and reusing presentations.

2.1 Contextual Constraints and Their Management in Presentations

The idea of paired constraints originated in art and cognitive science [129, 155]. This framework combines problem-solving precedents with restructuring of problem stages, where constraints both *preclude* search in some areas and *promote* search in others [155]. Similar principles have been applied in HCI research, where constraints guide design and problem-solving in areas such as interaction design [57] and user-centered design [121], and in adaptive user interfaces that treat contextual factors as constraints shaping interface design [47].

We now turn to how contextual constraints have been recognized and managed in presentations. Prior work highlights factors such as time limits [154, 92, 153], audience expectations [154, 153], and presentation settings [58]. Presentation textbooks and training materials describe these as “issues” or “trade-offs” to be manually checked after slides are completed [137, 130, 49, 116, 31]. Building on the “restrict principle” [155], we distinguish precluding constraints (i.e., strict time limits) and promoting constraints (i.e., tailoring content to a specific audience group). Our work also examines how such constraints interact with each other during iterative authoring.

Time is a recurring constraint. A common principle was to estimate time per slide [116, 137] with systems supporting time management during slide preparation [154, 92] or during delivery [136]. Most slide creation tools adopted a slide-based time estimation method [92, 154], assigning each slide a similar amount of time while tracking authoring progress. However, though ensuring flexibility and freedom, this approach is cumbersome, especially when multiple slides support a single idea, making slide-level estimates inaccurate. TalkZone used section-level tracking during rehearsal and delivery [136], but still required fully authored slides, lacking support during authoring. Overall, though recognizing the value of multi-level time management, prior work treated time mainly as an external constraint to be monitored. In contrast, our work positions time as a design factor that actively shapes the presentation narrative and guides content structuring during authoring.

Another recurring constraint is the *knowledge gap* between presenters and audience: expressions natural to a presenter may be difficult for an audience with different backgrounds. Studies of biomedical speeches showed that slides could help audiences process dense technical content with lower cognitive load [46], while corpus analyses revealed mismatches between academic vocabulary and general comprehension [41]. Similar challenges in scientific reading have motivated in-situ knowledge support for out-of-domain audiences, suggesting opportunities to streamline content on presentation slides. For example, Qlarity [52], CiteSee [32], QuickRef [124], and Semantic Reader Project [97] simplified text through inline explanations, citation cards, or personalized visual cues, reducing search effort and aiding sense-making. Recent large language model (LLM)-based systems extend this approach by adapting content to audience needs through concept simplification, personalized jargon identification, and de-jargonization, enabling quicker understanding of difficult terms [14, 63, 120]. In addition to direct explanation, adequate metaphors could support content understanding and engagement [10]. ParseJargon [152] also showed the effectiveness of personalized explanations for jargon during an online meeting. While these works demonstrated the value of personalized support, they focused on *information receivers* after content was produced. In presentations, however, audiences have little real-time control over materials. Since presenters typically know the audience characteristics in advance, there is an opportunity for *author-side support* that enables proactive tailoring during narrative creation. Prior work, such as SlideSpace [49] and presentation textbooks [137, 130], emphasized audience awareness. Our work extends this by embedding audience adaptation directly into the authoring workflow, framing presentation construction as a user-oriented design practice.

2.2 Systems for Supporting Narrative Creation

A presentation is ultimately defined by its underlying narrative, with slides serving as a flexible medium for holding and communicating ideas. Researchers explored a range of systems to support narrative planning and storyline generation across different media beyond conventional tools such as PowerPoint¹, Apple Keynote², and Google Slides³. DataParticles introduced a block-based editor linking text, data, and visualization to support iterative narratives [27]. CodeToon demonstrated code-driven storytelling by generating comics from source code [158]. Narrative Visualization highlighted design strategies, such as ordering, messaging, transitions, and highlighting, to surface stories within data [140]. As slide decks are also multimodal storytelling media, our work examines how these principles can be translated to presentation authoring.

Another line of research has explored techniques for automatic slide generation from rich content documents. Examples include generating slides from topics for improvised TED talks [176], or converting documents into slide decks with only text elements [37] or both text and images [101, 18, 56]. Research in design and sense-making emphasizes that content creation is iterative: authors cycle between “exploration” and “exploitation” [17, 43, 126], continually reflecting on and revising ideas [87, 139]. Interactions between the material and evolving perception support new ideas and improved plans [86]. Acknowledging this, several systems proposed more controllable and workflow-integrated slide generation. Slide4N let data scientists generate slides from Jupyter notebooks by interactively selecting code cells and slide layouts [169]. OutlineSpark offered a user-controllable text outline for structuring generated decks [168], while KnowledgeDecks captured user interactions in visual analytic tools to generate slides documenting data science knowledge-discovery processes [35]. While these slide generation mechanisms gradually see the need for iterative slide generation, they typically place generation at the end of the pipeline, requiring manual editing and limiting feasibility for real presentation authoring, which also opens new opportunities for our work.

Prior works explored diverse authoring interfaces, automatic slide generation, and workflow-specific tools, but largely treated slide generation as a one-shot conversion task or outline-driven approaches. Few systems directly supported presenters in iteratively shaping narratives under contextual constraints such as time, audience, and communicative intent. Our work addresses this gap by introducing an intermediate authoring layer

¹<https://powerpoint.cloud.microsoft/>

²<https://www.apple.com/ca/keynote/>

³https://workspace.google.com/intl/en_ca/products/slides/

beyond individual slides, enabling more deliberate narrative construction with contextual awareness.

2.3 Techniques for Reusing Presentations

Rather than starting from scratch, presenters often adapt and modify existing slides into a new presentation with a different narrative [143, 25]. With conventional presentation tools, this process largely relies on manual copy-and-paste, which requires presenters to recall and locate slides by themselves. This workflow is inefficient and error-prone, motivating the development of tools supporting more effective and flexible reuse [143]. Commercial tools such as SlideLizard⁴, SlideCamp⁵, Shufflr⁶, and Empower⁷, offer centralized libraries where presentation decks can be stored and shared, while users maintain local copies. Updating to a centralized deck propagates the central deck to branch versions, allowing synchronization and version tracking across teams [82]. However, these tools mainly focused on organizational reuse and collaborative maintenance rather than supporting individual presenters during narrative construction.

On the other hand, a large body of research focused on slide retrieval and comparison. Several algorithms were proposed to help users more efficiently locate relevant slides [122, 181, 143]. Some other systems enabled users to track how slides evolve over multiple versions [45, 42]. Most recent tools expanded this space: Slide Gestalt [125] extracted structural relationships for accessibility; Elastica [30] adapted reusable visuals in live performance; Synergi [81] enabled block-level reuse in literature synthesis; and Semantic Reader Project [97] explored modular reuse in scholarly reading. However, these works focused only on retrieval, without exploring modular reuse in presentations or integrating components into evolving narratives across authoring cycles.

Several systems assisted authors in adapting existing content into new presentations. For example, NextSlidePlease [154] detected narrative structure in imported decks to support reuse during authoring. However, most approaches treated reuse as retrieving or comparing individual slides, disconnected from constructing a coherent presentation. Our work instead embeds reuse as an iterative, narrative-driven process through an intermediate authoring unit: a “section” that may span multiple slides to convey an idea. This enables presenters to save, organize, and adapt content at multiple granularities (single

⁴<https://slidelizard.com/en>

⁵<https://slidecamp.io/>

⁶<https://shufflr.com/>

⁷<https://www.empowersuite.com/en/>

slide, multi-slide idea, or full deck), aligning reuse with contextual constraints while supporting narrative construction.

Chapter 3

Formative Study

We conducted semi-structured interviews with ten experts (p1-10) to understand (1) how presenters tackled different contextual constraints in their natural workflow across multiple authoring cycles with the slide-creation tools they are familiar with and (2) what challenges they encountered during this process. The following section discusses the study details, followed by the insights and challenges, and ends with design goals rooted in the challenges.

3.1 Study Setup

We recruited ten participants through a mailing list of a local university (5 male, 5 female, between 22 to 34 years old, average 28.3 years). All participants had prior experience in presentation materials preparation and adaptation under different constraints, with an average of 7.2 years of presenting experiences after post-secondary education. Nine were doctoral students and one was a postdoctoral researcher. All participants reported their discipline as human-computer interaction (HCI), related sub-fields (e.g., information visualization, AR, VR, Haptics), or interaction design.

The study consisted of a one-hour interview session with open-ended questions about participants' general presentation workflows, their challenges and workarounds, and envisioned helper tools. The study sessions were audio recorded, which were then manually transcribed, coded, and analyzed through a thematic analysis. This study was approved by our institutional research ethic review board.

3.2 Challenges in Current Presentation Authoring Process

Based on a thematic analysis of the transcripts, we distilled challenges and design goals for supporting a more structured and smoother presentation authoring experience under different contextual constraints. Our findings are organized around the challenges that arose throughout the process. Interviewees are referred to as *p1-10* in the following.

3.2.1 C1: Fragmented and ad-hoc constraints handling.

Participants consistently noted that while starting by considering constraints, they often lost track of these factors while authoring was underway, since their attention narrowed to developing individual ideas or slides. As a result, constraints either resurfaced during rehearsal or were overlooked completely, leading to more rounds of iteration or ad-hoc adaptations during delivery. This fragmentation reduced the overall efficiency of the authoring cycle. The common contextual constraints identified were *time*, *audience*, and *communicative intent* (or the *goal* of the presentation).

Time. Participants described timing as especially important but difficult to monitor during authoring. Several explicitly called time management “*the hardest thing*” (p9, p10). Participants mentioned relying on heuristics such as the “*two slide per minute*” rule (p9) and repeatedly rehearsed in later stages to trim content (p1, p4, p7, p9, p10). Yet even with those strategies, participants still found it hard to balance the amount of content with the overall story flow (p9). This was caused by a lack of overview of the storyline, which resulted in redundant work throughout the authoring cycle (p7). These strategies demonstrated a lack of adequate time monitoring support, forcing presenters to manage those iterations manually, which made the presentation authoring experience inefficient and disrupted the story creation flow.

Audience. All participants emphasized that “*having your audience in mind is very important*” (p4), recognizing doing a presentation as inherently an audience-oriented task [49]. Presenters needed to cater the presentation to a specific audience’s needs (p1, p6), especially their knowledge background (p2). Failing to consider these could lead to comprehension challenges (p2) and misalignment between the intended message and the audience’s understanding. However, the cognitive load of balancing content creation with adaptation often lead to either overlooked the intended adaptations altogether or postponed them to later iterations—such as simplifying language, reducing jargon or technical words (p2, p3), and proactively tailoring content to the audience (p2, p5, p8). This tension highlighted the

need for maintaining awareness for the audience throughout the entire authoring workflow, not just at the planning stages.

Communicative intent. Finally, some presenters emphasized the importance of aligning the slides with an overarching communicative intent. They deliberately prioritized certain parts of the storyline over others to achieve *“the overall main goal”* of the presentation (p7), such as making a *“physical connection with people”* (p10). Participants also pointed out how slide content distribution might affect how the audience perceived *“which part was more important”* (p3). Altogether, these showed that communicative intent was not only about setting the message at the beginning, but also about actively and consistently weighing and distributing content to support that message.

3.2.2 C2: Difficulty in maintaining narrative coherence.

Maintaining a coherent narrative flow during presentation authoring was a major challenge for multiple participants. Several participants emphasized the difficulty of *“keeping in mind the main storyline”* (p7) and double-checking *“if it flows nicely or not”* (p9), even with the rough initial outline developed before building slides. In addition, some also mentioned the one-idea-per-slide strategy (p2, p6), but it quickly became unmanageable as slide counts grew large (p6). Despite those strategies, participants still found it difficult to decide what to include, exclude, or emphasize (p9).

Maintaining coherence also required attention to distributing and transitioning across sections (p4). As one noted, *“I need to make sure that the transition between sections is consistent”* (p9). Additionally, participants found the content distribution should correctly signal importance, as *“[inappropriate allocation] will cause the audience to lose their sense”* (p3). Achieving this required carefully *“go[ing] through all the slides and examining content”* (p1), which caused *“a lot of redundant work”* (p7) and drifting away from the main point (p10). Lacking visibility support offloaded the consideration to presenters, resulting in a long iteration cycle or overlooked factors during delivery. To gain *“a good overview of the storyline”* (p7), they expressed the need to visualize and manage both high-level and lower-level transitions (p2, p7).

3.2.3 C3: Inefficient content management and reuse.

“Borrow[ing] slides from the [prior] presentation slides” (p8) was a common strategy across all participants. Participants described reuse as a multi-level, multi-stage action while carefully crafting a new storyline. Some started with *“copying and pasting my entire old*

slide deck” (p7, p9), then selectively kept those that could be reused or readapted, even though sometimes *“the storyline may not align that well”* (p7). Others explained reuse as an on-demand action, reusing smaller units like groups of slides (p7, p8, p9), a single slide that *“conveys one piece of [context independent] information”* (p7, p9), or a single element (p10). This demonstrated that reuse happened throughout the workflow, not defined and fixed at the beginning.

At the same time, participants maintained a long-term slide repository such as *“the main deck”* (p8), or a form of *“version controlling”* (p10) to support future reuse. One described using Apple Keynote as a *“kind of like external library”* (p10). However, manually synchronizing the newer slide version back to the main deck was difficult: *when I make changes to the specific versions, I hope it applies to the main deck as well* (p8). Others also highlighted the need for certain flexible granularity, like *“slide-based”* (p9) or *“element-based”* (p7, p10), which underscored both the effort required and the limitations of using the current tools.

Together, these challenges revealed that adequate constraint management and content reuse in single presentation and across multiple authoring cycles are important to how presenters adapt their narratives in response to various contextual constraints over time. Yet current tools treated these as external factors to presentation authoring, forcing presenters to rely on ad-hoc workarounds that were often redundant, time-consuming, or inconsistent. The key is not whether reuse happens, but how to integrate it seamlessly at multiple granularities within the authoring workflow.

3.3 Design Goals

To tackle the identified challenges, we derived the following key design goals, each corresponding to a challenge (C1-3), to inform the design of our presentation authoring tool, RESLIDE.

DG1: Enable proactive management of contextual constraints along presentation authoring. To improve the overall efficiency, authoring tools should actively assist presenters in externalizing and organizing their thoughts while accounting for relevant constraints. This may involve providing subtle reminders, visualizing the current authoring status, and helping presenters manage these factors as needed. By bringing those constraints to the front, we aim to increase awareness and encourage presenters to plan and revise narratives with more consistent alignment to those contextual constraints.

DG2: Facilitate structured narrative construction based on idea pieces and

relationships. Tools should enable fluid navigation across multiple levels of abstraction, allowing presenters to decompose a narrative into sub-sections, group slides together, and assign appropriate emphasis level to each section. Supporting such structured mappings would help presenters maintain coherence between evolving slides and the main intended narrative flow.

DG3: Support flexible reuse and synchronization of contents at different granularities. To address the issue of adapting slides from prior presentations, slide creation tools should be able to assist content reuse and synchronize at multiple granularities, such as slides, text, and graphics. It should also facilitate modular management, and offer smooth synchronization across different versions. Enabling such support would effectively integrate both a single authoring cycle and long-term content management into a natural slide authoring workflow.

Chapter 4

ReSlide

In this section, we introduce **RESLIDE**, a presentation authoring tool that enables a smooth and considerate experience by aligning contextual constraints with evolving presentation narratives, developed based on the above design goals.

4.1 System Overview

RESLIDE was implemented with a React frontend and Node.js backend. The frontend user interface consists of three views: 1) a *Timeline Overview* (Figure 1.1A), a *Slide Repository* (Figure 1.1B), and a *Slide Editor* (Figure 1.1C). **RESLIDE** introduces *Section* (Figure 1.1-1) as the core unit of authoring, which is a group of consecutive slides exhibiting a semantically coherent idea. This structure gives presenter more flexibility in managing content while focusing on building ideas (**DG2**). The Timeline Overview presents the narrative structure of the presentation, including all created sections with accumulated time (**DG1**, **DG2**). Below the timeline, each section is visualized as a vertical bar (Figure 4.1), with its associated slides grouped inside (**DG2**). The Slide Repository lists reusable resources at multiple granularities (presentation, section, slide), enabling content retrieval and reuse (**DG3**). This promotes long-term flexible usage by offering five options to synchronized reused slides with prior versions (Figure 1.1D). The Slide Editor provides a workspace organized by sections and incorporates jargon checking (Figure 1.1-2) to adapt content to target audiences, supporting both constraint awareness and narrative construction (**DG1**, **DG2**). Together, these features enhance presenters' awareness of contextual constraints alongside the natural narrative construction workflow across multiple presentation authoring cycles.




Figure 4.1: Each Section is visualized as a vertical bar below the timeline, displaying its associated slide thumbnails.

The backend of RE_SLIDE comprises two main components: an *API layer* and a *Jargon Detection Module*. The API layer manages the storage and retrieval of reusable presentation content by communicating with MongoDB for slide content and ASW S3 for image data. The Jargon Detection Module analyzes the audience information to identify potential jargon in the current slide content. This component embeds GPT-4o-mini via the OpenAI API, applying prompt engineering to generate simplified alternatives tailored to the specified audience. The processed outputs are relayed to the frontend components for visualization and interaction.

4.2 Define Unit of Authoring with Section

To help presenters effectively manage contextual constraints (**DG1**) and organize content while constructing a storyline (**DG2**), we introduce the idea of Section. This is inspired by agenda or overview slide, a common practice where presenters provide the audience with a roadmap of the presentation. RE_SLIDE’s Section aligns with the presenters’ natural mental model of how a presentation flows, where one section represents a single idea and




Figure 4.2: Sections serve as a core unit of authoring in REsLIDE. Each Section represents one idea and can contain multiple slides, with time duration and emphasis level set as parameters.

may contain multiple slides (Figure 4.2). REsLIDE proposes a section-first workflow, in which the users first create sections, and then adds slides within them. This mirrors how presenters mentally group content and supports aligning ideas with constraints.

Presentations usually come with a certain time limit, which is one key contextual constraint we identified. Although the given limit applies to the entire presentation, our formative study found that only considering the overall limit is too vague for presenters, resulting in more iterations in later stages. When the user creates a new presentation, REsLIDE requires them to set the total time (Figure 4.3). Additionally, when creating a section, the user can set a time duration, estimating the duration this part may occupy in the total presentation, otherwise a default time of 2-minute is applied.

Presenters sometimes do not know how long a section should take in the actual presentation at the beginning. But they usually have a rough sense about whether this idea is important to the presentation, or to the audience, which is related to the audience and the communicative intent constraints (Figure 4.3). To better assist the users in mapping their intended storyline with the actual presentation content, REsLIDE also allows them to set an emphasis level to a section. There are three levels of emphasis—low, medium, high—indicating how important that section is over the entire narrative.

The time duration and emphasis level of sections work together. Sometimes there might be a conflict between these two settings when a high-priority section gets less time than a low-priority section. For example, if the Key Result section (high priority) is allocated two minutes while the Conclusion section (low priority) is given four minutes, this creates




Figure 4.3: User-defined constraints in RESLIDE. At the presentation level (A), presenters specify overall duration and audience characteristics. At the section level (B), presenters define section duration and emphasis level.

a conflict. Those sections are then be highlighted, as a soft indication, reminding presenters to reconsider the balance between timing and narrative flow (Figure 4.4). There are three levels of indication represented by three colors, ranging from yellow to red, indicating an increasing conflict; blue means no conflict. A yellow mark signals low conflict when an important section receives about the same time as a less important section; an orange mark indicates medium conflict when an important section receives 75% of the less important ones; a red mark shows high conflict when an important section receives only 50% or less than the less important ones. These indicators ensure that time allocation and communicative intent remain aligned during presentation authoring.

4.3 Support Narrative Construction with Timeline and Jargon Detection

Based on authoring units exhibited as sections, RESLIDE allows for a proactive constraint management (**DG1**) and a structured narrative building (**DG2**) with the interactive Timeline Overview and jargon checking feature.

To accommodate different workflows, RESLIDE allows sections to be added individually,




Figure 4.4: The Timeline Overview highlights constraint conflicts. Sections with mismatched emphasis levels and time durations are color-coded: blue indicates no conflicts, while conflicts ranging from yellow to red. If the total duration exceeds the time limit, the exceeding section(s) are marked in dark red.

in bulk, or as placeholders for rapid outlining. Once created, sections are displayed in the Timeline Overview (Figure 1.1A), providing a high-level context of the narrative structure while surfacing potential conflicts. Presenters can refine the narrative by dragging to reorder or adjusting their duration directly. Slides can also be organized and moved across different sections by drag-and-drop. Together, the timeline visualization allows presenters to manage detailed slide content without losing control of the overall structure. Following the presenters' mental model, the Timeline Overview (Figure 1.1A) and the Slide Editor (Figure 1.1C) effectively externalize the presenters' mental workload while supporting a clearer storyline development experience.

Adapting slide content to the audience was another recurring challenge identified in the formative study. Upon creating a new presentation, RESLIDE also requires users to describe the intended audience (Figure 4.3). RESLIDE models the audience in two ways (Figure 4.5A-a1):

- *Expertise level*: Presenters can categorize their audience on a five-point scale, ranging from 1 (Novice) to 5 (Expert). A higher number indicates greater prior knowledge about the topic.
- *Audience description*: Presenters can write a short description to capture extra infor-

mation not represented by the expertise level.

These settings guide the jargon detection in RESLIDE. The lower the audience expertise the presenter sets, the more detailed and strict the system checks for potentially difficult expressions, languages, etc., correspondingly. RESLIDE first expands the user’s input audience description and expertise rating into a detailed profile that includes background, inferred expertise, known concepts, likely jargon areas, and domain context. This profile is then used in a second analysis of slide content, detecting jargon and suggesting tailored alternatives or definitions while considering user characteristics and presentation context. Both steps prompts are implemented using ChatGPT API calls, and the prompts are listed in [Section .1](#).

The jargon checker can be applied at the slide level and triggered on demand ([Figure 4.5B](#)). When activated, the system checks the current slide and reports each potential jargon by: 1) highlighting the jargon in its original position ([Figure 4.5B-b2](#)), and 2) displaying the expression, the on-demand definition, and two suggested simpler alternatives tailored to the modeled audience on a side panel ([Figure 4.5B-b3](#)). Presenters are also allowed to hide individual words or all flagged terms if they find the suggestions not useful, or if they have already resolved the suggestions.

4.4 Enable Multi-Layered Presentation Reuse and Synchronization with Central Repository

To support long-term efficiency and reuse (**DG3**), RESLIDE introduces a central Slide Repository. Content can be managed at multiple granularities—presentation, section, and slide—where each has consistent saving and reusing behaviors ([Figure 4.6](#)). This reflects how presenters naturally manage their presentation structure while providing flexibility across different knowledge scales. Saving the higher-level structure preserves its connection with the lower-level structure.

At the presentation level, the complete deck, including all sections with their corresponding metadata (e.g., duration, emphasis), is stored in the repository. A saved presentation can be re-imported into the workspace as its full narrative, enabling presenters to adapt to new contextual constraints. Similarly, at the section level, sections and their metadata can be stored, and restored as a copy to a new workspace. Thus, presenters can then adjust the newly-imported section to fit a new storyline. Finally, at the slide level, each slide is linked to all of its previous versions. Presenters can update individual components (e.g., text boxes, figures) on a slide. Upon saving, the system checks for




Figure 4.5: At the start of an authoring cycle, the presenter configures the *target audience* (A) by selecting expertise level and providing a brief description, which is then displayed in the Editor View (a1-a2). When the jargon checker is triggered (b1), the system scans the current slide for terms that may exceed the audience’s modeled expertise. If jargon is detected, the system 1) highlights the term in its original location (b2) and 2) lists all flagged terms in a Side Panel (b3), showing each term’s expression with an on-demand definition and two suggested simpler alternatives tailored to the modeled audience.

changes and asks the presenter to synchronize slide versions via a Slide Comparison Panel (Figure 1.1D). There are five options to resolve the synchronization across versions:

- *Ignore Changes*: discard all changes and keep the prior version.
- *Set as Origin Slide*: save the new slide as the starting point of a new lineage.
- *Keep Both Versions*: save the new slide while keeping lineage.
- *Replace Content*: overwrite the selected prior version(s) with the new slide.

Later, an individual slide can be retrieved through a keyword search. Selected slides should be placed into a chosen section. Hovering over slides in the repository shows an enlarged preview that assists with browsing and selection.




Figure 4.6: RESLIDE organizes content in a hierarchical structure where a presentation consists of sections, and each section consists of multiple slides. This structure allows content in multiple granularities to be stored in and retrieved from the central Slide Repository.

4.5 Example Workflow

Below we present an example workflow to demonstrate some of the features described above. Alice, a psychology PhD student, is preparing a ten-minute talk on the topic “Why multitasking feels productive,” a topic relevant to one of her publications. Her audience includes middle to high school students, parents, and community members. Alice begins by launching RESLIDE, setting her total presentation time (10 minute) (Figure 4.3A), entering the audience description—high school students, parents, and community members who are curious about this topic but may have a range of backgrounds in it (Figure 4.5 a1-a2)—with expertise level to three, reflecting their limited prior knowledge.

Then, she outlines the storyline by creating sections, assigning emphasis levels, and adjusting their duration (Figure 4.2). For example, she includes an Introduction and Conclusion with no emphasis, as these are not central to her main point. Between them, she adds three body sections (Figure 4.1): one defining multitasking (low emphasis, 2 minutes), one explaining the illusion of increased productivity (high emphasis, 3.5 minutes), and one discussing daily implications (medium emphasis, 2.5 minutes).

To construct her slides efficiently, Alice strategically reuses the Introduction section




Figure 4.7: The workflow involves switching between the Timeline Overview and the Slide Editor, enabling users to alternate between high-level structuring and detailed content editing.

from a prior presentation to peers that is stored in her repository (Figure 1.1B). She also incorporates several content slides from that presentation into her explanation section. While drafting, she activates the jargon checking, which flags the term “Heavy Media Multitaskers (HMMs)” and suggests two alternatives: “frequent media users” and “people who multitask with media” (Figure 4.5 b1-b3). She thus integrates these revisions to ensure content clarity for her audience. Once she done with a section, she switches back to the Timeline Overview to ensure she stayed aligned with the constraints (Figure 4.7). Once all slides are completed, Alice saves and synchronizes (Figure 1.1D) them back to the central slide repository (Figure 1.1B), making them available for future reuse.

Chapter 5

Evaluation

We conducted two in-lab user studies to evaluate how RESLIDE supports presentation authoring with contextual constraints, aiming to answer the following research questions (RQs):

- **RQ1:** How does RESLIDE support presenters in *managing contextual constraints* during slide authoring?
- **RQ2:** How does RESLIDE support presenters in *structuring and maintaining narrative coherence* during authoring?
- **RQ3:** How does RESLIDE’s support for *modular reuse and synchronization* influence presenters’ authoring practices?

Study 1 employed a within-subjects design, assessing how RESLIDE supports presentation authoring compared to a baseline system. It primarily focused on **RQ1** and **RQ2**, and partially explored reuse behaviors relevant to **RQ3**. Study 2 was an exploratory study where participants used RESLIDE across two authoring scenarios, which simulated multi-session authoring practice, allowing us to further examine **RQ 1-3**.

5.1 General Study Setup

We conducted both studies in person after receiving approval from the university’s ethics review board. Study sessions were audio and video recorded, lasted around 100-130 minutes, during which participants completed assigned tasks, post-task questionnaires, and a

semi-structured interview. Participants were compensated with \$40 CAD. For our presentation authoring tasks, we created a pool of four topics that was accessible to most people, preventing potential presumptions about participants’ backgrounds:

- Understand everyday inflation behind rising grocery prices?
- How does social media impact mental health?
- How AI and personalized suggestions shape what you see online?
- Why multi-tasking feels productive, and is it really the case?

5.2 Study 1: Comparative Study with Single-Session Authoring

We compared RESLIDE with a baseline system that mimics current slide authoring tools (e.g., Google Slides, Microsoft PowerPoint, Apple Keynote). To ensure fairness between the two systems and minimize bias from participants’ varying familiarity with existing tools, we built a customized baseline that provided only basic slide authoring functionalities (i.e., text and image) and slide reuse. The authoring functions were identical in both systems. In the Baseline, participants could reuse slides via copying and pasting, which is similar to how slide reuse works in commercial presentation authoring tools.

5.2.1 Participants

We recruited 12 experienced presenters (three males, nine females; between 23 and 33 years old, average 26 years) via social media and mailing lists. Participants were asked to self-report their slide-preparation and presentation skills through a set of screening questions. All participants were familiar with at least 1 slide creation tool. Ten out of 12 participants have received training in public speaking or presentation skills, while the 2 who have not received training reported frequent presentation experiences in the past 3 months or had adequate experience of presenting to mixed audiences. The screening questions and the detailed demographic information are listed in [Section .2](#).

5.2.2 Task

Participants were asked to create a presentation based on a given scenario under a specified topic selected from the topic pool. During the task, participants had access to: a short scenario description, some supplementary materials to help them familiarize themselves with the topic, a slide deck previously created for a different scenario on the same topic, and space for notes or drafting. We encouraged presenters to adapt the provided slide deck while creating the new storyline. However, if the participants felt an extra burden while adapting that slide deck into the current slide, they could also not reuse any of the reusable slides.

5.2.3 Procedure

After signing the consent form, participants were given an overview of the study procedure, duration, and data collection details. The study was conducted in person, with participants accessing both systems via a research computer. A brief tutorial, followed by a brief training session that included a set of warm-up tasks, going through all the features, was provided to demonstrate the use of both systems. Participants could familiarize themselves with the tools before starting the tasks.

Participants were instructed to pick two topics from the topic pool that they were more familiar with. Then they used the two systems (RESLIDE and Baseline) to perform two presentation authoring tasks, each corresponding to a different topic they selected. The order of the systems was counter-balanced across participants. Participants were given 30 minutes to complete one task. A countdown timer was provided, and the researcher also gave time indications every 10 minutes after the task started. After each task, participants rated the system usability, cognitive load, intent expression, and slide authoring experience with the system on a 7-point Likert Scale. Upon completing two tasks, a semi-structured interview was conducted to gather feedback on participants' experiences. Participants were encouraged to share general comments on any aspects of the study and then prompted on specific aspects, including interface usability, pros and cons of the systems, as well as prior difficulties and strategies tackling the contextual constraints and reuse slides. They were also asked about the differences between systems' impact on their approaches to solving their raised strategies and suggestions for improvement. Observations noted by the researcher during the session were also discussed.

5.3 Study 2: Exploratory Study with Multi-Session Authoring

This study was designed to explore the multi-session authoring practice of RESLIDE. For this purpose, we created two scenarios: (1) a public talk at a local library (time: 5-10 min, audience: general audience, communicative intent: explaining concept and application), and (2) an academic conference presentation (time: 10-15 min, audience: domain experts, communicative intent: discussing technical details).

5.3.1 Participants

We recruited eight experienced presenters (four males, four females; between 25 and 39 years old, average 30 years) via social media and mailing lists. All participants were familiar with at least one presentation authoring tool. They were more experienced presenters who satisfied at least two of the following: Have done 6–10 presentations in the past 3 months; Have had formal training before; Have been presenting for 5+ years since starting post-secondary; Have usually presented to a large audience (e.g., more than 50 people). Their detailed demographic information is listed in [Section .2](#).

5.3.2 Task

Participants created slides on one topic of their choice from the topic pool. Unlike Study 1, participants were instructed to create presentations for the above two scenarios on the same topic, enabling the reuse of their own slides in the second round. During the first task, participants had access to: a short scenario description, some supplementary materials to help them be familiar with the topic, and space for notes or drafting. In the second task, besides the above items, participants also had access to the slides they created.

5.3.3 Procedure

The beginning of the procedure for Study 2 was the same as that for Study 1, until the actual study tasks. Participants were instructed to build two presentation slides on the same topic, each under different constraints as described above. They were encouraged to think-aloud during the presentation authoring. After each authoring task, participants rated their experience with the system on a 7-point Likert Scale. Upon completion of

two tasks, we conducted a semi-structured interview to gather feedback on participants' experiences with RESLIDE. The questions in the interview were similar to those in Study 1, and observation notes were discussed.

5.4 Data Collection and Analysis

Evaluating presentation quality is inherently subjective, especially when participants are instructed to focus on storyline construction. Therefore, we utilized a combination of standardized surveys, custom self-report ratings, interaction logs, and interviews to build a comprehensive understanding. To assess usability and cognitive effort while creating a presentation, we leveraged the UMUX-Lite and NASA-TLX in Study 1 and analyzed the results with the Wilcoxon signed-rank test. To evaluate support for contextual constraints (**RQ1**), narrative coherence (**RQ2**), and reuse workflows (**RQ3**), we designed a set of custom questions. Those questions fell into three categories (constraints, coherence, and reuse), capturing participants' subjective impression while doing tasks (see [Section .2](#)). In both studies, we logged participants' interaction data to understand how they managed their constraints and structured their slides. We analyzed the interview data using a thematic analysis, employing both inductive and deductive approaches. We open-coded the transcribed interviewees' responses, employing affinity diagramming to sort the initial codes. Through iterative discussions and the organization of these codes, we identified a number of recurring patterns and themes.

Chapter 6

Results

We begin by reporting the questionnaire results in Study 1 ([Section 6.1](#)), and then present the qualitative themes from both studies that examine how RESLIDE supports presenters in managing *contextual constraints* (RQ1), maintaining *narrative coherence* (RQ2), and enabling *modular reuse and synchronization* (RQ3), which elaborate five themes ([Section 6.2-Section 6.6](#)). Participants in Study 1 are referred to as *c1-12* and those in Study 2 are referred to as *e1-8*.

6.1 Questionnaire Results from Study 1

Although seven of 12 participants reported lower perceived workload on the NASA-TLX for RESLIDE, the difference between Baseline and RESLIDE ($Mdn_B = 3.33$ vs. $Mdn_R = 3.42$, $p = 0.414$) was not significant. This suggests that our system did not add cognitive or physical burden while making presentations, even though it offered many new features. The system usability scores computed from UMUX-LITE were significantly greater ($p = 0.021^*$) for RESLIDE ($Mdn_R = 6$) compared to Baseline ($Mdn_B = 5.25$). Participants consistently reported significantly better results with RESLIDE on all subjective metrics ([Figure 6.1](#)), from Contextual Constraints ($Mdn_B = 4.4$ vs. $Mdn_R = 5.9$, $p = 0.0024$, $r = 0.815$), Narrative Coherence ($Mdn_B = 4.16$ vs. $Mdn_R = 6.16$, $p = 0.0143$, $r = 0.706$), to Reuse ($Mdn_B = 5.167$ vs. $Mdn_R = 6.667$, $p = 0.0203$, $r = 0.667$).




Figure 6.1: Participants’ responses in Study 1 for comparing presentation authoring experiences with Baseline and ReSLIDE.

6.2 Externalizing Constraints Made Intent More Tangible and Actionable (RQ1 + RQ2)

Baseline led to a freeform authoring workflow similar to participants’ prior experience with existing tools, which prioritized content over constraints. Participants relied on their own judgment to keep content aligned with constraints, resulting in inefficiencies or ignoring the constraints altogether during delivery.

With ReSLIDE, participants actively externalized constraints by “*setting [them] as input for the interface*” (e7), assigning initial time estimates to each section in advance (c12), and later performing a final check to ensure alignment (e1). Although some noted difficulties in estimating time or importance until the slide deck was complete, they found that introducing constraints early effectively highlighted this information during planning (c7, e1, e7). This set a solid foundation for the subsequent storyline development.

The Timeline Overview further supported planning by combining “*importance of each part*” (c1, c11) with the time allocation. It provided an intuitive and effective way to help participants “*stay mentally aware and structure slides within the time limit*” (c1, c5, e3, e7), allowing them to quickly “*glance [at] which part [takes] how many minutes*” (c8, e8) while “*breaking down the topic into subsections*” (c8, c9). Some participants admitted that in Baseline, they unintentionally ignored the audience information, whereas setting the audience in ReSLIDE effectively reminded them throughout the authoring process (c7, e7).

6.3 Structuring with Sections Shifted Presenters’ Authoring Focus from Slides to Ideas (RQ2)

No matter which tool they used, participants generally started with *“a rough outline with high-level ideas”* (c1, c4, c7, c10, e2, e3, e4, e6, e7). They perceived creating a presentation as the process of connecting ideas, where an idea served as a basic building block in their mental workflow (c10).

Baseline worked well when participants already had a moderately defined storyline in their mind (c8). While ensuring smooth transitions between ideas was important, participants also emphasized the need to include all the necessary ideas in the presentation (c4). With RESLIDE, participants used Section to bridge their mental narratives and slides, enabling them to connect and transit between ideas at different granularities (c10, c12, e2). They noted that Section helped them to *“think from the section level, not just the slide level”* (c12, e6), which *“matched their mental model”* (c10, e8) and helped them visualize the flow of high-level ideas in the presentation (c1, c8, c12).

6.3.1 The Timeline Overview Assisted Idea Reflection and Organization

Beyond the section structure, participants valued the ability to move smoothly and cohesively between high-level ideas and low-level details (c6). Defining sections with descriptions and later reorganizing them at both the section level and the slide level was a *“process of sorting out mental models of how the story should flow”* (e3, e5, e7). The Timeline Overview supported this process by making the section distributions visible along editing, helping participants track how ideas are constructed and flowed (c8, e5). They found it useful to keep thoughts organized across two separated views (e5), where one focused on structure (Figure 1.1A) while the other focused on details (Figure 1.1C). This structure made the idea *“visually easier to comprehend”* (c5, c8), especially when the slides was *“too complex to just have [the structure] in [their] head”* (c3, c8, c12).

By contrast, the absence of a clear layout and visual indicators in Baseline made navigation (c12) difficult, even when they developed their own workarounds such as *“manually defin[ing] the section slides”* (e6, c12) or drafting outlines (c11, c12). These strategies could not help them distinguish sections in the middle part of a presentation. To improve the expressivity, participants suggested a *“multi-layered grouping”* (e1, e5) with a restricted *“tree structure”* (e5, c7) and a timeline to enable clearer connections between ideas, enhancing flexibility while maintaining restrictions. However, manually switching between




Figure 6.2: Four representative user-interaction sequences, selected from both the comparative and exploratory studies. These illustrate (1) different patterns of revisiting Timeline Overview, (2) different timing of jargon checking, and (3) different strategies of content reuse (exploratory vs. intentional). The Other category denotes residual actions such as text editing, image editing, or window switching.

these two views increased mental workload and disrupted authoring flow (c7, c11, c12, e5). One suggested workaround was to show high-level timeline information and layered storyline structure on the Editor view, while also allowing access to detailed notes for specific sections (c8).

6.3.2 Section-first Workflow Encouraged Narrative-centered Authoring

Starting from creating a section, participants adopted different strategies when developing their narrative within this section-first workflow. Some started from a top-down (c1, e6, e5) approach, while others used a bottom-up approach (c6, c12, e7); then they reflected on the storyline as a whole by seeing from each other’s perspective (c4, c6, e5). Participants who were used to starting from big pictures appreciated that they could start off from “*high-level story chunks*” then “*put details in*” (c1) while adjusting and reorganizing chunks later (c1). Although participants admitted that the sections might be changed later, they still liked the ability to externalize and manage their initial thoughts at a higher level (e3, e4). For participants who went from bottom-up, they listed all the available ideas as sections first (c12), then organized ideas into a coherent storyline: “*Section actually comes very late in my presentations*” (c6). In this case, instead of constraining their thoughts, the role of the Timeline Overview was to lay out the existing ideas so that they could then reorganize based on their evolving storyline (c7, e2).

6.4 Visualizing Time and Emphasis Effectively Guided Presenters to Resolve Constraints (RQ1 + RQ2)

Building up from where externalizing constraints at the beginning increased the awareness of constraints, RESLIDE continued to assist participants in managing constraints through the Timeline Overview over Baseline. Initializing audience characteristics directly influenced participants to consider the appropriate emphasis for each section with narratives, which enabled participants to “*estimate how many slides I should make*” (e6). The Timeline Overview helped participants to tailor slides to the target audience (c4, c5) by aligning the slide content with their intended strategy as a subtle reminder (c7). Interaction logs confirmed the pattern, showing that multiple participants (Figure 6.2) set an initial timeline at the start, regularly returned to the overview during authoring, and performed a final adjustment at the end to ensure their deliverable satisfied all constraints.

6.4.1 Reflecting on Sections Enables Accurate and Flexible Time Estimation

Considering the total time limit, when participants had a longer section, they “*had to remove some of the other content for the sake of the [total presentation] time*” (e5). In Baseline, where participants did not have access to the Timeline Overview, they estimated their total presentation time on a per-slide basis (c2, c8), where “*one slide is one minute of content*” (c8). However, they also reflected that this approach did not work in reality based on their own experiences (c2, c8, c9), as sometimes it was hard to “*split [the idea] out into slides*” (c2). Participants executed the story flow by connecting multiple ideas. Instead, RESLIDE “*forced [presenters] to think more about how long groups of slides should take*” (c2, c5), adding another time estimation approach from a higher-level point of view, which “*improved time efficiency*” (c9). Meanwhile, it still preserved the slide-level time estimation via the “*vertical visualization for each slide*” within each section (Figure 4.1), which triggered presenters to “*have the awareness of this [section] within how many minutes*” (c3, c5, c11).

6.4.2 Leveraging Visual Indicators in Timeline Overview to Prevent Overloading and Unbalanced Idea Development

Besides the structure, visual highlights (Figure 4.4) effectively reminded presenters to stay mindful of time and idea distribution when the current slide status conflicts with the pre-

defined restrictions (e5, e7), allowing participants to prioritize more important sections and keep content aligned with their planned timing (e2, e8). If there were too many sections that could not fit within the time limit, the highlights from Timeline Overview triggered participants to revise their initial plan, such as “*reducing section time*” (c11), “*reorganizing sections*” (c11), “*merging two sections together*” (e5), or simply “*deleting sections that are not so important*” (c2, c4, e4). If there were not enough content to make up for the total presentation time, participants would consider “*adding new sections*” (c1, c2, e3), “*redistributing time to different sections*” (e4), or simply “*revising the total presentation time*” (c8) within an appropriate range. The Timeline Overview also effectively assisted participants in reviewing “*how many pictures are [in the presentation]*” (c1), helping them with alignment between visuals in the presentation and audience needs within total time constraints (c9). Participants mentioned that it was hard to “*control the timeline and the level of importance by [themselves] in [their] mind*” (c3) using Baseline (c2, c8, c9) without these higher-level cues. Since they found it easy to lock into developing one idea and ignore other ideas (c7, c8), they “*consistently underestimated the time spent on each idea*”, resulting in much longer presentation time or rushed delivery.

6.5 Jargon-checking Prompted Reflection and Inspiration (RQ1 + RQ2)

Participants described different strategies for integrating the jargon checking into their workflow. In some cases, participants utilized jargon-checking as an estimation of whether the slides were too difficult for the audience, then they manually made adjustments in both content and language (c6) without taking the suggestions. For an audience with lower perceived knowledge, participants utilized the jargon checking to “*simplify words*” (c5, c8, c9, c11, c12, e2, e3, e5), reducing their own effort of manually checking while keeping the context. When the audience was at a similar or higher knowledge level than themselves, they either felt adjusting terminology was unnecessary (e2), or wanted it to provide “*more professional terms*” (c9, e5, e7). Others found it helpful as a source of inspiration when they located the jargon by themselves, but could not remember the term (e5) or struggled to rephrase in their own words (e1, c7). Some applied it selectively to specific slides, while others preferred to apply it to all slides at once if possible to maintain their mental flow during editing (c7). Interaction logs revealed two distinct usage patterns. Some participants ran jargon checking immediately after drafting a slide (Figure 6.2, c9, e7), resulting in sparse checks distributed throughout the authoring cycle. Others postponed it until the entire draft was complete (Figure 6.2, c10), leading to more frequent checks at

the end.

While appreciating this feature, participants expressed a need to expand its functionality. Suggestions included broadening the context to make recommendations—all previous slide content and other specified constraints (e2, e5, e8)—and tailor them to audience characteristics such as language proficiency (e3, e5), cultural background (e3, e5), general education level (c5, e6), age (c5, c9), intent, and interest (c6). Some participants envisioned a “*two-way jargon checking*” that compared their own knowledge level with that of their audience (e7). Participants also asked for suggestions beyond text difficulty, including tone, phrasing, section allocation, story structure, topic choice, and even the slide content such as alternative images (c12, e2, e3, e4, e5, e7). For instance, younger kids required playful language and accessible visuals such as a “*nice, pretty cat*” to sustain engagement (c12, e7), a general audience preferred more visually appealing images (c1, c12, e3), and experts wanted more complicated and well-illustrative diagrams or figures from paper (c12, e8). Similarly, some audiences desired longer discussion sections with shorter introductions and conclusions (c4), whereas children’s limited attention required shorter, more engaging presentations (c12).

6.6 Modular Reuse and Synchronization Supported Flexible Authoring Across Versions (RQ2+ RQ3)

Participants found that RESLIDE effectively supported them in reusing slides. While reuse was possible in Baseline, several participants ended up with copying all existing slides into the workspace regardless of the content, then deleting them all afterwards (c5, c6). In contrast, they explicitly expressed their preferences towards the reuse in RESLIDE (c5, c6, c10). In reusing prior decks, the interaction logs showed distinct strategies that informed participants’ long-term workflows. (1) Exploratory reuse: participants browsed existing slides for ideas and inspiration (Figure 6.2, c9, e7, e8). This was characterized by multiple filter actions without actual reuse, which is either concentrated at the start (e7, e8), or spread throughout the authoring cycle (c9). (2) Intentional reuse: participants adapted targeted content for the new presentation (Figure 6.2, c9, c10, e8), typically involving multiple filtering actions followed by one or more reuse actions. (3) Versioning and synchronization: participants saved and aligned materials across multiple iterations to manage evolving content. We elaborate on these strategies below.

6.6.1 Exploratory Reuse for Idea Generation and Inspiration

“Cold start” was a common problem participants encountered (c5, c8, c12, e5), which reduced their slide authoring efficiency. Previous slides were a resource that they treated as a base of the new presentation, helping “*reducing the mental burden*” (c5, e3, e4). This cold start problem happened at the beginning when facing the entire presentation (c5, c8, c12, e4, e5), or when they lost inspiration for “*the middle sections*” (c12). Depending on which stages participants felt lost, the units they reused were also different. For the entire presentation, some participants were unsure about the content, so they simply added all existing slides before planning (c12, e3, e5). They also generally reused the “*most recent presentations because they reflected more of my taste at the current moment*” (e3). For the meat of a presentation, they often reused a specific section that represented one particular idea, differentiating the content only in terms of phrasing (c12). In some cases, participants skimmed through existing relevant content only to “*look for idea collections*” for inspiration for their own storyline and structure (c8, c9).

6.6.2 Intentional Reuse for Targeted Content Adaptation

In other cases, participants started with a rough mental model and reused content with a clear target. Participants closely examined whether the existing presentations were helpful in terms of how much effort was needed to adapt the content, and used the ones “*[they] didn’t have to change much*” (c3, e1). This reuse happened at multiple granularity levels. At the narrative level, participants reused the previous presentation, or the entire structure if they found the current storyline was “*exactly the same or similar to*” what they had before (c6, e7, e8), then adapted the diverging part to fit into their intended story flow and section allocations (e5, e7). If they saw that the storylines were extremely different using the Timeline Overview, they instead started from scratch. In this case, the unit of reuse came down to sections or slides (e2, e8). Participants started with adding section(s) into the high-level storyline (c10, e3, e8), roughly configured the section parameters in a way that fitted their mental model, then restructured the sections in detail, such as adding or adjusting slides (c10, e6, e7), or changing expressions to be more appropriate to the target audience (e6). Moving down to the slide level, the reusable elements became more specific, such as the introduction slides (c11), graphics (c9, c12), and textual components (e6). During authoring, participants also regularly opened the Central Slide Repository to skim the previous slides, but used it as “*a guideline to double-check if what [they were] doing is correct*” (c9).

6.6.3 Saving and Synchronizing Across Versions to Manage Evolving Content

Participants saved and synchronized slide content across the entire authoring process, which also included multiple granularities. When participants finished slides for one presentation, they intentionally saved their presentation content to the central slide repository for future reuse. While the provided four options gave participants enough flexibility (i.e., Ignore Changes, Save as Origin, Keep Both Versions, Replace Content), they generally found it difficult to decide which versions to replace. They thus chose “*to keep both versions*” as a “*safe fallback*” (e5), or ignoring changes if they found that changes were made by accident via the highlighted differences in the slide comparison window. Some also expressed their willingness to save sections (c10) and “*particular slides*”, such as a core idea, introduction page, or conclusion page (e4), if those were important and general enough to be reused in the future.

Chapter 7

Discussion

Rooted in the “restrict principle” introduced by Reitman [129] and Stokes [155], ReSlide demonstrates how constraints can be treated as a design factor in presentation authoring. Specifically, our work shows how the three identified contextual constraints—time, audience, and communicative intent—can serve as active design factors in author-side communication support. In practice, ReSlide increased presenters’ awareness of constraints and improved their efficiency and effectiveness of presentation authoring.

Two key mechanisms supported this process. First, by asking presenters to explicitly define constraints at the beginning of authoring, ReSlide anchored these considerations in their minds. Second, through the Timeline Overview and highlight reminders, presenters could visualize authoring status. This enabled earlier revisions and reduced the risk of overlooking constraints, improving both authoring efficiency and presentation quality.

Additionally, building on Horvitz’s concept of mixed-initiative interaction [71], our work demonstrates how this paradigm can be instantiated in presentation authoring. ReSlide enables the AI to proactively surface the potential jargon use, while still leaving presenters in control of narrative construction. This extends mixed-initiative design into a previously less studied domain and offers new opportunities for human–AI collaboration.

Below, we discuss a set of design implications to inform future HCI research as well as the limitations of our study.

7.1 Design Implications

7.1.1 Future Tools Should Support Context-aware Authoring

In ReSlide, presenters treated time, audience, and communicative intent not just as constraints to check against, but as design factors that guided presentation authoring. By externalizing these constraints at the start, presenters carried them throughout the process, shaping story construction without adding extra burdens. For example, presenters frequently adjusted section durations or emphasis levels in the Timeline Overview to redistribute intended attention across their story, echoing how sense-making frameworks describe iterative refinement [87]. In contrast, prior tools and textbooks often stopped at a slide-level time estimate or treated time separately from other factors [130, 13], which increased mental overhead during authoring and made it harder for presenters to see the bigger picture. This resonates with the Universal Design for Learning framework [134, 19], which integrates learner characteristics (i.e., background, ability, goals) as active design constraints in shaping curricula. The TA framework [12], similarly, highlights how classroom tools can surface and integrate pedagogical constraints like pacing, curricular goals, and student diversity to support real-time teaching. Future systems should therefore integrate constraints directly into the design process, embedding them as authoring units such as Sections with constraints and Timelines, so that authors can more easily reflect on their progress and make timely adjustments.

7.1.2 Future Tools Should More Carefully Balance Between AI-automation and Author Controls

Our findings also show that the author-side audience-oriented support is more effective when embedded into the authoring process rather than offered as post-hoc repair. In ReSlide, jargon checking was used for flagging difficult terms and also acted as a reflective trigger that encouraged presenters to refine explanations, anticipate questions, and adjust expressions for different audiences. While some participants appreciated jargon checking as an on-demand interaction (c7, c10, e8), some participants desired the jargon checking to work more proactively, automatically giving them more timely suggestions continuously during authoring (e5, e7). This points to a broader design implication: effective authoring tools should balance proactive AI suggestions while preserve the author’s agency, which is especially important in mixed-initiative writing and code generation [70, 74, 131]. For example, Hoque et. al’s HaLLMARK system [70] highlights the significance of retaining

the sense of control and ownership of the text through visualizing interaction provenance, while Hutson [74] points out AI’s assistive roles during writing without impinging human agency. Prior works in code assistance have introduced ‘control mechanisms’ in AI-powered code generation tools to help manage such complexity [33, 172], both highlighting the need to make control options explicit and provide suggestions based on context. Future systems should carefully balance automation with adequate author control, supporting efficiency based on context and reflections without reducing creativity or audience-oriented decision making.

7.1.3 Future Tools Should Include Support for Longitudinal Workflows

Presenters engaged in reuse throughout the authoring process, reusing existing materials at multiple granularities, from individual elements, whole slides, to entire Sections or presentations, depending on their goals at different authoring stages. Synchronization further extended reuse across sessions, allowing presenters to reflect on and push changes from current slides back into the repository for future adaptation. While prior work has already well-leveraged centralized libraries on collaborative, multi-user scenarios [84, 108, 55], our findings show that reuse is equally valuable for individual usage, especially when it is embedded in the storyline constructions. Our findings resonate with the prior works, which also highlight the value and various roles of saving and reusing designs as inspiration resources, communication foundations, and essential components across multiple design sessions [21, 142]. Future systems should support reuse as an *iterative, narrative-driven activity* that evolves alongside storyline development, enabling content to be adapted both within a single project and across multiple projects.

7.2 Limitations and Future Work

Our study was conducted in a lab setting with limited time. These constraints enabled us to capture users’ adaptation strategies in a controlled and concise format, but they did not fully reflect the longer, more iterative workflows that occur in real-world practice. Although we encouraged more naturalistic engagement in the exploratory study by having participants create slides for potential reuse, the short familiarization period limited opportunities for deeper reflection and iteration.

Similarly, synchronization behaviors were only examined in the second session of the exploratory study. With longer-term use, participants may develop more nuanced strategies

for when and how to synchronize content across iterations. Future work should therefore examine RESLIDE over extended periods of time to better capture real-world authoring and reuse practices.

Moreover, in our study, we deliberately limited the available slide authoring functionalities (i.e. shapes, animations, and audio), to isolate and study core interactions in a controlled setting. While this enabled precise observations, richer authoring functionalities might change how participants construct, reuse, and adapt Sections. For example, animations or transitions could influence how presenters structure content across Sections, or how they reuse elements at different granularities. Future work should investigate how expanded authoring features intersect with and are shaped by constraint-sensitive authoring and reuse.

In addition, our current system lacks multi-layered grouping with a restricted structure and a timeline, which participants noted could further enhance flexibility in structuring ideas while maintaining necessary constraints. Implementing such structure with smooth transitions could reduce cognitive overhead and improve expressivity in future iterations. Similarly, our current jargon checking only considers expertise level, which does not fully capture other important audience characteristics. Future system should expand the scope of suggestions, to consider other audience characteristics, narrative status, and contextual constraints such as time and communicative intent. Incorporating slide-level content suggestions, such as specific images and section allocation, could extend benefit beyond text-only phrase suggestions. Future work should focus on integrating these extensions—multi-layered grouping, expanded consideration scope, and richer suggestion support—into a natural content authoring process, where constraints as active design factors.

Chapter 8

Conclusion

In this paper, we introduced RESLIDE, a slide authoring tool designed to support presentation authoring under contextual constraints for a single presentation and long-term authoring cycles. RESLIDE integrated contextual constraints into the iterative presentation authoring process through Section, a unit for grouping slides based on ideas, together with Timeline Overview, Slide Editor, and a centralized library, enabling flexible reuse and making constraints active design factors in storyline construction. RESLIDE also incorporated jargon checking to help presenters tailor their content for audience understanding. Through a comparative within-subject study and an exploratory study, experienced presenters found RESLIDE valuable for aligning contextual constraints with their evolving presentation content across multiple authoring sessions. The design implications derived from this work shed light on future research to design effective authoring environments that better support iterative and constraint-driven content creation.

References

- [1] [2304.02623] Beyond Summarization: Designing AI Support for Real-World Expository Writing Tasks.
- [2] Better Presentations: A Guide for Scholars, Researchers, and Wonks 9780231542791.
- [3] Content Management with empower®.
- [4] (PDF) The Structure of PowerPoint Presentations: The Art of Grasping Things Whole. *ResearchGate*.
- [5] (PDF) The Structure of PowerPoint Presentations: The Art of Grasping Things Whole. *ResearchGate*.
- [6] (PDF) Toward a better understanding of PowerPoint deck design. *ResearchGate*.
- [7] PowerPoint reuse slides.
- [8] Reuse (import) slides from another presentation - Microsoft Support.
- [9] Slide Analysis Method for Editing Lecture Materials based on Hierarchical Structures of Subject Terminologies | IEEE Conference Publication | IEEE Xplore.
- [10] Towards Dialogic and On-Demand Metaphors for Interdisciplinary Reading | Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems.
- [11] Shaaron Ainsworth. DeFT: A conceptual framework for considering learning with multiple representations. *Learning and Instruction*, 16(3):183–198, June 2006.
- [12] Pengcheng An, Kenneth Holstein, Bernice d’Anjou, Berry Eggen, and Saskia Bakker. The TA Framework: Designing Real-time Teaching Augmentation for K-12 Classrooms. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–17, Honolulu HI USA, April 2020. ACM.

- [13] Robert Rene Henri Anholt. *Dazzle 'em with style: the art of oral scientific presentation*. Elsevier Academic Press, Philadelphia, 2nd ed edition, 2006.
- [14] Sumit Asthana, Hannah Rashkin, Elizabeth Clark, Fantine Huot, and Mirella Lapata. Evaluating LLMs for Targeted Concept Simplification for Domain-Specific Texts. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6208–6226, Miami, Florida, USA, 2024. Association for Computational Linguistics.
- [15] Sumit Asthana, Hannah Rashkin, Elizabeth Clark, Fantine Huot, and Mirella Lapata. Evaluating LLMs for Targeted Concept Simplification for Domain-Specific Texts. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6208–6226, 2024. arXiv:2410.20763 [cs].
- [16] Tal August, Kyle Lo, Noah A. Smith, and Katharina Reinecke. Know Your Audience: The benefits and pitfalls of generating plain language summaries beyond the "general" audience. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, pages 1–26, New York, NY, USA, May 2024. Association for Computing Machinery.
- [17] Alex Baker and André Van Der Hoek. Ideas, subjects, and cycles as lenses for understanding the software design process. *Design Studies*, 31(6):590–613, November 2010.
- [18] Sambaran Bandyopadhyay, Himanshu Maheshwari, Anandhavelu Natarajan, and Apoorv Saxena. Enhancing Presentation Slide Generation by LLMs with a Multi-Staged End-to-End Approach, June 2024. arXiv:2406.06556.
- [19] Thomas Baumann and Insa Melle. Evaluation of a digital UDL-based learning environment in inclusive chemistry education. *Chemistry Teacher International*, 1(2):20180026, December 2019.
- [20] CHARLES BAZERMAN. Physicists Reading Physics: Schema-Laden Purposes and Purpose-Laden Schema. *Written Communication*, 2(1):3–23, January 1985. Publisher: SAGE Publications Inc.
- [21] Karim Benharraq and Amy Pavel. HistoryPalette: Supporting Exploration and Reuse of Past Alternatives in Image Generation and Editing, January 2025. arXiv:2501.04163 [cs].

- [22] Lawrence Bergman, Jie Lu, Ravi Konuru, Julie MacNaught, and Danny Yeh. Outline wizard: presentation composition and search. In *Proceedings of the 15th international conference on Intelligent user interfaces*, IUI '10, pages 209–218, New York, NY, USA, February 2010. Association for Computing Machinery.
- [23] Oliver Brdiczka and Doron Kletter. SlideDeckFinder: identifying related slide decks based on visual appearance and composition patterns. In *Proceedings of the 4th Information Interaction in Context Symposium*, IIIX '12, pages 306–309, New York, NY, USA, August 2012. Association for Computing Machinery.
- [24] Matthew Brehmer. Video-Conferencing Beyond Screen-Sharing and Thumbnail Webcam Videos: Gesture-Aware Augmented Reality Video for Data-Rich Remote Presentations, 2025. Version Number: 1.
- [25] José H. Canós, María Isabel Marante, and Manuel Llavador. SliDL: A Slide Digital Library Supporting Content Reuse in Presentations. In Mounia Lalmas, Joemon Jose, Andreas Rauber, Fabrizio Sebastiani, and Ingo Frommholz, editors, *Research and Advanced Technology for Digital Libraries*, pages 453–456, Berlin, Heidelberg, 2010. Springer.
- [26] José H. Canós, María Isabel Marante, and Manuel Llavador. SliDL: A Slide Digital Library Supporting Content Reuse in Presentations. In David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, Mounia Lalmas, Joemon Jose, Andreas Rauber, Fabrizio Sebastiani, and Ingo Frommholz, editors, *Research and Advanced Technology for Digital Libraries*, volume 6273, pages 453–456. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. Series Title: Lecture Notes in Computer Science.
- [27] Yining Cao, Jane L E, Chen Zhu-Tian, and Haijun Xia. DataParticles: Block-based and Language-oriented Authoring of Animated Unit Visualizations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, pages 1–15, New York, NY, USA, April 2023. Association for Computing Machinery.
- [28] Yining Cao, Jane L E, Chen Zhu-Tian, and Haijun Xia. DataParticles: Block-based and Language-oriented Authoring of Animated Unit Visualizations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, pages 1–15, New York, NY, USA, April 2023. Association for Computing Machinery.

- [29] Yining Cao, Jane L E, Chen Zhu-Tian, and Haijun Xia. DataParticles: Block-based and Language-oriented Authoring of Animated Unit Visualizations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, pages 1–15, New York, NY, USA, April 2023. Association for Computing Machinery.
- [30] Yining Cao, Rubaiat Habib Kazi, Li-Yi Wei, Deepali Aneja, and Haijun Xia. Elastica: Adaptive Live Augmented Presentations with Elastic Mappings Across Modalities. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–19, Honolulu HI USA, May 2024. ACM.
- [31] Shirley Carter-Thomas and Elizabeth Rowley-Jolivet. Analysing the scientific conference presentation (CP), a methodological overview of a multimodal genre. *ASp - La revue du GERAS*, (39-40):59–72, 2003. Publisher: revue.org (en ligne) / Bordeaux : GERAS (imprimé).
- [32] Joseph Chee Chang, Amy X. Zhang, Jonathan Bragg, Andrew Head, Kyle Lo, Doug Downey, and Daniel S. Weld. CiteSee: Augmenting Citations in Scientific Papers with Persistent and Personalized Historical Context. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, pages 1–15, New York, NY, USA, April 2023. Association for Computing Machinery.
- [33] Valerie Chen, Alan Zhu, Sebastian Zhao, Hussein Mozannar, David Sontag, and Ameet Talwalkar. Need Help? Designing Proactive AI Assistants for Programming, 2024. Version Number: 2.
- [34] Valerie Chen, Alan Zhu, Sebastian Zhao, Hussein Mozannar, David Sontag, and Ameet Talwalkar. Need Help? Designing Proactive AI Assistants for Programming, October 2024. arXiv:2410.04596 [cs] version: 1.
- [35] Leonardo Christino, Taylor Hill, and Fernando Paulovich. Knowledge-Decks: Automatically Generating Presentation Slide Decks of Visual Analytics Knowledge Discovery Applications, 2022. Version Number: 1.
- [36] Christopher Clark and Santosh Divvala. PDFFigures 2.0: Mining figures from research papers. In *2016 IEEE/ACM Joint Conference on Digital Libraries (JCDL)*, pages 143–152, June 2016.
- [37] Maria João Costa, Hugo Amaro, and Hugo Gonçalo Oliveira. SmartEDU: Accelerating Slide Deck Production with Natural Language Processing. In Elisabeth Métais, Farid Meziane, Vijayan Sugumaran, Warren Manning, and Stephan Reiff-Marganiec,

editors, *Natural Language Processing and Information Systems*, volume 13913, pages 109–123. Springer Nature Switzerland, Cham, 2023. Series Title: Lecture Notes in Computer Science.

- [38] Anita Crescenzi, Yuan Li, Yinglong Zhang, and Rob Capra. Towards Better Support for Exploratory Search through an Investigation of Notes-to-self and Notes-to-share. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR’19, pages 1093–1096, New York, NY, USA, July 2019. Association for Computing Machinery.
- [39] Anita Crescenzi, Yuan Li, Yinglong Zhang, and Rob Capra. Towards Better Support for Exploratory Search through an Investigation of Notes-to-self and Notes-to-share. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR’19, pages 1093–1096, New York, NY, USA, July 2019. Association for Computing Machinery.
- [40] Anita Crescenzi, Yuan Li, Yinglong Zhang, and Rob Capra. Towards Better Support for Exploratory Search through an Investigation of Notes-to-self and Notes-to-share. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR’19, pages 1093–1096, New York, NY, USA, July 2019. Association for Computing Machinery.
- [41] Thi Ngoc Yen Dang. A corpus-based study of vocabulary in conference presentations. *Journal of English for Academic Purposes*, 59:101144, September 2022.
- [42] Laurent Denoue, Scott Carter, and Matthew Cooper. SlideDiff: Animating Textual and Media Changes in Slides. In *Proceedings of the ACM Symposium on Document Engineering 2018*, DocEng ’18, pages 1–4, New York, NY, USA, August 2018. Association for Computing Machinery.
- [43] Tomás Dorta, Annemarie Lesage, Edgar Pérez, and J.M. Christian Bastien. Signs of Collaborative Ideation and the Hybrid Ideation Space. In Toshiharu Taura and Yukari Nagai, editors, *Design Creativity 2010*, pages 199–206. Springer London, London, 2011.
- [44] Graham Dove, Nicolai Brodersen Hansen, and Kim Halskov. An Argument For Design Space Reflection. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*, pages 1–10, Gothenburg Sweden, October 2016. ACM.
- [45] Steven M. Drucker, Georg Petschnigg, and Maneesh Agrawala. Comparing and managing multiple versions of slide presentations. In *Proceedings of the 19th annual*

ACM symposium on User interface software and technology, UIST '06, pages 47–56, New York, NY, USA, October 2006. Association for Computing Machinery.

- [46] Betty Lou Dubais. The Use of Slides in Biomedical Speeches. *The ESP Journal*.
- [47] Mateusz Dubiel, Bereket Abera Yilma, Kayhan Latifzadeh, and Luis A. Leiva. A Contextual Framework for Adaptive User Interfaces: Modelling the Interaction Environment, March 2022. arXiv:2203.16882 [cs].
- [48] Darren Edge, Joan Savage, and Koji Yatani. HyperSlides: dynamic presentation prototyping. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 671–680, New York, NY, USA, April 2013. Association for Computing Machinery.
- [49] Darren Edge, Xi Yang, Yasmine Kotturi, Shuoping Wang, Dan Feng, Bongshin Lee, and Steven Drucker. SlideSpace: Heuristic Design of a Hybrid Presentation Medium. *ACM Trans. Comput.-Hum. Interact.*, 23(3):16:1–16:30, June 2016.
- [50] Quanfu Fan, Arnon Amir, Kobus Barnard, Ranjini Swaminathan, and Alon Efrat. Temporal Modeling of Slide Change in Presentation Videos. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing - ICASSP '07*, pages I–989–I–992, Honolulu, HI, April 2007. IEEE.
- [51] Temiloluwa Femi-Gege, Matthew Brehmer, and Jian Zhao. VisConductor: Affect-Varying Widgets for Animated Data Storytelling in Gesture-Aware Augmented Video Presentation, 2024. Version Number: 1.
- [52] Raymond Fok, Joseph Chee Chang, Tal August, Amy X. Zhang, and Daniel S. Weld. Qlarify: Recursively Expandable Abstracts for Dynamic Information Retrieval over Scientific Papers. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, UIST '24, pages 1–21, New York, NY, USA, October 2024. Association for Computing Machinery.
- [53] Raymond Fok, Hita Kambhamettu, Luca Soldaini, Jonathan Bragg, Kyle Lo, Marti Hearst, Andrew Head, and Daniel S Weld. Scim: Intelligent Skimming Support for Scientific Papers. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*, IUI '23, pages 476–490, New York, NY, USA, March 2023. Association for Computing Machinery.

- [54] Adam Fourney, Michael Terry, and Richard Mann. Gesturing in the Wild: Understanding the Effects and Implications of Gesture-Based Interaction for Dynamic Presentations. 2010.
- [55] Ulrich Frank. A Conception of a Presentation Development and Management System Featuring ‘Smart Slides’. *Enterprise Modelling and Information Systems Architectures (EMISAJ)*, 6(1):47–69, 2011. Number: 1.
- [56] Tsu-Jui Fu, William Yang Wang, Daniel McDuff, and Yale Song. DOC2PPT: Automatic Presentation Slides Generation from Scientific Documents. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(1):634–642, June 2022. Number: 1.
- [57] Verena Fuchsberger, Martin Murer, Ilhan Aslan, Alexander Meschtscherjakov, Manfred Tscheligi, Petra Sundström, and Daniela Petrelli. Contextual constraints: consequences for interaction design. In *Proceedings of the 2014 companion publication on Designing interactive systems*, DIS Companion ’14, pages 227–230, New York, NY, USA, June 2014. Association for Computing Machinery.
- [58] Jacqueline Gartner-Schmidt. The New Normal – Virtual and Hybrid Presentations: Developing Content, Designing Slides, and Delivery Guidelines. *Ear, Nose & Throat Journal*, 101(9_suppl):20S–28S, November 2022. Publisher: SAGE Publications Inc.
- [59] Sreyan Ghosh, Sonal Kumar, Harsh Jalan, Hemant Yadav, and Rajiv Ratn Shah. Cisco at AAAI-CAD21 shared task: Predicting Emphasis in Presentation Slides using Contextualized Embeddings, February 2021. arXiv:2101.11422 [cs].
- [60] Frederic Gmeiner, Nicolai Marquardt, Michael Bentley, Hugo Romat, Michel Pahud, David Brown, Asta Roseway, Nikolas Martelaro, Kenneth Holstein, Ken Hinckley, and Nathalie Riche. Intent Tagging: Exploring Micro-Prompting Interactions for Supporting Granular Human-GenAI Co-Creation Workflows, February 2025. arXiv:2502.18737 [cs].
- [61] Frederic Gmeiner, Nicolai Marquardt, Michael Bentley, Hugo Romat, Michel Pahud, David Brown, Asta Roseway, Nikolas Martelaro, Kenneth Holstein, Ken Hinckley, and Nathalie Riche. Intent Tagging: Exploring Micro-Prompting Interactions for Supporting Granular Human-GenAI Co-Creation Workflows, February 2025. arXiv:2502.18737 [cs].
- [62] Graduate School of Information Science, Nagoya University, Nagoya, Japan, Jie Zhang, Chuan Xiao, Graduate School of Information Science, Nagoya University,

- Nagoya, Japan, Toyohide Watanabe, Nagoya Industrial Science Research Institute, Nagoya, Japan, Yoshiharu Ishikawa, and Graduate School of Information Science, Nagoya University, Nagoya, Japan. Detecting Reused Elements in Presentation Slides. *International Journal of Knowledge Engineering-IACSIT*, 1(2):154–159, 2015.
- [63] Yue Guo, Joseph Chee Chang, Maria Antoniak, Erin Bransom, Trevor Cohen, Lucy Lu Wang, and Tal August. Personalized Jargon Identification for Enhanced Interdisciplinary Communication, November 2023. arXiv:2311.09481 [cs].
 - [64] Brian D. Hall, Lyn Bartram, and Matthew Brehmer. Augmented Chironomia for Presenting Data to Remote Audiences. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pages 1–14, Bend OR USA, October 2022. ACM.
 - [65] Bo He, Jun Wang, Jieli Qiu, Trung Bui, Abhinav Shrivastava, and Zhaowen Wang. Align and Attend: Multimodal Summarization With Dual Contrastive Losses. pages 14867–14878, 2023.
 - [66] Andrew Head, Kyle Lo, Dongyeop Kang, Raymond Fok, Sam Skjonsberg, Daniel S. Weld, and Marti A. Hearst. Augmenting Scientific Papers with Just-in-Time, Position-Sensitive Definitions of Terms and Symbols. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI ’21, pages 1–18, New York, NY, USA, May 2021. Association for Computing Machinery.
 - [67] Lucas Torroba Hennigen, Shannon Shen, Aniruddha Nrusimha, Bernhard Gapp, David Sontag, and Yoon Kim. Towards Verifiable Text Generation with Symbolic References, April 2024. arXiv:2311.09188.
 - [68] Terje Hillesund. Digital reading spaces: How expert readers handle books, the Web and electronic paper. *First Monday*, April 2010.
 - [69] E. Hoque, P. Kavehzadeh, and A. Masry. Chart Question Answering: State of the Art and Future Directions. *Computer Graphics Forum*, 41(3):555–572, 2022. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14573>.
 - [70] Md Naimul Hoque, Tasfia Mashiat, Bhavya Ghai, Cecilia D. Shelton, Fanny Chevalier, Kari Kraus, and Niklas Elmqvist. The HaLLMark Effect: Supporting Provenance and Transparent Use of Large Language Models in Writing with Interactive Visualization. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–15, Honolulu HI USA, May 2024. ACM.

- [71] Eric Horvitz. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, CHI '99, pages 159–166, New York, NY, USA, May 1999. Association for Computing Machinery.
- [72] Kate Y. Huang, Helene H. Fung, and Peifeng Sun. The effect of audience–character similarity on identification with narrative characters: A meta-analysis. *Current Psychology*, 43(8):7026–7043, February 2024.
- [73] Mina Huh, Dingzeyu Li, Kim Pimmel, Hijung Valentina Shin, Amy Pavel, and Mira Dontcheva. VideoDiff: Human-AI Video Co-Creation with Alternatives, February 2025. arXiv:2502.10190 [cs].
- [74] James Hutson. Human-AI Collaboration in Writing: A Multidimensional Framework for Creative and Intellectual Authorship. *International Journal of Changes in Education*, February 2025.
- [75] Naoto Inoue, Kotaro Kikuchi, Edgar Simo-Serra, Mayu Otani, and Kota Yamaguchi. Towards Flexible Multi-Modal Document Models. pages 14287–14296, 2023.
- [76] Farnaz Jahanbakhsh, Elnaz Nouri, Robert Sim, Ryen W. White, and Adam Fourney. Understanding Questions that Arise When Working with Business Documents. *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW2):341:1–341:24, November 2022.
- [77] Vijay Jaisankar, Sambaran Bandyopadhyay, Kalp Vyas, Varre Chaitanya, and Shwetha Somasundaram. PostDoc: Generating Poster from a Long Multimodal Document Using Deep Submodular Optimization, May 2024. arXiv:2405.20213.
- [78] P. N. Johnson-Laird. Book Review: Cognition and Thought: An Information-Processing Approach. *Quarterly Journal of Experimental Psychology*, 18(4):386–387, November 1966.
- [79] Liwei He Jonathan Grudin. Comparing Presentation Summaries: Slide vs. Reading vs. Listening. *CHI 2000*.
- [80] Hyeonsu Kang, Joseph Chee Chang, Yongsung Kim, and Aniket Kittur. Threddy: An Interactive System for Personalized Thread-based Exploration and Organization of Scientific Literature. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, UIST '22, pages 1–15, New York, NY, USA, October 2022. Association for Computing Machinery.

- [81] Hyeonsu B Kang, Tongshuang Wu, Joseph Chee Chang, and Aniket Kittur. Synergi: A Mixed-Initiative System for Scholarly Synthesis and Sensemaking. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–19, San Francisco CA USA, October 2023. ACM.
- [82] Yudai Kato, Shun Shiramatsu, Tadachika Ozono, and Toramatsu Shintani. Slide Repository: A Repository for Slide Reuse Using Research Episode Network. In *2014 IIAI 3rd International Conference on Advanced Applied Informatics*, pages 198–203, August 2014.
- [83] Robin H. Kay. Exploring the use of video podcasts in education: A comprehensive review of the literature. *Computers in Human Behavior*, 28(3):820–831, May 2012.
- [84] Ali Khalili, Sören Auer, Darya Tarasowa, and Ivan Ermilov. SlideWiki: Elicitation and Sharing of Corporate Knowledge Using Presentations. In Annette ten Teije, Johanna Völker, Siegfried Handschuh, Heiner Stuckenschmidt, Mathieu d’Acquin, Andriy Nikolov, Nathalie Aussenac-Gilles, and Nathalie Hernandez, editors, *Knowledge Engineering and Knowledge Management*, pages 302–316, Berlin, Heidelberg, 2012. Springer.
- [85] Tae Soo Kim, Matt Latzke, Jonathan Bragg, Amy X. Zhang, and Joseph Chee Chang. Papeos: Augmenting Research Papers with Talk Videos. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, UIST ’23, pages 1–19, New York, NY, USA, October 2023. Association for Computing Machinery.
- [86] G. Klein, B. Moon, and R.R. Hoffman. Making Sense of Sensemaking 1: Alternative Perspectives. *IEEE Intelligent Systems*, 21(4):70–73, July 2006.
- [87] Jon Kolko. Sensemaking and Framing: A Theoretical Reflection on Perspective in Design Synthesis. July 2010.
- [88] Merlijn Koupprie, , and Froukje Sleeswijk Visser. A framework for empathy in design: stepping into and out of the user’s life. *Journal of Engineering Design*, 20(5):437–448, October 2009. Publisher: Taylor & Francis .eprint: <https://doi.org/10.1080/09544820902875033>.
- [89] Mina Lee, Katy Ilonka Gero, John Joon Young Chung, Simon Buckingham Shum, Vipul Raheja, Hua Shen, Subhashini Venugopalan, Thiemo Wambsganss, David Zhou, Emad A. Alghamdi, Tal August, Avinash Bhat, Madiha Zahrah Choksi,

- Senjuti Dutta, Jin L. C. Guo, Md Naimul Hoque, Yewon Kim, Simon Knight, Seyed Parsa Neshaei, Agnia Sergeyuk, Antonette Shibani, Disha Shrivastava, Lila Shroff, Jessi Stark, Sarah Stermann, Sitong Wang, Antoine Bosselut, Daniel Buschek, Joseph Chee Chang, Sherol Chen, Max Kreminski, Joonsuk Park, Roy Pea, Eugenia H. Rho, Shannon Zejiang Shen, and Pao Siangliulue. A Design Space for Intelligent and Interactive Writing Assistants, March 2024. arXiv:2403.14117.
- [90] Guy Lev, Michal Shmueli-Scheuer, Jonathan Herzig, Achiya Jerbi, and David Konopnicki. *TalkSumm: A Dataset and Scalable Annotation Method for Scientific Paper Summarization Based on Conference Talks*. June 2019.
- [91] Junze Li, Yue Zhang, Chengbo Zheng, Dingdong Liu, Zeyu Huang, and Xiaojuan Ma. InsightBridge: Enhancing Empathizing with Users through Real-Time Information Synthesis and Visual Communication. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, pages 1–15, New York, NY, USA, April 2025. Association for Computing Machinery.
- [92] Leonhard Lichtschlag, Thorsten Karrer, and Jan Borchers. Fly: a tool to author planar presentations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '09, pages 547–556, New York, NY, USA, April 2009. Association for Computing Machinery.
- [93] M. Liesse, G. Lhoest, A. Trouet, and P. Tulkens. Uptake and intracellular localization of morphine in lysosomes and cell sap of cultured fibroblasts. *Archives Internationales De Physiologie Et De Biochimie*, 84(3):638–639, 1976.
- [94] Chin-Yew Lin. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics.
- [95] Yefeng Liu, Darren Edge, and Koji Yatani. SidePoint: a peripheral knowledge panel for presentation slide authoring. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 681–684, New York, NY, USA, April 2013. Association for Computing Machinery.
- [96] Yuhan Liu, Aadit Shah, Jordan Ackerman, and Manaswi Saha. Exploring the Design Space of Real-time LLM Knowledge Support Systems: A Case Study of Jargon Explanations. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, pages 1–20, New York, NY, USA, April 2025. Association for Computing Machinery.

- [97] Kyle Lo, Joseph Chee Chang, Andrew Head, Jonathan Bragg, Amy X. Zhang, Cassidy Trier, Chloe Anastasiades, Tal August, Russell Authur, Danielle Bragg, Erin Bransom, Isabel Cachola, Stefan Candra, Yoganand Chandrasekhar, Yen-Sung Chen, Evie Yu-Yen Cheng, Yvonne Chou, Doug Downey, Rob Evans, Raymond Fok, Fangzhou Hu, Regan Huff, Dongyeop Kang, Tae Soo Kim, Rodney Kinney, Aniket Kittur, Hyeonsu Kang, Egor Klevak, Bailey Kuehl, Michael Langan, Matt Latzke, Jaron Lochner, Kelsey MacMillan, Eric Marsh, Tyler Murray, Aakanksha Naik, Ngoc-Uyen Nguyen, Srishti Palani, Soya Park, Caroline Paulic, Napol Rachatasumrit, Smita Rao, Paul Sayre, Zejiang Shen, Pao Siangliulue, Luca Soldaini, Huy Tran, Madeleine van Zuylen, Lucy Lu Wang, Christopher Wilhelm, Caroline Wu, Jiangjiang Yang, Angele Zamarron, Marti A. Hearst, and Daniel S. Weld. The Semantic Reader Project: Augmenting Scholarly Documents through AI-Powered Interactive Reading Interfaces, April 2023. arXiv:2303.14334.
- [98] Kyle Lo, Joseph Chee Chang, Andrew Head, Jonathan Bragg, Amy X. Zhang, Cassidy Trier, Chloe Anastasiades, Tal August, Russell Authur, Danielle Bragg, Erin Bransom, Isabel Cachola, Stefan Candra, Yoganand Chandrasekhar, Yen-Sung Chen, Evie Yu-Yen Cheng, Yvonne Chou, Doug Downey, Rob Evans, Raymond Fok, Fangzhou Hu, Regan Huff, Dongyeop Kang, Tae Soo Kim, Rodney Kinney, Aniket Kittur, Hyeonsu B. Kang, Egor Klevak, Bailey Kuehl, Michael J. Langan, Matt Latzke, Jaron Lochner, Kelsey MacMillan, Eric Marsh, Tyler Murray, Aakanksha Naik, Ngoc-Uyen Nguyen, Srishti Palani, Soya Park, Caroline Paulic, Napol Rachatasumrit, Smita Rao, Paul Sayre, Zejiang Shen, Pao Siangliulue, Luca Soldaini, Huy Tran, Madeleine Van Zuylen, Lucy Lu Wang, Christopher Wilhelm, Caroline Wu, Jiangjiang Yang, Angele Zamarron, Marti A. Hearst, and Daniel S. Weld. The Semantic Reader Project. *Communications of the ACM*, 67(10):50–61, October 2024.
- [99] Kyle Lo, Zejiang Shen, Benjamin Newman, Joseph Chang, Russell Authur, Erin Bransom, Stefan Candra, Yoganand Chandrasekhar, Regan Huff, Bailey Kuehl, Amanpreet Singh, Chris Wilhelm, Angele Zamarron, Marti A. Hearst, Daniel Weld, Doug Downey, and Luca Soldaini. PaperMage: A Unified Toolkit for Processing, Representing, and Manipulating Visually-Rich Scientific Documents. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 495–507, Singapore, 2023. Association for Computational Linguistics.
- [100] Zehao Lu. Unsupervised Paper2Slides Generation. Master’s thesis, 2024. Accepted: 2024-02-15T14:49:49Z.

- [101] Himanshu Maheshwari, Sambaran Bandyopadhyay, Aparna Garimella, and Anand-havelu Natarajan. Presentations are not always linear! GNN meets LLM for Text Document-to-Presentation Transformation with Attribution. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15948–15962, Miami, Florida, USA, November 2024. Association for Computational Linguistics.
- [102] Damien Masson, Young-Ho Kim, and Fanny Chevalier. Textoshop: Interactions Inspired by Drawing Software to Facilitate Text Editing, February 2025. arXiv:2409.17088 [cs].
- [103] Damien Masson, Sylvain Malacria, G  ry Casiez, and Daniel Vogel. Charagraph: Interactive Generation of Charts for Realtime Annotation of Data-Rich Paragraphs. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, pages 1–18, New York, NY, USA, April 2023. Association for Computing Machinery.
- [104] Damien Masson, Sylvain Malacria, G  ry Casiez, and Daniel Vogel. DirectGPT: A Direct Manipulation Interface to Interact with Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, pages 1–16, New York, NY, USA, May 2024. Association for Computing Machinery.
- [105] Damien Masson, Sylvain Malacria, Edward Lank, and G  ry Casiez. Chameleon: Bringing Interactivity to Static Digital Documents. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, pages 1–13, New York, NY, USA, April 2020. Association for Computing Machinery.
- [106] Martha J. Maxwell. Skimming and Scanning Improvement: The Needs, Assumptions and Knowledge Base. *Journal of Reading Behavior*, 5(1):47–59, March 1972. Publisher: SAGE Publications.
- [107] Richard E. Mayer, Sherry Fennell, Lindsay Farmer, and Julie Campbell. A Personalization Effect in Multimedia Learning: Students Learn Better When Words Are in Conversational Style Rather Than Formal Style. *Journal of Educational Psychology*, 96(2):389–395, 2004. Place: US Publisher: American Psychological Association.
- [108] Yelena Mejova, Klaar De Schepper, Lawrence Bergman, and Jie Lu. Reuse in the wild: an empirical and ethnographic study of organizational content reuse. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI

- '11, pages 2877–2886, New York, NY, USA, May 2011. Association for Computing Machinery.
- [109] Yan Meng, Liangming Pan, Yixin Cao, and Min-Yen Kan. FOLLOWUPQG: Towards Information-Seeking Follow-up Question Generation, September 2023. arXiv:2309.05007.
 - [110] Ine Mols, Elise Van Den Hoven, and Berry Eggen. Informing Design for Reflection: an Overview of Current Everyday Practices. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*, pages 1–10, Gothenburg Sweden, October 2016. ACM.
 - [111] Ine Mols, Elise Van Den Hoven, and Berry Eggen. Informing Design for Reflection: an Overview of Current Everyday Practices. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*, pages 1–10, Gothenburg Sweden, October 2016. ACM.
 - [112] Ishani Mondal, Shwetha S, Anandhavelu Natarajan, Aparna Garimella, Sambaran Bandyopadhyay, and Jordan Boyd-Graber. Presentations by the Humans and For the Humans: Harnessing LLMs for Generating Persona-Aware Slides from Documents. In Yvette Graham and Matthew Purver, editors, *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2664–2684, St. Julian’s, Malta, March 2024. Association for Computational Linguistics.
 - [113] Michael Mose Biskjaer and Kim Halskov. Decisive constraints as a creative resource in interaction design. *Digital Creativity*, 25(1):27–61, January 2014. Publisher: CAA Website eprint: <https://doi.org/10.1080/14626268.2013.855239>.
 - [114] Prasanth Murali, Javier Hernandez, Daniel McDuff, Kael Rowan, Jina Suh, and Mary Czerwinski. AffectiveSpotlight: Facilitating the Communication of Affective Responses from Audience Members during Online Presentations. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–13, Yokohama Japan, May 2021. ACM.
 - [115] Elisabeth Métais, Farid Meziane, Vijayan Sugumaran, Warren Manning, and Stephan Reiff-Marganiec, editors. *Natural Language Processing and Information Systems: 28th International Conference on Applications of Natural Language to Information Systems, NLDB 2023, Derby, UK, June 21–23, 2023, Proceedings*, volume 13913 of *Lecture Notes in Computer Science*. Springer Nature Switzerland, Cham, 2023.

- [116] Kristen M. Naegle. Ten simple rules for effective presentation slides. *PLOS Computational Biology*, 17(12):e1009554, December 2021.
- [117] Nagoya University, Japan, Jie Zhang, Chuan Xiao, Nagoya University, Japan, Sheng Hu, Nagoya University, Japan, Toyohide Watanabe, Nagoya Industrial Science Research Institute, Japan, Yoshiharu Ishikawa, and Nagoya University, Japan. Managing Presentation Slides with Reused Elements. *International Journal of Information and Education Technology*, 6(3):170–177, 2016.
- [118] Benjamin Newman, Luca Soldaini, Raymond Fok, Arman Cohan, and Kyle Lo. A Question Answering Framework for Decontextualizing User-facing Snippets from Scientific Documents. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3194–3212, Singapore, December 2023. Association for Computational Linguistics.
- [119] Benjamin Newman, Luca Soldaini, Raymond Fok, Arman Cohan, and Kyle Lo. A Question Answering Framework for Decontextualizing User-facing Snippets from Scientific Documents. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3194–3212, Singapore, December 2023. Association for Computational Linguistics.
- [120] Sachita Nishal, Eric Lee, and Nicholas Diakopoulos. De-jargonizing Science for Journalists with GPT-4: A Pilot Study, October 2024. arXiv:2410.12069 [cs].
- [121] Donald A. Norman, editor. *User centered system design: new perspectives on human-computer interaction*. Erlbaum, Hillsdale, NJ, 9. [print.] edition, 1986.
- [122] Yoshiaki Oida, Kazuo Hiekata, Taiga Mitsuyuki, Hiroki Kamba, and Isaac Okada. Development of presentation slide retrieval system based on visual information. *Journal of Industrial Information Integration*, 12:3–12, December 2018.
- [123] Yoshiaki Oida, Kazuo Hiekata, Taiga Mitsuyuki, Hiroki Kamba, and Isaac Okada. Development of presentation slide retrieval system based on visual information. *Journal of Industrial Information Integration*, 12:3–12, December 2018.
- [124] Sangjun Park, Chanhee Lee, Jieun Han, and Uichin Lee. QuickRef: Should I Read Cited Papers for Understanding This Paper? In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI EA ’23, pages 1–7, New York, NY, USA, April 2023. Association for Computing Machinery.

- [125] Yi-Hao Peng, Peggy Chi, Anjuli Kannan, Meredith Ringel Morris, and Irfan Essa. Slide Gestalt: Automatic Structure Extraction in Slide Decks for Non-Visual Access. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, pages 1–14, New York, NY, USA, April 2023. Association for Computing Machinery.
- [126] Peter Pirolli and Stuart Card. The Sensemaking Process and Leverage Points for Analyst Technology as Identified Through Cognitive Task Analysis.
- [127] Larissa Pschetz, Koji Yatani, and Darren Edge. TurningPoint: narrative-driven presentation planning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 1591–1594, New York, NY, USA, April 2014. Association for Computing Machinery.
- [128] Abhinav Ramesh Kashyap, Yajing Yang, and Min-Yen Kan. Scientific document processing: challenges for modern learning methods. *International Journal on Digital Libraries*, 24(4):283–309, December 2023.
- [129] Walter Ralph Reitman. Cognition and thought : an information-processing approach. In *Cognition and thought : an information-processing approach*. Wiley, New York, 1965. tex.lccn: 65025856.
- [130] Garr Reynolds. *Presentation zen: simple ideas on presentation design and delivery*. Voices that matter. New Riders, Berkeley, CA, 2nd ed., rev. & updated edition, 2012.
- [131] Mohi Reza, Jeb Thomas-Mitchell, Peter Dushniku, Nathan Laundry, Joseph Jay Williams, and Anastasia Kuzminykh. Co-Writing with AI, on Human Terms: Aligning Research with User Demands Across the Writing Process, August 2025. arXiv:2504.12488 [cs].
- [132] Reinout Roels and Beat Signer. MindXpres: An Extensible Content-Driven Cross-Media Presentation Platform. In Boualem Benatallah, Azer Bestavros, Yannis Manolopoulos, Athena Vakali, and Yanchun Zhang, editors, *Web Information Systems Engineering – WISE 2014*, pages 215–230, Cham, 2014. Springer International Publishing.
- [133] Reinout Roels and Beat Signer. A Conceptual Framework and Content Model for Next Generation Presentation Solutions. *Proc. ACM Hum.-Comput. Interact.*, 3(EICS):7:1–7:22, June 2019.

- [134] David H. Rose and Anne Meyer. *Teaching every student in the Digital Age: universal design for learning*. Association for Supervision and Curriculum Development, Alexandria, Va, 2002.
- [135] Elizabeth Rowley-Jolivet and Shirley Carter-Thomas. The rhetoric of conference presentation introductions : context, argument and interaction. *International Journal of Applied Linguistics*, 15(1):45, 2005.
- [136] Bahador Saket, Sijie Yang, Hong Tan, Koji Yatani, and Darren Edge. TalkZones: section-based time support for presentations. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services, MobileHCI '14*, pages 263–272, New York, NY, USA, September 2014. Association for Computing Machinery.
- [137] Jonathan A. Schwabish. *Better presentations: a guide for scholars, researchers, and wonks*. Columbia University Press, New York City, 2017.
- [138] D.A. Schön. Designing as reflective conversation with the materials of a design situation. *Knowledge-Based Systems*, 5(1):3–14, March 1992.
- [139] Donald A. Schön. *The reflective practitioner: how professionals think in action*. Basic Books, New York, 1983.
- [140] Edward Segel and Jeffrey Heer. Narrative Visualization: Telling Stories with Data. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1139–1148, November 2010. Conference Name: IEEE Transactions on Visualization and Computer Graphics.
- [141] Moushumi Sharmin and Brian P. Bailey. "I reflect to improve my design": investigating the role and process of reflection in creative design. In *Proceedings of the 8th ACM conference on Creativity and cognition*, pages 389–390, Atlanta Georgia USA, November 2011. ACM.
- [142] Moushumi Sharmin, Brian P. Bailey, Cole Coats, and Kevin Hamilton. Understanding knowledge management practices for early design activity and its implications for reuse. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2367–2376, Boston MA USA, April 2009. ACM.
- [143] Moushumi Sharmin, Lawrence Bergman, Jie Lu, and Ravi Konuru. On slide-based contextual cues for presentation reuse. In *Proceedings of the 2012 ACM international*

- conference on Intelligent User Interfaces*, IUI '12, pages 129–138, New York, NY, USA, February 2012. Association for Computing Machinery.
- [144] Moushumi Sharmin, Lawrence Bergman, Jie Lu, and Ravi Konuru. On slide-based contextual cues for presentation reuse. In *Proceedings of the 2012 ACM international conference on Intelligent User Interfaces*, IUI '12, pages 129–138, New York, NY, USA, February 2012. Association for Computing Machinery.
 - [145] Zejiang Shen, Tal August, Pao Siangliulue, Kyle Lo, Jonathan Bragg, Jeff Hammerbacher, Doug Downey, Joseph Chee Chang, and David Sontag. Beyond Summarization: Designing AI Support for Real-World Expository Writing Tasks, April 2023. arXiv:2304.02623.
 - [146] Yu-Zhe Shi, Hao-fei Hou, Zhang-qian Bi, Fan-xu Meng, Xiang Wei, Le-cheng Ruan, and Qining Wang. AutoDSL: Automated domain-specific language design for structural representation of procedures with constraints. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12177–12214, 2024. arXiv:2406.12324 [cs].
 - [147] Yu-Zhe Shi, Haotian Li, Lecheng Ruan, and Huamin Qu. Constraint Representation Towards Precise Data-Driven Storytelling. In *2024 IEEE VIS Workshop on Data Storytelling in an Era of Generative AI (GEN4DS)*, pages 4–12, St. Pete Beach, FL, USA, October 2024. IEEE.
 - [148] Donghoon Shin, Lucy Lu Wang, and Gary Hsieh. From Paper to Card: Transforming Design Implications with Generative AI. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, pages 1–15, New York, NY, USA, May 2024. Association for Computing Machinery.
 - [149] Amirreza Shirani, Franck Deroncourt, Nedim Lipka, Paul Asente, Jose Echevarria, and Thamar Solorio. SemEval-2020 Task 10: Emphasis Selection for Written Text in Visual Media. In Aurelie Herbelot, Xiaodan Zhu, Alexis Palmer, Nathan Schneider, Jonathan May, and Ekaterina Shutova, editors, *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1360–1370, Barcelona (online), December 2020. International Committee for Computational Linguistics.
 - [150] Amirreza Shirani, Giai Tran, Hieu Trinh, Franck Deroncourt, Nedim Lipka, Paul Asente, Jose Echevarria, and Thamar Solorio. Learning to Emphasize: Dataset and Shared Task Models for Selecting Emphasis in Presentation Slides, January 2021. arXiv:2101.03237 [cs].

- [151] Amirreza Shirani, Giai Tran, Hieu Trinh, Franck Dernoncourt, Nedim Lipka, Jose Echevarria, Thamar Solorio, and Paul Asente. PSED: A Dataset for Selecting Emphasis in Presentation Slides. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli, editors, *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4314–4320, Online, August 2021. Association for Computational Linguistics.
- [152] Yifan Song, Wing Yee Au, Hon Yung Wong, Brian P. Bailey, and Tal August. Personalized Real-time Jargon Support for Online Meetings, August 2025. arXiv:2508.10239 [cs].
- [153] Ryan Spicer, Yu-Ru Lin, Aisling Kelliher, and Hari Sundaram. NextSlidePlease: Authoring and delivering agile multimedia presentations. *ACM Trans. Multimedia Comput. Commun. Appl.*, 8(4):53:1–53:20, November 2012.
- [154] Ryan P. Spicer and Aisling Kelliher. NextSlidePlease: navigation and time management for hyperpresentations. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '09, pages 3883–3888, New York, NY, USA, April 2009. Association for Computing Machinery.
- [155] Patricia D. Stokes. Creativity from Constraints: What can we learn from Motherwell? from Modrian? from Klee? *The Journal of Creative Behavior*, 42(4):223–236, December 2008.
- [156] Sangho Suh, Martinet Lee, and Edith Law. How do we design for concreteness fading? Survey, general framework, and design dimensions. In *Proceedings of the Interaction Design and Children Conference*, IDC '20, pages 581–588, New York, NY, USA, June 2020. Association for Computing Machinery.
- [157] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. Sensecape: Enabling Multilevel Exploration and Sensemaking with Large Language Models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, UIST '23, pages 1–18, New York, NY, USA, October 2023. Association for Computing Machinery.
- [158] Sangho Suh, Jian Zhao, and Edith Law. CodeToon: Story Ideation, Auto Comic Generation, and Structure Mapping for Code-Driven Storytelling. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, UIST '22, pages 1–16, New York, NY, USA, October 2022. Association for Computing Machinery.

- [159] Edward Sun, Yufang Hou, Dakuo Wang, Yunfeng Zhang, and Nancy X. R. Wang. D2S: Document-to-Slide Generation Via Query-Based Text Summarization. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou, editors, *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1405–1418, Online, June 2021. Association for Computational Linguistics.
- [160] Tono Suwartono and Miftahul Furqon. Presentation Content Structuredness: How does It Affect the Audience? *Jurnal Imiah Pendidikan dan Pembelajaran*, 5(3):446, August 2021.
- [161] Amanda Swearngin, Chenglong Wang, Alannah Oleson, James Fogarty, and Amy J. Ko. 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*, pages 1–13, Honolulu HI USA, April 2020. ACM.
- [162] Kentaro Takahira, Wong Kam-Kwai, Leni Yang, Xian Xu, Takanori Fujiwara, and Huamin Qu. TangibleNet: Synchronous Network Data Storytelling through Tangible Interactions in Augmented Reality. 2025. Publisher: arXiv Version Number: 1.
- [163] Kentaro Takahira, Yue Yu, Takanori Fujiwara, Ryo Suzuki, and Huamin Qu. InSitu-Tale: Enhancing Augmented Data Storytelling with Physical Objects, 2025. Version Number: 3.
- [164] Ryota Tanaka, Kyosuke Nishida, Kosuke Nishida, Taku Hasegawa, Itsumi Saito, and Kuniko Saito. SlideVQA: A Dataset for Document Visual Question Answering on Multiple Images. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11):13636–13645, June 2023. Number: 11.
- [165] Prateksha Udhayanan, Suryateja Bv, Parth Laturia, Dev Chauhan, Darshan Khandelwal, Stefano Petrangeli, and Balaji Vasanth Srinivasan. Recipe2Video: Synthesizing Personalized Videos From Recipe Texts. pages 2268–2277, 2023.
- [166] Maria Vlahova-Takova and Milena Lazarova. Dual-Branch Convolutional Neural Network for Image Comparison in Presentation Style Coherence. *Engineering, Technology & Applied Science Research*, 15(2):21719–21727, April 2025.
- [167] David Wadden, Kejian Shi, Jacob Morrison, Aakanksha Naik, Shruti Singh, Nitzan Barzilay, Kyle Lo, Tom Hope, Luca Soldaini, Shannon Zejiang Shen, Doug

- Downey, Hannaneh Hajishirzi, and Arman Cohan. SciRIF: A Resource to Enhance Language Model Instruction-Following over Scientific Literature, August 2024. arXiv:2406.07835.
- [168] Fengjie Wang, Yanna Lin, Leni Yang, Haotian Li, Mingyang Gu, Min Zhu, and Huamin Qu. OutlineSpark: Igniting AI-powered Presentation Slides Creation from Computational Notebooks through Outlines. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–16, Honolulu HI USA, May 2024. ACM.
 - [169] Fengjie Wang, Xuye Liu, Ou Jing Liu, Ali Neshati, Tengfei Ma, Min Zhu, and Jian Zhao. Slide4N: Creating Presentation Slides from Computational Notebooks with Human-AI Collaboration. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–18, Hamburg Germany, April 2023. ACM.
 - [170] Leijie Wang, Kathryn Yurechko, Pranati Dani, Quan Ze Chen, and Amy X. Zhang. End User Authoring of Personalized Content Classifiers: Comparing Example Labeling, Rule Writing, and LLM Prompting. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, pages 1–21, New York, NY, USA, April 2025. Association for Computing Machinery.
 - [171] Ning Wang, Genaro Rebolledo-Mendez, Vania Dimitrova, Noboru Matsuda, and Olga C. Santos, editors. *Artificial Intelligence in Education. Posters and Late Breaking Results, Workshops and Tutorials, Industry and Innovation Tracks, Practitioners, Doctoral Consortium and Blue Sky: 24th International Conference, AIED 2023, Tokyo, Japan, July 3–7, 2023, Proceedings*, volume 1831 of *Communications in Computer and Information Science*. Springer Nature Switzerland, Cham, 2023.
 - [172] Ruotong Wang, Ruijia Cheng, Denae Ford, and Thomas Zimmermann. Investigating and Designing for Trust in AI-powered Code Generation Tools. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 1475–1493, Rio de Janeiro Brazil, June 2024. ACM.
 - [173] Ruotong Wang, Lin Qiu, Justin Cranshaw, and Amy X. Zhang. Meeting Bridges: Designing Information Artifacts that Bridge from Synchronous Meetings to Asynchronous Collaboration. *Proc. ACM Hum.-Comput. Interact.*, 8(CSCW1):35:1–35:29, April 2024.
 - [174] Xin Wang and Ning Li. A Presentation Generation Method Combining Topic Segmentation and Text Summarization. In *2024 IEEE Symposium on Computers and Communications (ISCC)*, pages 1–7, June 2024. ISSN: 2642-7389.

- [175] Yuanyuan Wang and Kazutoshi Sumiya. A Browsing Method for Presentation Slides Based on Semantic Relations and Document Structure for e-Learning. *Journal of Information Processing*, 20(1):11–25, 2012.
- [176] Thomas Winters and Kory W. Mathewson. Automatically Generating Engaging Presentation Slide Decks. In Anikó Ekárt, Antonios Liapis, and María Luz Castro Pena, editors, *Computational Intelligence in Music, Sound, Art and Design*, volume 11453, pages 127–141. Springer International Publishing, Cham, 2019. Series Title: Lecture Notes in Computer Science.
- [177] Peter Wright and John McCarthy. Empathy and experience in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’08, pages 637–646, New York, NY, USA, April 2008. Association for Computing Machinery.
- [178] Bohao Wu, Qingyun Wang, and Yue Guo. Explain Less, Understand More: Jargon Detection via Personalized Parameter-Efficient Fine-tuning, May 2025. arXiv:2505.16227 [cs].
- [179] JoAnne Yates. The PowerPoint Presentation and Its Corollaries: How Genres Shape Communicative Action in Organizations. In Mark Zachry and V Thralls, editors, *Communicative Practices in Workplaces and the Professions: Cultural Perspectives on the Regulation of Discourse and Organizations*. Baywood Publishing Company, Inc., 2007.
- [180] Edward Zanders and Lindsay MacLeod. *Presentation Skills for Scientists: A Practical Guide*. Cambridge University Press, 2 edition, October 2018.
- [181] Jie Zhang, Chuan Xiao, Toyohide Watanabe, and Yoshiharu Ishikawa. Content-Based Element Search for Presentation Slide Reuse. *IEICE Transactions on Information and Systems*, E97.D(10):2685–2696, 2014.
- [182] Litian Zhang, Xiaoming Zhang, Ziming Guo, and Zhipeng Liu. CISum: Learning Cross-modality Interaction to Enhance Multimodal Semantic Coverage for Multimodal Summarization, February 2023. arXiv:2302.09934.
- [183] Chengbo Zheng, Yuanhao Zhang, Zeyu Huang, Chuhan Shi, Minrui Xu, and Xiaojuan Ma. DiscipLink: Unfolding Interdisciplinary Information Seeking Process via Human-AI Co-Exploration. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, UIST ’24, pages 1–20, New York, NY, USA, October 2024. Association for Computing Machinery.

APPENDICES

.1 Jargon Checking prompts

.1.1 Prompt for Expanded Audience Profile

.1 Prompts for Audience Context Expansion

Listing 1: Prompt for audience context expansion used in ReSlide

```
You are an expert at understanding audience descriptions for
presentations.
Analyze this audience description and provide detailed context
that will help with jargon detection.
```

```
ORIGINAL AUDIENCE DESCRIPTION: "${originalDescription}"
USER-PROVIDED EXPERTISE LEVEL: ${userExpertiseLevel}/5
```

```
Your task: Expand this into a detailed profile that clearly
explains what this audience would and would not know.
```

```
Respond in JSON format:
```

```
{
  "expandedDescription": "Detailed 2-3 sentence description of
    their background and knowledge level",
  "inferredExpertiseLevel": number (1-5, can adjust user's
    estimate if clearly wrong),
  "knownConcepts": ["concept1", "concept2", "concept3"],
  "likelyJargon": ["term1", "term2", "term3"],
  "domainBackground": "Their field/industry background"
```

```
}
```

EXAMPLES:

Input: "NLP professor"

Output: {

```
  "expandedDescription": "Computer Science professor specializing
    in Natural Language Processing research. Has PhD-level
    expertise in machine learning, deep learning, linguistics,
    and computational methods. Familiar with all standard ML
    algorithms, programming concepts, and academic research
    terminology.",
  "inferredExpertiseLevel": 5,
  "knownConcepts": ["machine learning", "neural networks", "
    transformers", "BERT", "random forest", "SVM", "tokenization
    ", "embeddings", "algorithms", "APIs", "frameworks", "deep
    learning", "statistical models"],
  "likelyJargon": ["novel architectures from 2024", "proprietary
    model names", "company-specific tools", "bleeding-edge
    research terms"],
  "domainBackground": "Computer Science academia with focus on NLP
    /AI research"
```

```
}
```

Input: "undergrad freshman no programming experience"

Output: {

```
  "expandedDescription": "First-year undergraduate student with no
    prior programming or computer science background. Familiar
    with basic technology use (smartphones, apps, social media)
    but unfamiliar with technical concepts, programming
    terminology, or how software systems work.",
  "inferredExpertiseLevel": 1,
  "knownConcepts": ["apps", "websites", "social media", "AI tools
    like ChatGPT", "smartphones", "basic internet concepts"],
  "likelyJargon": ["programming terms", "algorithms", "databases",
    "machine learning", "neural networks", "APIs", "coding
    concepts"],
  "domainBackground": "General education, non-technical"
```

```
}
```

Be specific about what they would know vs. what would be jargon.
Focus on their domain expertise.

.1.1 Prompt for Jargon Detection

Listing 2: Prompt for jargon detection used in ReSlide

Analyze this slide content for jargon terms that would confuse the specified audience.

AUDIENCE PROFILE:

```
- Original Description: "${expandedContext.originalDescription}"
- Detailed Profile: ${expandedContext.expandedDescription}
- Expertise Level: ${expandedContext.inferredExpertiseLevel}/5
- Domain Background: ${expandedContext.domainBackground}
${presentationContext ? '- Presentation Context: ${
  presentationContext}' : ''}
```

WHAT THIS AUDIENCE KNOWS:

```
${expandedContext.knownConcepts.length > 0
  ? expandedContext.knownConcepts.map(concept => '- ${
    concept}').join('\n')
  : '- No specific known concepts provided'}
```

LIKELY JARGON FOR THIS AUDIENCE:

```
${expandedContext.likelyJargon.length > 0
  ? expandedContext.likelyJargon.map(jargon => '- ${
    jargon}').join('\n')
  : '- No specific jargon areas identified'}
```

SLIDE CONTENT TO ANALYZE:

```
Title: ${slideTitle || 'Untitled'}
Content: ${slideText}
```

CRITICAL INSTRUCTIONS:

1. Use the detailed audience profile to determine what would be jargon
2. If a term is in the "WHAT THIS AUDIENCE KNOWS" list, it's NOT jargon

3. If a term is similar to items in "LIKELY JARGON" list, it probably IS jargon
4. Consider the expertise level ($\text{\$}\{\text{expandedContext.inferredExpertiseLevel}\}/5$) carefully
5. Only flag terms that would genuinely prevent understanding

EXPERTISE LEVEL GUIDELINES:

- Level 1-2: Most technical terms are jargon, but common tech words (AI, app, website) are okay
- Level 3: Specialized and domain-specific terms are jargon
- Level 4: Only cutting-edge or highly specialized terms are jargon
- Level 5: Only the most novel, bleeding-edge, or extremely specialized terms are jargon

IMPORTANT:

- For professors/experts (Level 4-5): Very few terms should be jargon
- For beginners (Level 1-2): Many technical terms will be jargon
- Focus on the audience's specific domain background: $\text{\$}\{\text{expandedContext.domainBackground}\}$

Respond with JSON only (no markdown):

```
{
  "jargonTerms": [
    {
      "term": "exact term from text",
      "definition": "Clear explanation appropriate for this audience",
      "alternatives": ["simpler alternative 1", "accessible phrase 2"],
      "startIndex": number,
      "endIndex": number
    }
  ]
}
```

.2 Evaluation Study Details

.2.1 Screening Questions

Definition for presentation for the following screening questions:

- Included **planned content**
- Had a **clear and purposeful structure**
- Used **visual aids** (like slides)
- Were **tailored to a specific audience, setting, time limit, or goal**.

Presentation skill assessment

1. How often did you give presentations in the past 3 months?
 - 0
 - 1-2 times
 - 3-5 times
 - 6-10 times
 - More than 10 times
2. How confident are you in your overall presentation experiences?
 - Not at all confident
 - Slightly confident
 - Moderately confident
 - Very confident
 - Extremely confident
3. Have you received any formal training in public speaking or presentation skills?
 - No
 - Yes, informal trainings (e.g. workshops, guides, or feedback from others)

- Yes, formal instructions (e.g. a course, debate/public speaking club)
 - Other:
4. About how many years have you been giving presentations (e.g. in classes, workshops, or conferences) since post-secondary education?
5. Which best describes the typical audience for your past presentations ? (Choose the most common)
- Mostly peers or classmates
 - Mixed audiences (e.g. peers, instructors, or professionals)
 - Primarily professional audiences (e.g. clients, conferences, external stakeholders)
 - Other:
6. I have presented to a large audience (e.g., more than 50 people).
- Never true
 - Rarely true
 - Sometimes but infrequently true
 - Neutral
 - Sometimes true
 - Usually true
 - Always true
7. How often did you use slide creation software to prepare presentations?
- Never
 - Rarely
 - Sometimes
 - Often
 - Always
8. Which slide creation tools were you familiar with? (Select all that apply)
- Microsoft Power Point
 - Google Slides

Table 1: Demographic information for within-subject participants

Participant ID	Age	Gender	Occupation or Field of Study
C1	25	Woman	Software Developer
C2	24	Woman	NLP
C3	33	Woman	Data Science
C4	32	Woman	HCI-AI
C5	22	Woman	Student, PhD
C6	23	Man	Computer Science
C7	24	Woman	Computer Science
C8	24	Man	Computer Science
C9	25	Woman	Student, Master's
C10	24	Woman	Student, PhD
C11	25	Man	Student, PhD
C12	26	Woman	Data Science

- Apple Keynote
- Prezi
- Canva
- Other

.2.2 Demographic information for with-subject participants

.2.3 Demographic information for exploratory participants

.2.4 Custom Likert Scale Questions

Contextual Constraints related

1. I felt the slides I created captured constraints for time, audience, and presentation intent.
2. The system helped me stay mentally aware and structured my slides within the time limit.
3. The system helped me map my communicative intent to the presentation content.

Table 2: Demographic information for exploratory participants

Participant ID	Age	Gender	Occupation or Field of Study
E1	25	Man	Computer Science
E2	39	Man	Data Science
E3	28	Man	Computer Science
E4	34	Woman	Teaching Mathematics/Studied Finance
E5	34	Man	HCI
E6	25	Woman	Computer Science
E7	25	Woman	Computer Science
E8	29	Woman	Computer Science

4. The system increased my awareness of matching language to the target audience’s needs.
5. The system helped me express ideas in a way that’s clear and relevant to my audience.

Narrative Coherence

1. The system helped me structure and organize my thoughts while making the presentation.
2. I found it easy to plan my overall presentation flow.
3. I found it easy to switch between/adjust both the overall structure and detailed content organization during authoring.

Presentation Content Reuse

1. I found it convenient and efficient to reuse slides using this system.
2. It was easy to understand how to reuse slides with the system.
3. The system gave me control over how my existing slides were reused and updated.