

Malleable Overview-Detail Interfaces

Bryan Min

University of California San Diego
La Jolla, California, USA
bdmin@ucsd.edu

Yining Cao

University of California San Diego
La Jolla, California, USA
yic069@ucsd.edu

Allen Chen

University of California San Diego
La Jolla, California, USA
aschen@ucsd.edu

Haijun Xia

University of California San Diego
La Jolla, California, USA
haijunxia@ucsd.edu

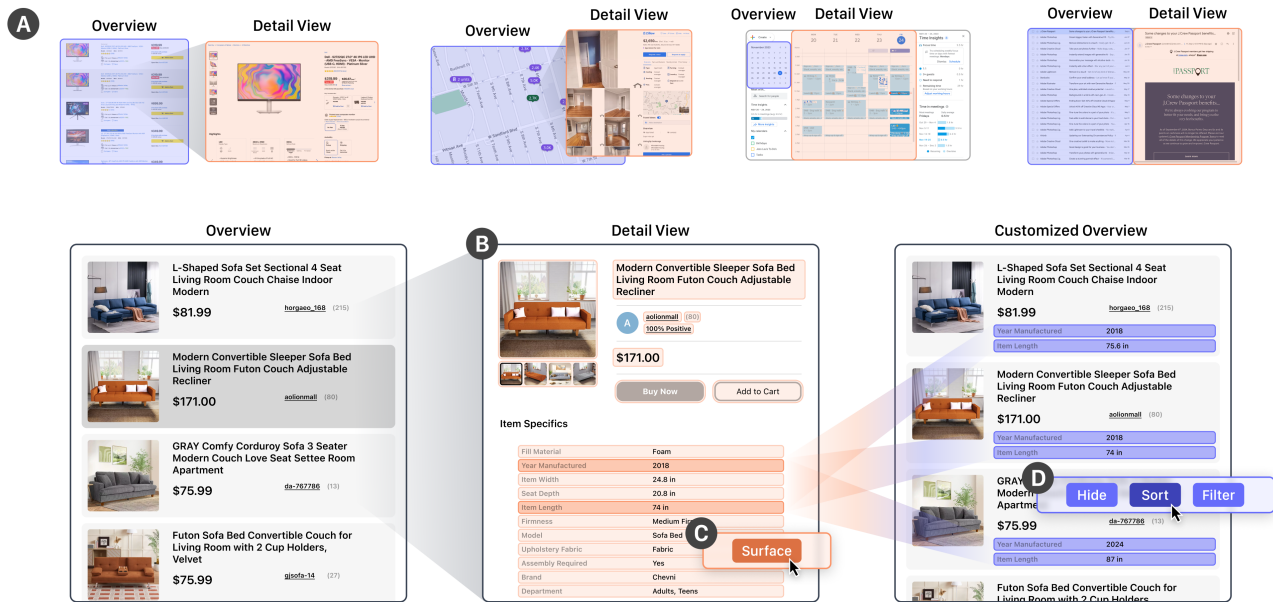


Figure 1: (A) We perform a content analysis of overview-detail interfaces (blue: overview; orange: detail view) to understand its variations in the wild, finding three key dimensions: *content*, *composition*, and *layout*. (B) Based on these dimensions, we provide a set of interactions for end-users to customize their overview-detail interface. Among them, we contribute a novel technique, *Fluid Attributes*, enabling users to (C) surface attributes to the overview that only exist in the detail view and (D) operate on attributes they surface.

Abstract

The overview-detail design pattern, characterized by an overview of multiple items and a detailed view of a selected item, is ubiquitously implemented across software interfaces. Designers often try to account for all users, but ultimately these interfaces settle on a single form. For instance, an overview map may display hotel prices but omit other user-desired attributes. This research instead explores the malleable overview-detail interface, one that end-users can customize to address individual needs. Our content analysis of

overview-detail interfaces uncovered three dimensions of variation: content, composition, and layout, enabling us to develop customization techniques along these dimensions. For content, we developed Fluid Attributes, a set of techniques enabling users to show and hide attributes between views and leverage AI to manipulate, reformat, and generate new attributes. For composition and layout, we provided solutions to compose multiple overviews and detail views and transform between various overview and overview-detail layouts. A user study on our techniques implemented in two design probes revealed that participants produced diverse customizations and unique usage patterns, highlighting the need and broad applicability for malleable overview-detail interfaces.



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

CHI '25, Yokohama, Japan

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1394-1/25/04

<https://doi.org/10.1145/3706598.3714164>

CCS Concepts

• **Human-centered computing** → **User interface management systems**; **Interaction techniques**.

Keywords

Overview-Detail Interfaces, End-User Customization, Malleable Interfaces, Interface Design Patterns

ACM Reference Format:

Bryan Min, Allen Chen, Yining Cao, and Haijun Xia. 2025. Malleable Overview-Detail Interfaces. In *CHI Conference on Human Factors in Computing Systems (CHI '25), April 26–May 01, 2025, Yokohama, Japan*. ACM, New York, NY, USA, 25 pages. <https://doi.org/10.1145/3706598.3714164>

1 Introduction

The overview-detail design pattern describes interfaces that typically consist of two views: an overview that presents a collection of items along with their key attributes and a detail view that provides an in-depth look at a selected item, including many, if not all, of its attributes [63, 69]. This pattern enables users to efficiently browse through a large collection of items and delve into specific ones of interest, facilitating one of the most essential operations in information tasks [25, 48]. This is perhaps why the overview-detail pattern has become ubiquitous throughout information systems, including shopping, booking, and rental sites, calendars, email clients, and numerous others (Fig. 1A).

Typically, it is the designers and developers that determine how the overview-detail pattern is instantiated for a specific context, considering factors such as the layout and interaction between the overview and the detail view, as well as the different sets of attributes to present in these views. While they strive to create the best design for their general user base, a single instantiation of the overview-detail pattern is inevitably not enough to match the diverse needs of all users and their usage contexts. For instance, separating the detail view to another page, such as in most shopping websites, can reduce clutter in the interface but incur repetitive view switching (Fig. 2A). Meanwhile, displaying only the price of hotels on a map improves clarity but omits other attributes—such as ratings or the number of beds—that some users may prioritize (Fig. 2B). Yet, few modern interfaces are *malleable*, forcing the end-users to work with an interface set by their developers, regardless of the end-users' own needs and preferences.

At its core, the overview-detail design pattern is intertwined with fundamental challenges in information foraging, sense-making, and decision-making [59]. First, the overview may not present the right abstraction of individual information entities to meet users' specific needs, forcing them to forage for details from other views [54]. Second, a system may hinder sense-making if the representation of the collection of items it provides does not adequately address users' needs [34, 63, 76]. Third, the spatial arrangement and interaction between multiple views may hinder necessary navigation among the different views. Therefore, the key question we explore is: **instead of letting developers dictate how information is abstracted and presented, can we make the overview-detail design pattern malleable?** That is, can we empower end-users to flexibly customize overview-detail interfaces to form their own abstractions and spatial arrangements of information?

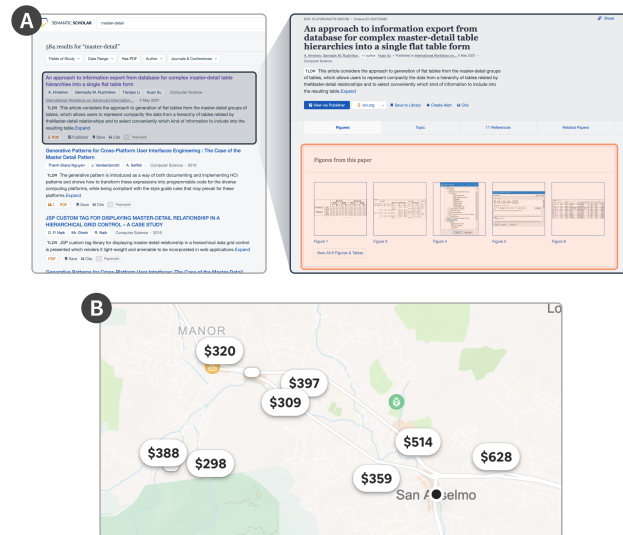


Figure 2: (A) Semantic Scholar includes a paper's title, authors, abstract, and others, but does not include figures in the overview. (B) Airbnb displays pins with prices on a map, representing the nightly prices of each location. Users must click each pin to view details not shown in the overview, such as the paper figures, ratings, or number of bedrooms, and navigate through each full page of details.

We break this goal into three research questions:

- [RQ1] What is the **design space** of various implementations of this design pattern in real-world contexts, and how can they inform the development of malleable overview-detail interfaces?
- [RQ2] What are the **interface designs and interaction techniques** that can enable end-users to flexibly transform an overview-detail interface?
- [RQ3] To what extent do end-users **prefer and utilize** malleable overview-detail interfaces to achieve their information tasks? What are their novel **usage patterns**? Is the malleability we provide enough?

We took the following steps to address these questions. First, we conducted a content analysis with a diverse sample of 303 instantiations of the overview-detail pattern and identified three key dimensions: *content*, *composition*, and *layout* (RQ1). Second, we developed interaction techniques that allow users to customize overview-detail interfaces along these three dimensions (RQ2). For the content dimension, we developed a set of interactions, which we label as *Fluid Attributes*, that enables users to surface new attributes from the detail view to the overview (Fig. 1C), hide unwanted attributes in the overview, and use any selected attribute to filter and sort items (Fig. 1D). Users can also prompt AI to perform these customizations automatically, reformat existing attributes, and create new ones. For the composition and layout dimensions, we developed various options for end-users to select, such as setting the number of overviews in the interface and selecting layout options. We demonstrate how our presented customizations could be designed and implemented through multiple applications.

Finally, we conducted a user study (N=12) to measure the utility and observe usage patterns of malleable overview-detail interfaces through two design probes in shopping and hotel booking tasks (RQ3). We found that participants, given the same tasks, desired and produced diverse customizations along all three dimensions. Furthermore, we discovered that participants' usage patterns resembled common tendencies in managing personal digital information and structuring tasks, demonstrating how malleable overview-detail interfaces integrate into broader, everyday workflows.

In summary, our work contributes:

1. A content analysis of the overview-detail design pattern that reveals three key dimensions of the design space.
2. Interaction techniques enabling end-users to customize overview-detail interfaces, offering diverse layouts and compositions of overviews and detail views, attribute manipulability across views, and AI-driven customizations.
3. A user study demonstrating that participants effectively utilized malleable overview-detail interfaces to address diverse needs and revealing emergent usage patterns that inform design suggestions for future implementations.

Our vision is to make all aspects of interfaces malleable. We approach this vision by tackling *one design pattern at a time*—analyzing variations of a single design pattern and exploring generalizable customization techniques that extend beyond their originally intended contexts. We hope this paper not only encourages future interfaces to implement malleable overview-detail interfaces, but also for future research to adopt this paper's approach to explore how more design patterns can become malleable.

2 Related Work

2.1 The Overview-Detail Design Pattern

The overview-detail design pattern describes interfaces that provide two views: an overview that contains a collection of items and a detail view that provides full details of a selected item [63]. This pattern, also referred to as “master-detail” [46, 47, 62] and “director-detail” [53], visualizes information for users to view a large collection of items first before uncovering their details. If overview-detail interfaces are designed with progressive disclosure in mind, they can offer a more approachable interface while still allowing users to access additional details [48, 63].

Prior research explored several variations of the overview-detail pattern in user interfaces. Case studies have explored how overview-detail interfaces can be nested to create hierarchical views [46, 62], how several detail views can be used simultaneously [60], and how different screen sizes may necessitate varying arrangements of the overview and detail views [47, 62]. Cockburn et al. have also explored how overviews and detail views can be arranged horizontally, vertically, and even along the “z-plane” by layering nested overviews and detail views along the z-axis, navigable by zooming. This z-axis arrangement also describes the structure of zoomable user interfaces such as Pad++ [9, 16]. A relevant but distinct design pattern also explored is the “focus+context” interface [36, 64], where focused content is placed seamlessly with the rest, oftentimes by reducing the view of the context space by scaling down items [8], skewing the views in 3D [43], and reducing the display's resolution [7].

These works demonstrate the diverse implementations of the overview-detail design pattern that researchers have developed, resulting in unique design choices tailored to specific contexts. However, our research aims to explore designs that enable users to flexibly customize the overview-detail interface according to their individual needs.

2.2 The Need for Malleable Overview-Detail Interfaces in Information Tasks

Many information tasks, such as foraging [54], sensemaking [59], and decision-making [13], involve an iterative process of searching, browsing, collecting information across various sources, making sense of collected information through various representations, and collecting new information based on new insights [55, 59]. Throughout this process, users need representations that help them effectively organize and preview a large number of information entities and quickly determine relevant ones by inspecting their details. The overview-detail pattern directly responds to such needs and, therefore, is ubiquitous in information systems, including information visualization systems [6, 12] and academic literature review tools [29, 50].

As mentioned earlier, a single instantiation determined by developers is not enough. Therefore, many research systems that have been created to support information tasks are essentially providing additional overview-detail instantiations to compensate what could not be supported with existing ones. For example, prior work has explored comparison tables as a means to better organize decision-making criteria [13], hierarchical organizations of collected snippets to facilitate information management and re-finding [28, 29, 37], spatial canvases that scaffold the user's search and synthesis processes [51, 57, 68], and node-link diagrams to contextualize users with their past browsing activities [5].

Another line of work, recognizing that a single instantiation is not enough, proposes providing multiple instantiations, such as those that enable multiple levels of details or transforming information representations. Pad++ introduced semantic zooming, which offers multiple levels of detail based on different zoom levels, providing an alternative to geometric scaling techniques used in traditional zoomable interfaces. Recent research has also built on this idea to generate these different levels of detail with AI [17, 65, 66]. Victor similarly illustrated multiple representations through the ladder of abstraction, demonstrating how users can move between representations of various details to gain a more comprehensive understanding of a system [70]. WritLarge, on the other hand, demonstrated techniques for moving between multiple axes of representation—semantic, structural, and temporal—on a digital whiteboard [74]. Graphologue also explored multiple representations, generating both text and node-link diagrams from an LLM in real-time to facilitate information comprehension [27]. We share the spirit of these systems that aim to provide multiple representations by providing end-users the ability to transform between different variations of overview-detail interfaces.

2.3 Designing End-User Customizable Information Systems

We share the research community’s long-standing vision of creating malleable information systems for end-users to easily and expressively customize them to suit their preferences and information needs. One research theme explores how to design the architecture of information systems to make them inherently customizable. For example, Haystack proposed a unified data model and a customizable interface layer to enable users to integrate diverse information sources and customize what information is presented to enable personalized information management [2, 30]. Smalltalk focused on enabling users to directly modify the appearance and behavior of any object in the GUI [32]. HyperCard explored enabling users to create dynamic and interactive content through a scripting language and graphical user interfaces [4]. Finally, Webstrates treated web documents as dynamic media, enabling users to create custom representations of the same DOM element while facilitating real-time, dynamic collaboration [35].

Another line of work explores the extension of existing systems to make them end-user customizable, particularly with web systems thanks to their accessible DOMs. For instance, research has explored enabling end-users to directly customize a web interface by modifying the CSS attributes for website prototyping [33, 38, 67], embedding annotations between DOM elements to facilitate in-situ note-taking [58], and filtering, sorting, and making spreadsheet computations to personalize lists of items shown on a webpage [26, 41]. Research has also explored the potential of web mashups and automations enabled by direct manipulation and programming-by-demonstration techniques to automate workflows and reduce context-switching [11, 21, 22, 40, 45, 77].

Most of the above work, however, focus on exploring the underlying information architecture or the technical approach to enable customization for specific activities, applications, and platforms, rather than concretely and conceptually scoping how broadly generalizable their customizations are. For example, it is not readily obvious how approaches taken by personal information management systems (e.g., Haystack [30], Notion [49]) should be generalized to shopping websites. Moreover, it is unclear how systems like Sifter [26], which recognize a common DOM structure for list views, can generalize to other list variations or even extend to different DOM structures. In contrast, our research focuses on investigating the overview-detail design pattern that is ubiquitous across diverse information systems and platforms. The design space of the overview-detail pattern that we propose allows interface designers to easily recognize the variations of overview-detail interfaces that exist in the wild and also serve as a comprehensive guide of what all the possible customizations any overview-detail interface variation could theoretically support.

2.4 Manipulable Attributes in Software Interfaces

We build upon extensive work that has explored how manipulable attributes can be utilized in a variety of domains. First, data visualization focuses on encoding data into visuals, and manipulating data attributes has been a core technique that enables users to

create, customize, and interact with data visualizations [42, 75]. Researchers have also used manipulable attributes to support creative tasks, enabling users to compose new graphical styles [72], search for related graphical objects based on attributes [73], and make color attributes manipulable across multiple representations [15]. Personal information management systems such as Haystack [31] have enabled users to manipulate attributes to modify how their personal information is presented, and database- and table-oriented software such as MS Access [44], Filemaker Pro [18], Notion [49], and Airtable [3] have also allowed the end-user to show and hide attributes on their interface. Additionally, our content analysis of overview-detail interfaces has found e-commerce sites that allow end-users to show or hide a fixed set of attributes (Section 3.3).

We see the overview and the detail view as two different containers of different sets of attributes, and by making both the attributes and views accessible and manipulable, we can enable end-users to flexibly customize them. Our formulation of the malleable overview-detail design pattern acts as an amplifier, allowing us to leverage all prior techniques for making attributes manipulable for the overview-detail pattern and making them generalizable beyond each of their originally intended scenarios and contexts. Additionally, we explore how AI can facilitate end-users in manipulating attributes within the overview-detail interface to achieve their diverse information needs, enabling users to prompt the interface to surface, hide, sort by, filter by, and select values in attributes.

3 Content Analysis

To better understand the various implementations of the overview-detail design pattern, we conducted a content analysis of overview-detail interfaces in desktop websites. Our content analysis was driven by the following questions:

- [Q1] What kinds of overview-detail interface arrangements exist in the wild?
- [Q2] What types of information are presented in the overview and detail view, respectively?
- [Q3] What aspects of overview-detail interfaces, if any, are already customizable by end-users?

We present our findings with a design space of three dimensions of variations in overview-detail interfaces (Q1, Q2) and report existing customizability options found in our collection (Q3).

3.1 Methods

3.1.1 Data Collection. The goal of our data collection was to collect a diverse and representative set of overview-detail interfaces that people regularly use. We initially collected the top 50 highest traffic websites in the United States based on data reported by Similarweb¹, a web analytics platform. However, as we analyzed this collection, we realized that many common websites were not included, such as flight booking or restaurant websites. To diversify our collection, we expanded to three more sources: 1) top 5 highest traffic websites in each of 33 web service categories² listed by Semrush³, another web analytics platform, 2) 20 randomly sampled

¹<http://similarweb.com/>. Accessed June 2024.

²Names of 33 categories listed by Semrush in Appendix A.1.

³<https://www.semrush.com/website/top/united-states/all/>. Data last updated May 2024. Accessed July 2024.

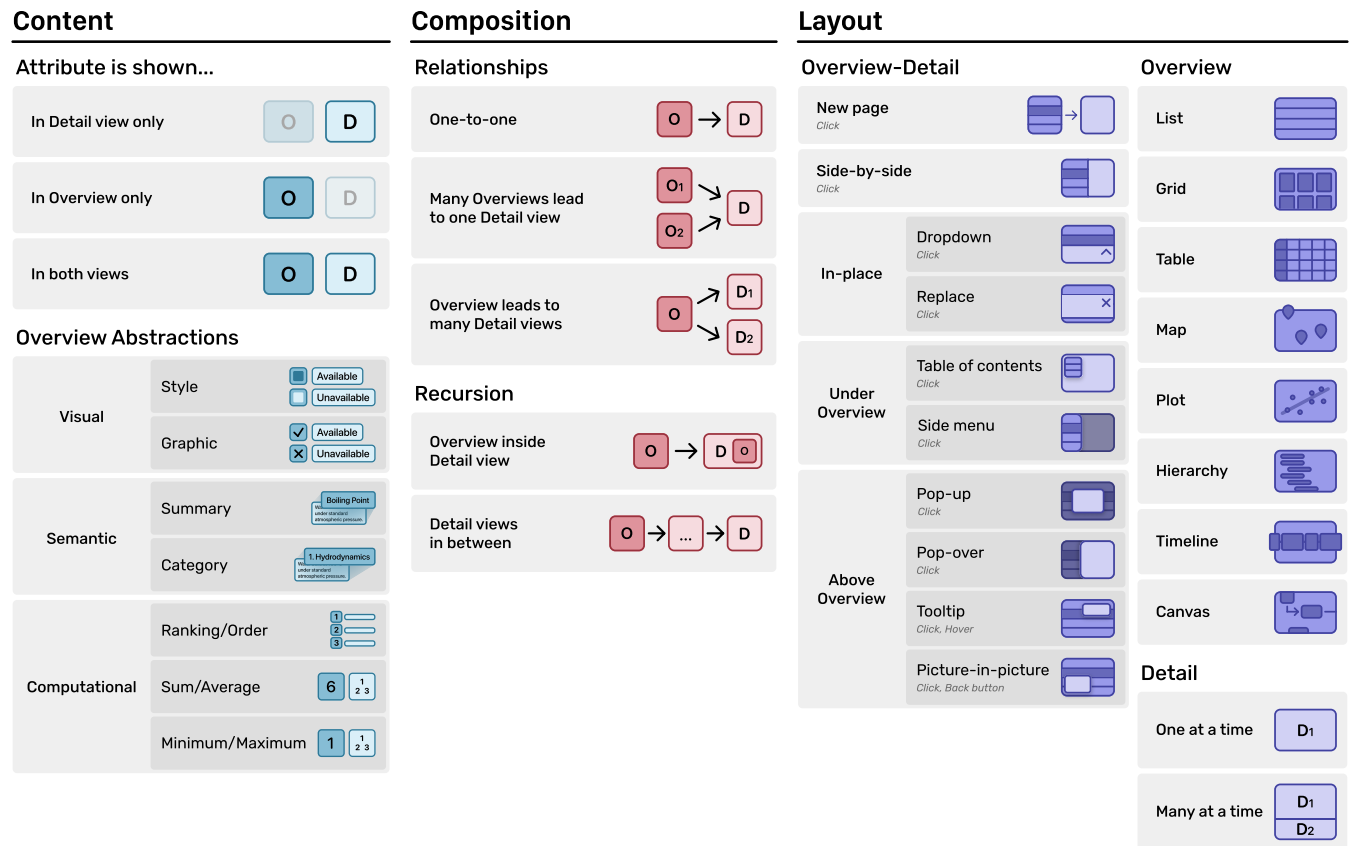


Figure 3: Design space of overview-detail interfaces on the Web: Content, Composition, and Layout.

websites from Wordpress’ showcase page to observe non-corporate websites such as blogs⁴, and 3) 20 randomly sampled websites from Awwwards’ gallery of “2024 Sites of the Year” to observe favorite designs from the UI/UX design community⁵. After filtering our list (e.g., removing duplicates), our final sample consisted of 156 websites. The first author examined all 156 websites, identifying 303 instances of overview-detail interfaces. Some websites featured multiple overviews linked to one detail view, or vice versa. For coding purposes, we defined each instance as a single overview-to-detail relationship. Thus, interfaces with multiple connections were recorded as separate instances.

3.1.2 Coding Procedure. We developed a codebook of overview-detail interfaces through an iterative open coding process involving collaboration among the first three authors. In the first round of open coding, we randomly selected 20 overview-detail instances from our dataset of 303, and the first three authors individually analyzed all 20 instances. The coders discussed and clustered their categories together to develop an initial codebook. In the next round, the first two authors independently coded a second sample of 20 instances, achieving an inter-rater reliability of 94.8% agreement between the codes, measured by Cohen’s kappa. They discussed

the remaining differences and revised the dimensions to reach a consensus. Given the high agreement rate, the authors determined that it was sufficient for the second author to complete the remaining coding, continuing to code in random samples of 20. After each round, the first two authors reviewed and discussed new potential dimensions and variations, revising the codebook as necessary. This process was repeated four times, stopping when no new dimensions or variations were found in the last three stages. In total, the authors analyzed 120 overview-detail instances.

3.2 Design Space

Our design space covers the following three dimensions:

- **Content** describes whether the information is presented in the overview, the detail view, or both, and how the attributes in the overview are abstracted from those in the detail view.
- **Composition** describes how multiple overviews and detail views are connected together.
- **Layout** describes the spatial arrangement of overview and detail views within the interface.

We describe notable cases found in these three dimensions and present the full design space in Figure 3.

⁴<https://wordpress.org/showcase/archives/>. Accessed July 2024.

⁵https://www.awwwards.com/websites/sites_of_the_year/. Accessed July 2024.

3.2.1 Content. Throughout our content analysis, we realized that some attributes were displayed only in the overview, some were displayed only in the detail view, and some were displayed in both. For example, the search results page for job openings on Indeed showed job titles, company names, and salary ranges in both views, while the overview omitted attributes like qualifications and benefits (Fig. 4A). Although we initially coded attributes as static information elements, we also observed this pattern in interactive elements such as ‘Save’ buttons, text fields, and menus, with some sites including them in both the overview and detail views (Fig. 4C), while others displaying them only in the detail view (Fig. 4D). Moving forward, we considered interactive elements to be attributes as well. We also found that most rental and housing platform overviews presented prices for each place as pins on a map, but in Zillow, if a place had more than one unit available, the pin displayed the number of units available instead (Fig. 4B). In text-heavy overview-detail interfaces such as messaging apps and online discussion forums, we noticed long blocks of text in the overview were either cut short via ellipses (Fig. 4E) or a fade-out gradient (Fig. 4F).

We further observed that some attributes in the overview served as an abstraction of information in the detail view, subsequently identifying three forms. First, we found visual abstractions, which were encoded into the style and visuals of the item. For instance, Ticketmaster used color shading to indicate the number of seats available per section, with sections that had more seats shown in a saturated blue, while sections with fewer seats appeared in lighter shades. Second, we found semantic abstractions that represented the contents in the detail view with summaries or categories. For instance, ChatGPT’s conversation interface generated a title that represented each conversation. Lastly, we found computational abstractions that computed attributes from the detail view into a value. For instance, Southwest’s website presents a calendar view of the cheapest flight prices for each day.

3.2.2 Composition. Additionally, we identified several overview-detail interfaces where multiple instances were composed together.

First, we found overview-detail relationships beyond just one-to-one. Many-to-one relationships presented multiple overviews such as a list overview and a map overview presented side-by-side (Fig. 5A), as well as different scopes such as a five-day list (Fig. 5B) and a 30-day calendar (Fig. 5C). We also found instances that synchronized navigation across the multiple overviews, where selecting an item in one overview would navigate users to the corresponding item in the other.

Second, we found instances of recursion—overview-detail interfaces nested within each other. For instance, additional overviews such as a list of articles appeared in the detail view of New York Times (Fig. 5D). This finding aligns with prior research that recognizes the variation of nested overview-detail views in hierarchical menu interfaces [16, 62]. We also observed cases of “intermediate” detail views, in which a link or button inside one detail view opened another detail view containing more or all attributes of the item. For instance, clicking on a user in a Discord server’s list of users opened a tooltip profile card, and clicking on the profile picture in the tooltip opened a full-screen popup, revealing further details about mutual friends and servers (Fig. 5E).

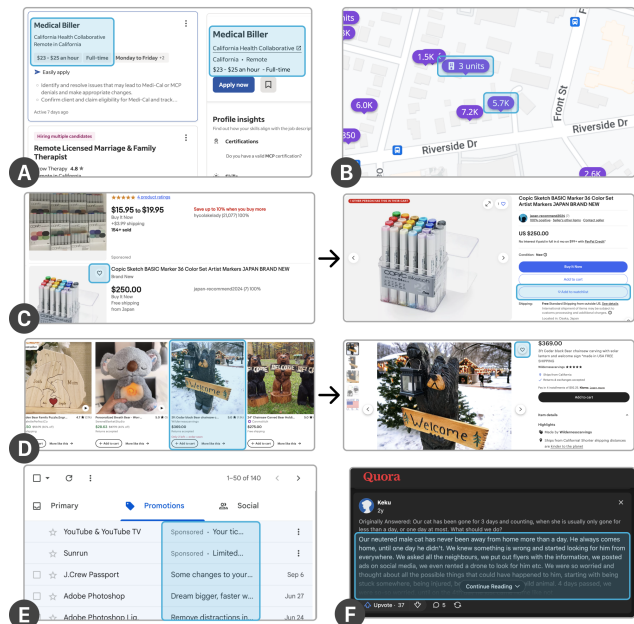


Figure 4: Example variations in the content dimension. (A) Indeed displays key attributes in both the overview and detail view. (B) Zillow’s map shows prices, but if multiple units are present, it instead shows the number of units. (C) eBay provides the “Add to Watchlist” button in both views, whereas (D) Etsy provides it only in the detail. (E) Gmail cuts off long text with ellipses while (F) Quora does so with a gradient.

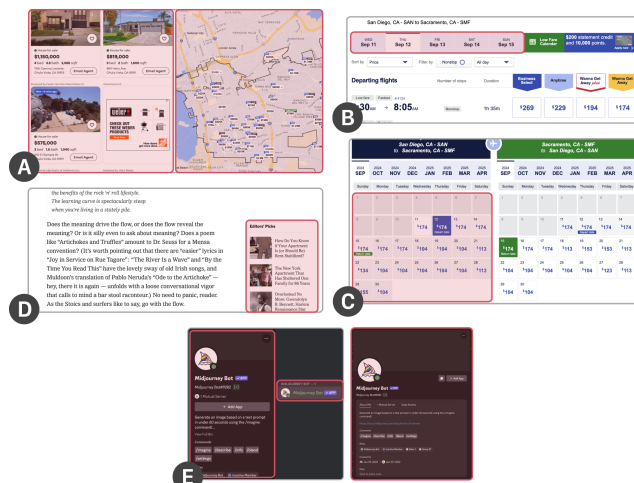


Figure 5: Example variations of the composition dimension. (A) Zillow shows a view composed of both a grid and a map. Southwest presents flight information through both (B) a 5-day table view and (C) a 30-day calendar view. (D) New York Times nests an overview inside a detail view with suggested articles. (E) Discord first shows a tooltip detail view of the user profile, and upon expanding, shows the full detail view.




Figure 6: Example variations in the layout dimension. (A) MLB’s player profile page maps batting points onto the baseball field; clicking a plot point opens a detail view of a video replay of the point. **(B)** Ticketmaster’s seating selection page also shows a map overview; when a particular seating section is selected, the overview list changes to match the corresponding section. **(C)** ESPN shows each player’s stats in a game, which can be opened as a dropdown. **(D)** Weather.com uses a similar dropdown layout, but unlike ESPN, multiple detail views can be opened simultaneously.

3.2.3 Layout. Our analysis revealed variations in the arrangement of overviews and detail views, the organization of items within overviews, and the number of items displayed in detail views.

Overview-Detail: First, we observed variations in the layout of overviews and detail views, where some websites arranged overviews and detail views side by side, while others opened detail views in a pop up or a dropdown.

Overview: We also observed variations of overview layout in the forms of lists, grids, and tables. Notable variations also included visualizations that plotted points that could be opened into detail (Fig. 6A) and sports stadiums to select seating sections (Fig. 6B).

Detail: We finally observed that detail views varied by how many were displayed at a time. Some instances only opened one detail view at a time, while others allowed users to open many at a time. For instance, ESPN allowed only one dropdown to be open at a time (Fig. 6C), while weather.com supported multiple (Fig. 6D).

3.3 Existing Customizability Features

In addition to the previously mentioned dimensions, we separately collected and coded existing customizability features for overview-detail interfaces. We found 8.33% (13/156) of websites provided customizations for overview-detail interfaces, with 2.56% (4/156) providing content customization, 0% (0/156) providing composition customization, and 7.69% (12/156) providing layout customization. All of the instances we collected were limited in customization, with the most providing eight layout variations.

We found Craigslist offered the most variations for customizing the overview-detail interface, with five options for customizing the overview layout and three for the overview-detail layout. The

choice of layout also indirectly influenced which attributes were displayed. For example, the grid view included an image thumbnail, while the list view did not, likely to maintain a more compact presentation. In the content dimension, we observed four websites that allowed users to customize which content attributes they wanted to view. Among these, eBay offered the most customization options, though it only allowed users to toggle six attributes.

3.4 Summary

In summary, our content analysis of 303 overview-detail instances in 156 websites identified three dimensions of variation: content, composition, and layout. We found that only 8.33% (13/156) of websites in our collection supported customization features, with those offering only limited options. Our collection of samples may not be fully representative of the entire set of information systems, and we may not have captured every instance of overview-detail interfaces within each sample. Thus, we do not claim the comprehensiveness of our design space and the frequencies of customizability features beyond our collection. Nevertheless, our analysis does provide insights into variations in real-world examples, offering contributions in understanding overview-detail interfaces.

4 Interaction Techniques for Malleable Overview-Detail Interfaces

In this section, we describe interaction techniques to customize overview-detail interfaces along the three dimensions identified from our content analysis.

4.1 Content

To customize overview-detail interfaces along the content dimension, we introduce Fluid Attributes, a set of interaction techniques that allows end-users to directly interact with detail attributes in the overview-detail interface. This enables users to control what, where, and how detail attributes are shown, as well as directly sort and filter by attributes. Fluid Attributes also unlocks further customizations in overview-detail interfaces when paired with generative AI tools by allowing users to query, generate, and reformat attributes. To support such interactions, overview-detail interfaces can provide a dedicated “Attributes Mode”, available to toggle on and off in a global toolbar on the website (Fig. 13D). Upon toggling on, the website highlights all attributes in the overview-detail interface, presenting users with affordances to manipulate and operate on the attributes.

4.1.1 Surfacing and Hiding Attributes (Fig. 7). In order to find details of items not shown in the overview, users must open the detail view of each individual item and locate the attribute in the interface. Fluid Attributes lets users select attributes from the detail view and directly surface them by clicking “Surface” (Fig. 7A). The selected attributes then populate in every item of the overview. On the other hand, if users want to hide attributes they are not interested in seeing in the overview, they can instead “Hide” those attributes (Fig. 7B). In essence, this interaction gives users the control to *get an overview of only the things they want* by manipulating attributes on their own, without the help from the developers.

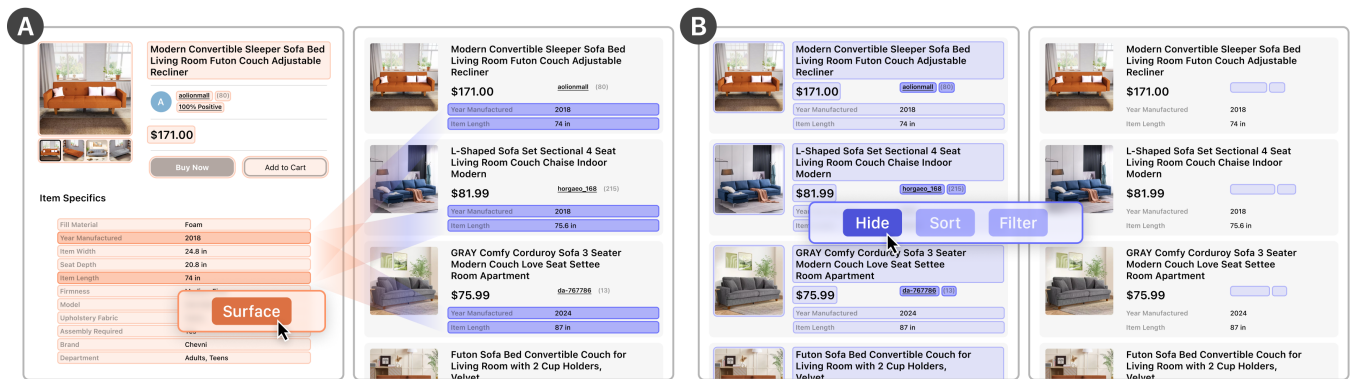


Figure 7: Upon entering “Attributes Mode”, (A) users can select attributes in the detail view to surface to the overview, and (B) hide selected attributes from the overview. Users can modify which attributes are shown and hidden through a context menu or by directly dragging selected attributes to their desired view.

Fluid Attributes is not limited to static attributes; it also includes dynamic and interactive elements such as widgets, buttons, and components. For instance, users can surface an image gallery component into the overview to view more images and even the *Add to Wishlist* button for each item on a shopping interface.

4.1.2 Sorting and Filtering with Attributes (Fig. 10). Rather than relying on a predetermined set of sort and filter functionalities given by the interface, Fluid Attributes provides the opportunity to make any attribute sortable and filterable. For example, users can surface the *Year Manufactured* attribute of a couch (Fig. 10.1), then directly sort them by the latest manufactured year (Fig. 10.2). Since Fluid Attributes also includes dynamic attributes, users can, for instance, sort by the boolean value associated with the *Add to Wishlist* button to “pin” items to the top of the search results.

4.1.3 AI-assisted Fluid Attributes. Fluid Attributes, when assisted by AI’s ability to process natural language queries to functionalities, demonstrates further interaction possibilities to customize the content in the overview-detail interface. Users can leverage AI through a prompt box provided by the interface.

Fluid Attributes via Natural Language Queries (Fig. 8). To customize attributes with natural language queries, users can for instance ask AI in a prompt box “How popular are these items?”, to which AI will identify relevant attributes (i.e., *Number of Watchers*, *Items Sold*) and surface them to the overview (Fig. 8A). This extends to the other operations in Fluid Attributes, such as adding a filter upon asking for “only couches with a linen texture” (Fig. 8B). Users can also ask to see “the back of the couches”, triggering AI to surface the *Image Gallery* attribute and select images that contain a back view of the couch (Fig. 8C). If AI fails to show this image, the user can always manually look through the images with the surfaced *Image Gallery* component.

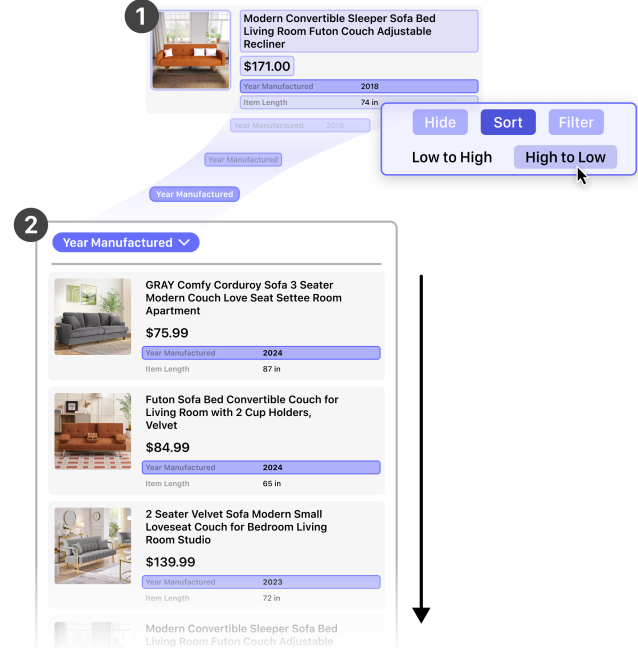


Figure 10: Since users can surface any attribute to the overview (1), this also gives them the ability to sort and filter by the surfaced attributes directly (2).

Generating New Attributes (Fig. 9). AI is also suited to generate new attributes from user queries. Users can ask for an attribute that may require computing values for each item in the overview. For instance, if a user asks for the “total cost” of an item, AI can sum values of attributes such as *Price* and *Shipping Fee* to create a new *Total Cost* attribute (Fig. 9.1). Users can hover over the newly generated attribute to reveal the source attributes the AI referenced (Fig. 9.2), and then surface these sources as well (Fig. 9.3-4). If AI cannot provide the value for an item, it responds “Not Specified”.

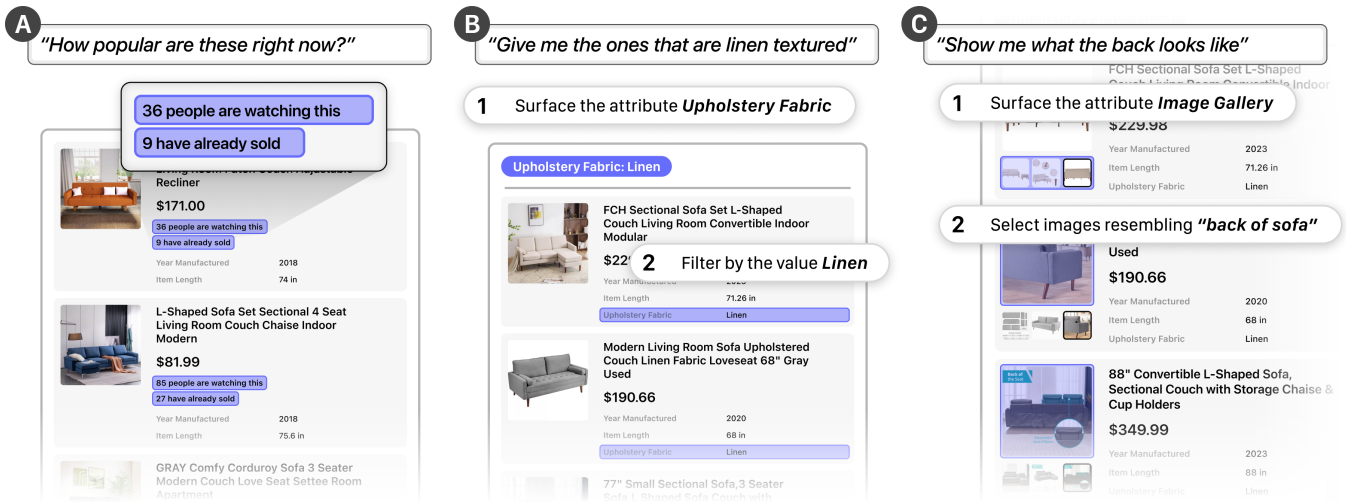


Figure 8: Users can additionally prompt AI to manipulate the attributes in the overview-detail interface. (A) Users can prompt AI to surface multiple attributes relating to the popularity of each product, (B) sort and filter for the values of a surfaced attribute, (C) and even surface specific values such as an image from the image gallery.

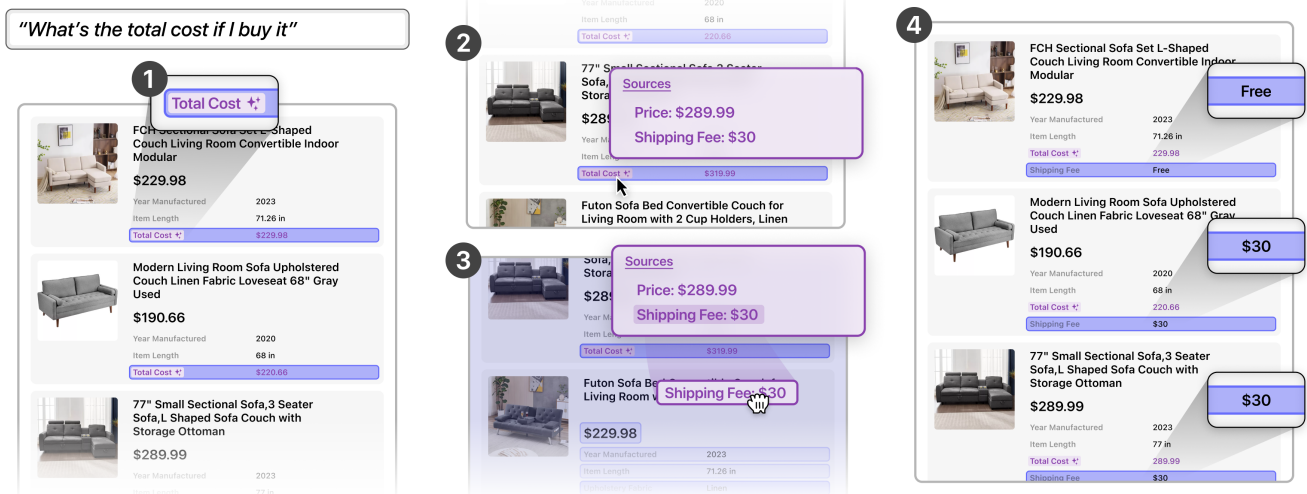


Figure 9: (1) Users can also generate new attributes with AI, for instance by asking for the total cost of the item, which sums the values of the *Price* and *Shipping Fee* attributes. (2) They can hover over the value to view its source attributes, and (3) also drag them into the overview, (4) surfacing the source attribute alongside the newly generated one.

Establishing Consistent Attribute Values (Fig. 12). Many attributes of items in online marketplaces are set not only by developers, but also by vendors. This results in many inconsistent attributes and values, such as different labels, currencies, units, and title formats. If users surface an attribute that is not found in other items, AI can scan through the item’s data to identify the attribute from other sources and fill them automatically (Fig. 12B). Users can also reformat attributes to make them inconsistent across all items. For instance, users can select the title attribute and either ask to “make them consistent” or make precise changes and ask to “show only the brand and the couch’s condition” (Fig. 12A).

4.2 Composition and Layout

To customize composition and layout, malleable overview-detail interfaces can provide variations identified by our the design space from Section 3. To make transforming between such variations most familiar to end-users, we provide dropdown menus for overview and overview-detail layouts and tabbed views to present and help manage multiple overviews (Fig. 11A, Fig. 13). By implementing these familiar interaction techniques, we aim to keep the interface customizations user-friendly.

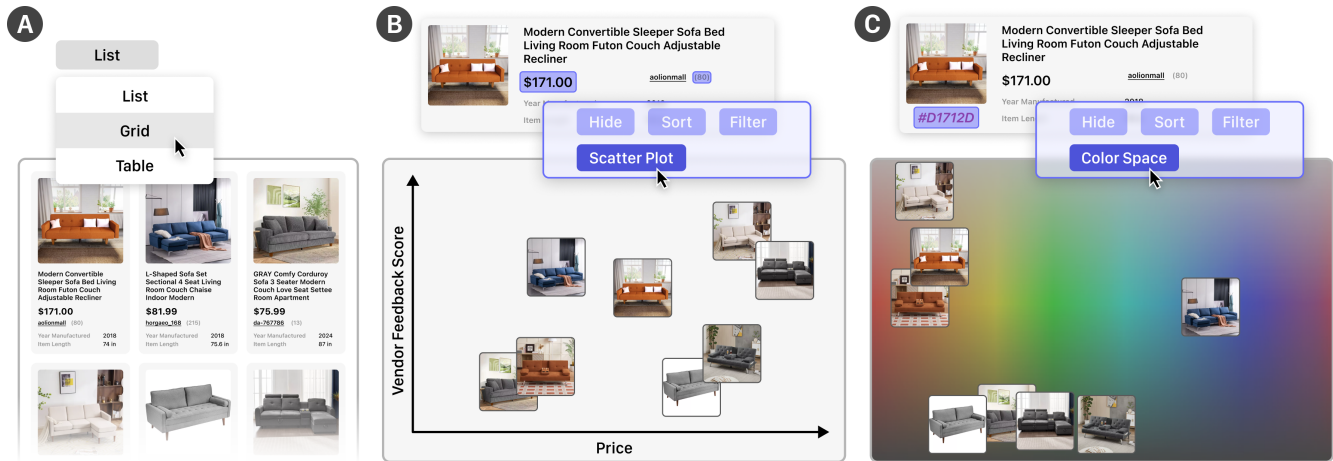


Figure 11: (A) Overview variations from the design space are available for the end-user through a dropdown menu. Fluid Attributes additionally allows users to invoke relevant overview representations through attributes. (B) For instance, users can select the *Price* and *Vendor Feedback Score* attributes to plot items in a scatter plot, or (C) generate a color value from the image to map items onto a color space.

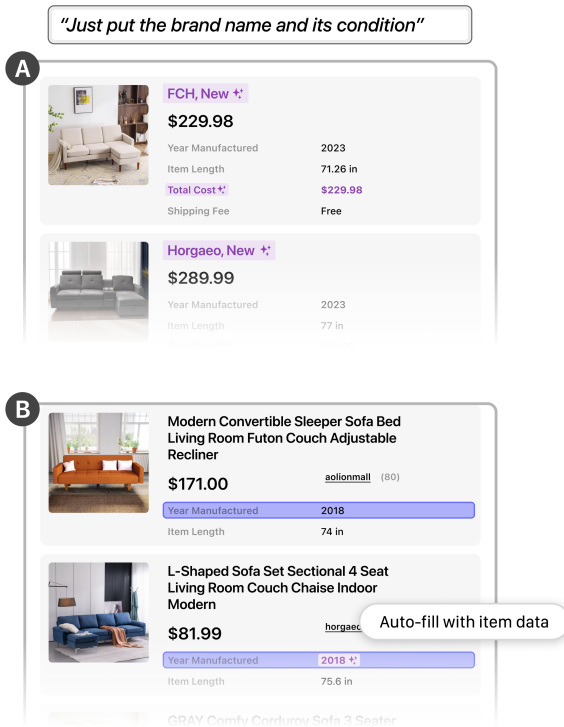


Figure 12: (A) Users can select the title and prompt AI to reformat the titles of items to make them more consistent. (B) Additionally, surfacing attributes from items that do not contain them will trigger AI to search and fill in the missing values to that attribute. If not found, AI provides the value “Not Specified”.

However, other overview-detail systems may support a greater variety of overview representations to meet diverse needs, which, when listed in a dropdown menu, can become large and overwhelming for the user. We recognize that attributes are closely associated with their overview representations, as they are mapped to specific visual elements or data points within the overview. Fluid Attributes leverages this property by enabling users to select desired properties and reveal relevant overviews in various representations. For instance, users can select the *Price* and *Vendor Feedback Score* attributes to view items in a scatter plot to consider the reliability of the vendors while comparing prices (Fig. 11B). Users can also generate a color attribute from the item image to map items onto spatial overviews such as a color space, allowing users to cluster and explore items by color, similar to Histomages [15] (Fig. 11C). A scatter plot and color space are just two examples of overviews; other malleable overview-detail interfaces can offer additional representations and more nuanced variations of scatter plots and color spaces. Instead of aiming to provide all possible overview representations, our goal is to offer a solution that allows users to intuitively and flexibly modify how their overview is represented.

5 Applications of Malleable Overview-Detail Interfaces

To demonstrate the generalizability of our presented customization techniques, we developed five overview-detail interfaces for various contexts and platforms: online shopping, hotel booking, managing emails, reading live articles, and ordering food on a mobile app. We developed the shopping and booking websites with higher fidelity than the others to use them as design probes for our user study. These two probes are described in more detail in Section 6.1.

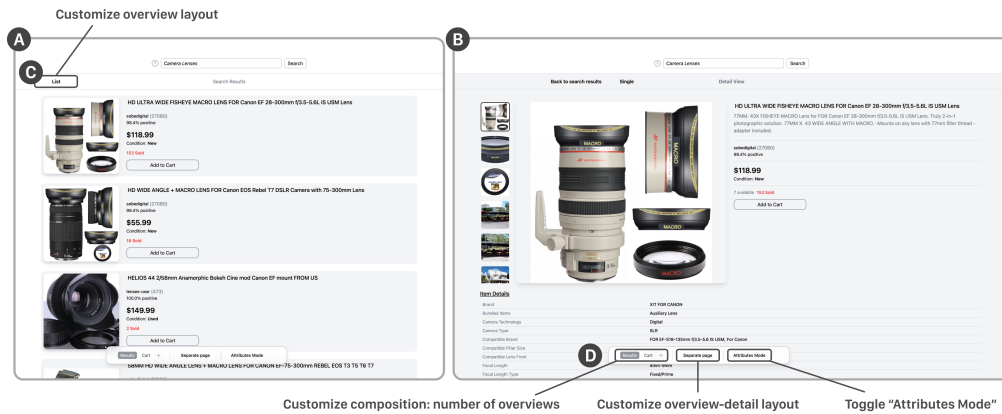


Figure 13: A screenshot of the shopping design probe, displaying (A) the overview list and (B) the detail view opened in separate pages. (C) Overview layout can be customized through a dropdown menu in the corner of each overview. (D) The toolbar provides customizations for the overview-detail, including adding and opening multiple overviews, changing overview-detail layout, and customizing attributes.




Figure 14: A screenshot of the hotel booking design probe, displaying (A) an overview list of bookmarked hotels and (B) an overview map presenting locations as pins. (C) A detail view of a selected hotel is opened as a popover on the left side of the map. (D) The toolbar presents a way for users to jump between different searched locations and customize attributes. (E) Surfaced attributes on the map populates the pin.

5.1 Scenarios

5.1.1 *Online Shopping (Fig. 13).* Online shopping websites commonly present a search result of items as a list overview with predefined attributes (e.g., item thumbnail, “Add to Cart” button) (Fig. 13A) and a detail view in a new page (Fig. 13B). However, this default requires users to repeatedly switch between both views to either view specific details or fully browse through individual items. A malleable overview-detail interface can instead provide a toolbar UI that presents options for the user to open detail views side-by-side or in-place with a dropdown (Fig. 13C). Additionally, users can add more overviews to save items in different collections (Fig. 13D). This can be used to compare sets of choices such as combinations of camera bodies and camera lenses to purchase. Furthermore, users can surface attributes in the overview by toggling on “Attributes Mode” to directly manipulate attributes and form custom abstractions of the interface (Fig. 13D).

5.1.2 *Booking Hotels (Fig. 14).* Hotel booking, home rental, and apartment search websites often provide a list view alongside a map view (Fig. 14A and B) to provide rich spatial information. However, these maps often present zero or one attribute for each location, with the only way to view these details being to click on the pin one at a time (Fig. 14C). Instead, a malleable overview-detail interface can enable the user to surface more details into the map, such as *Rating*, *Ranking*, and *Amenities* (Fig. 14D). Additionally, the user may customize different attributes for different overviews to leverage the benefits and affordances each overview representation provides. For instance, users can surface details about the number of reviewers in the map view to gauge which areas might be more popular (Fig. 14E), while they can surface details about amenities in the list view to easily sift through preferences by scrolling.

5.1.3 Reading and Managing Emails (Fig. 15). Oftentimes, emails contain important details in the text of their messages such as a response deadline. However, foraging through a large inbox for these details can grow extremely tedious. Using AI-assisted Fluid Attributes, users can prompt the interface to generate an attribute *Response Deadline* (Fig. 15A), surface it in the overview, and sort the emails by upcoming deadlines (Fig. 15B). Users can also choose to open emails in-place rather than side-by-side or on a new page, allowing them to fully read through important ones and quickly glance at less important ones.

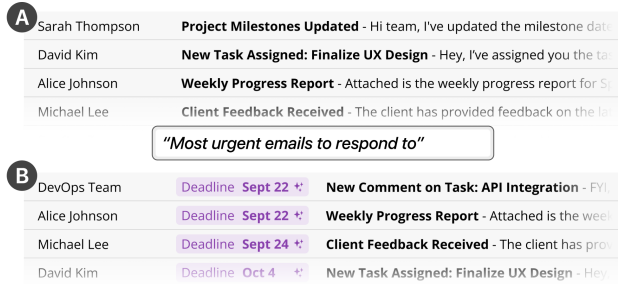


Figure 15: (A) Users can prompt AI to sort by emails with the most urgent deadlines. (B) Upon doing so, the attribute *Response Deadline* is surfaced and sorted by latest date.

5.1.4 Reading Live Articles (Fig. 16). Live news articles report on long-standing events such as presidential elections and sports games. These interfaces appear to aggregate shorter articles in a scrollable stack (Fig. 16A). However, it can be challenging for users joining in the middle of the event to quickly catch up on the timeline with a long document. To address this, users can select the *Published Date* attribute and bind it to a timeline overview, organizing all articles into a timeline. Users can also select which details to show in the timeline before diving into individual articles (Fig. 16B).

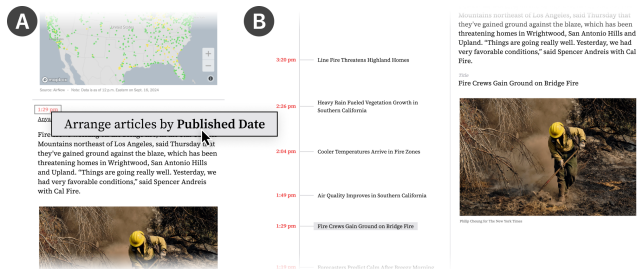


Figure 16: (A) Instead of scrolling through a long list of live updates, users can gain a broader overview of the event by selecting the *Published Date* attribute and transforming the layout from a list to a timeline. (B) Users can additionally generate a *Title* attribute to gain a summary of the events.

5.1.5 Mobile Applications: Food Delivery (Fig. 17). The overview-detail design pattern is also present across diverse platforms such as mobile devices. We recognize that our presented customization techniques can translate to these platforms as well. When ordering food on a mobile delivery app, some attributes are not shown in

the overview. For instance, Doordash shows prices and thumbnails of menu items to show what is popular (Fig. 17A), but do not show descriptions—which includes the calories. An implementation of Fluid Attributes in this app could allow calorie-conscious individuals to check which items best suit their needs by surfacing the calorie count into the overview. Users can easily invoke Fluid Attributes on mobile devices by long-pressing an attribute (Fig. 17B), which opens a context menu presenting available customizations, such as the “Surface” option (Fig. 17C). Tapping this option instantly reveals the calories for all items, eliminating the need to further switch between views (Fig. 17D).




Figure 17: Our presented techniques generalize across platforms such as mobile devices. (A) Upon tapping on a menu item on a mobile delivery interface, (B) calorie-conscious users can hold down on the description text of the item and (C) tap “Surface”. (D) This instantly reveals the ingredients and calories for every option.

5.2 Implementation Details

We implemented our desktop web applications in Next.js, a React web framework, while we implemented our mobile application in SwiftUI. We used eBay’s API to provide results in the shopping website and used Tripadvisor’s API to provide results in the hotel booking website. Our high-fidelity probes stored all customization settings on Firebase⁶ to log them for our user study. Other applications stored customizations locally on the application’s client.

The pipeline of the system diagram of our malleable overview-detail implementations is illustrated and detailed in Figure 18. The end-user can make customizations across the three dimensions, corresponding to the system’s three components. In our applications, content customization is implemented such that all data is fetched to the front-end, and the management of visible attributes is entirely client-side. However, other implementations can opt to fetch data attributes dynamically to minimize unnecessary database calls. Both composition and layout customizations are also managed on the client side.

We implement AI to customize the content dimension only, but it can also manipulate the other two dimensions. Since all customization actions are performed through individual settings within the system, these settings can be made accessible for AI to modify based on their relevance to the user’s prompt. This would also enable AI to sequence multiple customization operations together across different dimensions. The prompts used in our specific AI implementations are provided in Appendix A.2.

⁶<https://firebase.google.com/>. Accessed December 10, 2024.

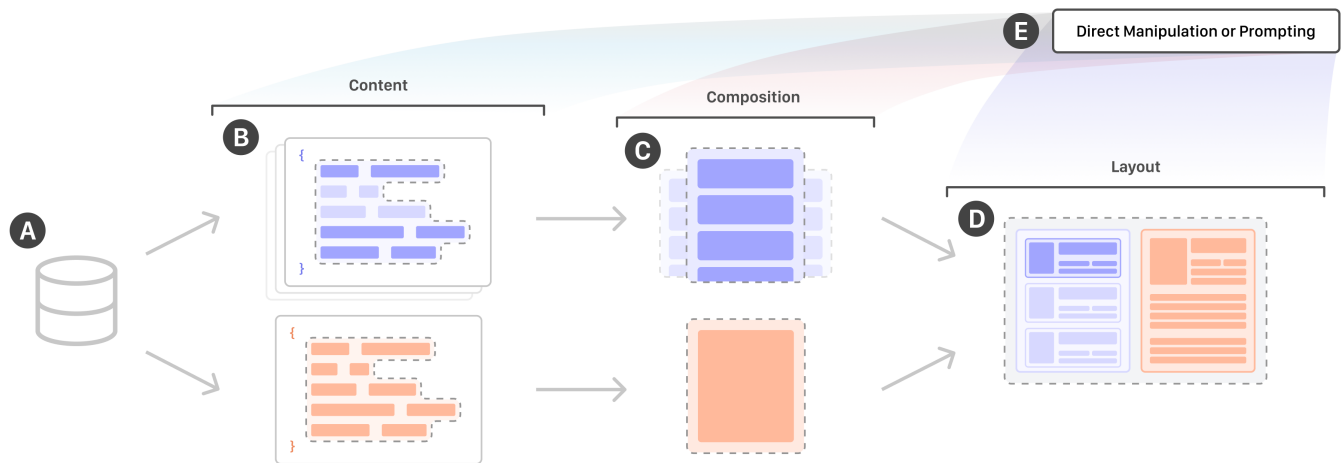


Figure 18: A system diagram of malleable overview-detail interfaces. (A) First, data for the overview and detail view is fetched from the system’s database. (B) The system then determines which attributes are included in the overview and detail view. This can either be determined from fetching specific attributes from the database or from the front-end client. (C) Next, the system determines the composition of the interface, such as the number of overviews to display. (D) Finally, the set of overviews and detail views are presented in the determined layout. (E) The end-user is able to customize these three stages through the interface by either directly interacting with the settings on the interface or by prompting AI.

6 User Study

We used our two high-fidelity applications (shopping and hotel booking) to understand how effectively users might customize malleable overview-detail interfaces and to observe emergent usage patterns as they utilize overview-detail interface customizations to best fit their needs and preferences. We aimed to answer the following questions:

- [RQ1] How effectively can users perform customizations on malleable overview-detail interfaces?
- [RQ2] What kinds of customization workflows and usage patterns emerge when interacting with malleable overview-detail interfaces?
- [RQ3] What additional customizations did users seek that were not supported by our design probes?

6.1 Design Probes

We selected shopping and hotel booking websites in order to approximate a common yet complex, real-world scenario of using malleable overview-detail interfaces. To focus our study on observing the use of the key customization features, we selected a subset of features that were demonstrated in Section 4.

In the shopping website’s content dimension, we included surfacing, hiding, sorting, and filtering by attributes, prompting AI to surface and generate new attributes, and auto-fill attribute values. Users can also sort and filter by attributes of each item. For composition, we enabled users to add, remove, and rename multiple overviews. For overview layouts, we provided list, grid, and table views, and for overview-detail layouts, we provided new page, side-by-side, dropdown, and pop-up layouts. For the booking website, we supported only the content dimension to further investigate

how users surface and hide attributes and to observe potential workflow differences compared to the shopping website. Upon surfacing attributes on the booking website, users were given the option to surface them to the map, the bookmarks list, or both. Screenshots of our probes in different configurations are in Appendix A.3.

6.2 Default Configurations

We configured our overview-detail interface probes with default settings to be intuitive for participants to quickly get accustomed to. Our shopping site displayed the search results as a list overview of items. The overview-detail layout was by default a new page, and the detail view displayed one item at a time. We additionally provided a second list overview named “Cart”, which displayed a grid layout of the items that were added to it. Each item in the overview displayed the following attributes: *Title*, *Thumbnail Image*, *Vendor Username*, *Vendor Feedback Score*, *Vendor Feedback Percentage*, *Product Price*, *Product Condition*, *Number of Products Sold*, and an *Add to Cart* button. Our booking site utilized two different overviews: a map overview and an overview list of bookmarks. Upon entering the interface, the map view displayed the *Rating* attribute of each hotel, and upon adding items the bookmark list, the *Hotel Name* attribute was displayed by default.

6.3 Procedure

We conducted a 90-minute study, which began with a brief introduction of overview-detail interfaces to familiarize participants with the concepts and terminology used throughout the study (<1 min). We then ran two 35-minute sessions, first on the booking task and second on the shopping task. We began with the booking task because it contained a smaller feature set which could more easily acclimate users to both websites. During each session, we first gave users a tutorial of our system, which included explanations of

each feature, as well as mini-tasks to acclimate participants to our probes and discover preferred customizations (~10 min). We then proceeded onto the main task, in which participants were given a task prompt and asked to think out loud [39] (~20 min). After each task, participants completed a questionnaire (7-point Likert scale) to share their experience with the probe (~5 min).

Our hotel booking task description and criteria was as follows:

You are traveling to the Bay Area for meetings. You will stay in two different locations, one night per stay and will need to find a hotel for each: San Francisco and San Jose.

- Your total budget is 350 USD.
- You want to eat at the hotel (or hotel room) as much as possible (strong preference for hotels with included breakfast or room service).
- You prefer modern and new hotels over traditional or old ones.
- You would like to stay one night at a premium or luxurious hotel.

Our shopping task description and criteria was as follows:

You are choosing two monitors: one for your home and one for your workplace.

- Your total budget for two monitors is 250 USD.
- Criteria for monitor at home:
 - You want a decent deal.
 - You want the highest resolution possible, 4k or maybe better.
 - The monitor should take up minimum space on your desk.
- Criteria for monitor at work:
 - The condition of the monitor is most important.
 - You want a display that's decently sized, but also with high pixel density (resolution / display size).
 - You don't need a monitor stand—assume the office already provides one.

At the end of the session, participants were interviewed to reflect and share their overall experience (~20 min). These interviews helped us understand their thoughts on customization preferences, strategies, limitations, and opportunities.

6.4 Participants

We recruited 12 participants (5 male, 7 female) from a local university through internal communication channels. Of the participants, five were graduate students and seven were undergraduate students. All participants reported having prior experience with online shopping and booking.

Initially, we asked participants to use their own laptops for the study so it is most familiar and comfortable for them, but after four participants, we deemed users may prefer a better setup. We offered the remaining participants an external monitor (4k resolution, 31" screen size), a keyboard, and a mouse. When offered, 7/8 participants chose to use the monitor. Participants were compensated 30 USD for 90 minutes of their time.

6.5 Measurements

To observe patterns in our tasks, we logged the types of customizations users made, the frequency with which they made specific customizations, and the attributes they surfaced or hid from the overviews. To capture context and insights into participants' behaviors, we recorded both screen activity and audio.

7 User Study Results

First, we observed that our participants made diverse customizations even when given the same shopping and booking tasks and the same starting interface configurations, suggesting that the customizations provided by malleable overview-detail interfaces are beneficial and desired. Many participants surfaced attributes that others had not, as illustrated in Figure 22. Some participants used multiple overviews for saving monitors into work and home lists (P1, P2, P6, P7, P12), while others created one intermediate space for candidate items (P3, P4, P8, P9, P10). While a majority of participants used the side-by-side overview-detail layout (P1, P2, P4-P12), P3 preferred to use the pop-up view, while P12 preferred both the dropdown and side-by-side views. Additionally, participants relied on a broad variety and combination of overview layouts between the list (P1~4, P7, P9, P10), grid (P2~5, P8, P10, P11), and table (P6, P11, P12). For instance, P4 used both grid and list views, while P6 used both table and list views together.

In the following sections, we report findings for each of our research questions.

7.1 RQ1. How effectively can users perform customizations on malleable overview-detail interfaces?

All participants completed both tasks within the provided time. They reported that the interfaces of both probes were easy to learn (booking: 11/12; shopping: 12/12) and operate (booking: 9/12; shopping 12/12), and most agreed it was useful to customize content (booking: 6/12 strongly agree; shopping: 10/12 strongly agree), composition (11/12), and layout (overview-detail: 11/12; overview: 12/12) (Fig. 19). All participants surfaced attributes on both websites, while 9/12 participants hid attributes on at least one of them, and 10/12 participants sorted or filtered by attributes on the shopping website. Most participants also customized the composition (10/12) and layout (11/12) of the shopping website (Fig. 20).

We also measured how often participants made customization operations. On average, participants made 1.80 customizations per minute during the booking task (SD=0.74) and 0.86 customizations per minute during the shopping task (SD=0.48). During their task, participants made an average of 22.08 content customizations (SD=12.29), 1.33 composition customizations (SD=0.84), and 2.17 layout customizations (SD=1.52). These results suggest that our participants were able to effectively utilize the customizations provided to them and understood the value of malleable overview-detail interfaces. P1 described how the customizability features allowed them to feel more in control over the interface, saying: "The ability to choose exactly what I wanted on my screen and close everything else was super nice. You can't really do [that] with a lot of sites. Like, it is what it is."




Figure 19: Participants’ responses to the utility and usability of interactions to customize the content, composition, and layout the hotel booking and shopping websites.

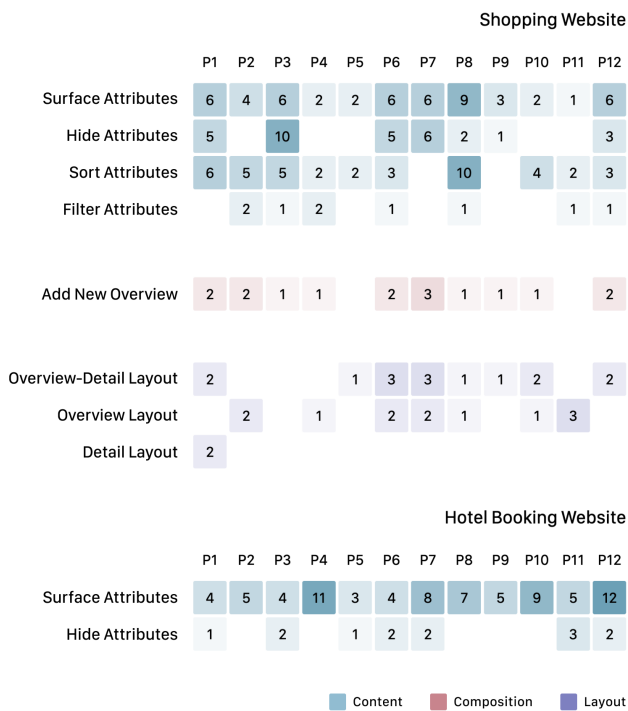


Figure 20: The number of customization operations made by each participant for each dimension. The shopping website supports customizations for all three dimensions, while the booking website supports surface and hide attributes. For the content dimension, we count how many new attributes users interacted with. For other dimensions, we count the total number of operations made.

Fluid Attributes reduces excessive context-switching. Participants explained that the utility of customizations came from its ability to avoid unnecessary navigation by bringing all details of interest into a single space. P3 stated, “with a lot of websites that you can’t really customize, you end up having to open the products one by one to look at all the specifications, which takes longer. I think it’s nice to be able to just see it all as you’re scrolling through.” P4 similarly described their experience with the booking website compared to other commercial websites, saying, “Websites like Booking.com, Expedia, and Airbnb don’t have these options. You have to literally look for each and everything. On the booking website, where only the options you want are displayed, it’s very easy.”

Fluid Attributes helps users align the interface with their decision criteria. Additionally, some participants shared that customizing content helped align the information they wanted to see on the website with their decision-making criteria. P12 said: “After reading the information, I kind of have some criteria in my mind of which features I probably want, so I will only concentrate on these features. [With the hotel booking website], I can just customize what information I really want to focus on. For example, I personally really prefer the ratings.”

AI-assisted Fluid Attributes reduces the need to sift through detail. We also observed almost all participants surfaced or generated attributes with AI (11/12). P12, who surfaced 5/6 attributes by prompting AI, found prompting the interface the most convenient as it gave them direct answers, stating: “I’m a very lazy person, so I’ll rely on this feature because I don’t want to click on each [item] and check whether they are 4k monitors.” P10 also appreciated the help from AI because they no longer felt “confined to certain characteristics” provided by each item. They further described how they found one shopping item that contained the *Material* attribute, but noticed it did not exist for all items. Instead of manually verifying whether items are missing the attribute or have it inconsistently named, they could use AI to verify it in a single action.




Figure 21: Logs of customizations made by participants throughout the shopping task for surfacing and hiding attributes, adding new overviews, and transforming overview and overview-detail layouts. While many participants made initial configurations, many also continued to make more throughout their decision-making processes, such as creating more overviews.

7.2 RQ2. What are the different workflows and usage patterns of malleable overview-detail interfaces?

Findings from RQ1 demonstrate that malleable overview-detail interfaces are both effective and preferred by our participants, as the supported customizations allowed them to achieve their desired configurations. In addition to enabling these configurations, we are also concerned with how to enable users to reach these configurations more quickly. Therefore, we aim to identify the types of overview-detail configurations users prefer, which could be provided as defaults to different user groups in future interfaces to reduce the number of customization they need to perform.

Customizing at the start and throughout. Logs of participants' customizations made throughout their shopping task are illustrated in Figure 21. While many participants began their task by making initial configurations of their overview-detail interface (P1–P5, P8–P12), many also continued to make customizations throughout (P1–P4, P6–P8, P10, P12). For instance, P3 initially surfaced attributes that best aligned with their criteria, but changed their set of surfaced attributes as they compared between more choices. Additionally, some participants added multiple overviews from the start (P1–P3, P6–P9, P12), while some added them ad hoc when they needed to save and directly compare between choices (P1, P4, P7, P10). P1, who continuously transformed their interface throughout their task, expressed the desire to stay in Attributes Mode, while P11, who preferred to “customize once and be done”, suggested a

dedicated customizations page to make their configurations in one place. Furthermore, P6 even expressed a desire to use their table and list configuration from the shopping website for the booking website as well. They noted that, for tasks like these, they almost always default to using a spreadsheet application as their go-to tool for managing and organizing information. To support users who continuously make customizations throughout their task, interfaces could reduce the friction of making these customization options with more direct manipulation interactions, keyboard shortcuts, or in-situ settings on the interface [56]. On the other hand, those who make initial configurations could benefit from a central settings page with available customization options, popular default configurations that users can access at any point, and if possible, saved customizations across different tasks and websites.

Common sequences of customizations. Most common sequences that we observed involved semantically related customizations, such as surfacing multiple attributes all relating to hotel food services or creating a new overview and immediately changing its layout. We also observed cases when participants sorted or filtered by values of attributes they had immediately surfaced (P2–4, P6, P8, P10, P11), suggesting a common *surface-then-sort* and *surface-then-filter* workflow. Interfaces can reduce the number of clicks by enabling interactions for creating and using composable operations, where developers can set defaults, end-users can construct their own, and AI can generate them as well.

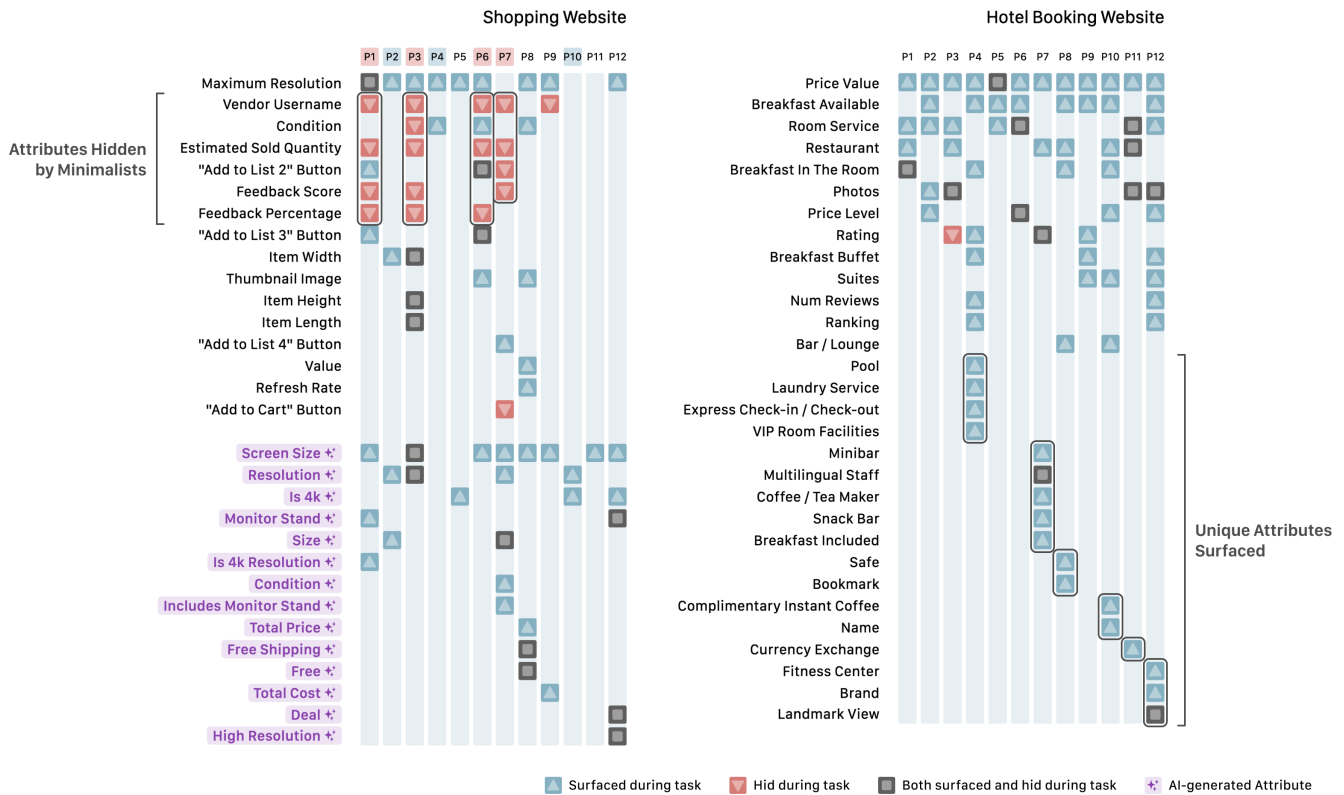


Figure 22: Visualization of all attributes that participants interacted with in the shopping and hotel booking website. The logs of the attributes that participants interacted with revealed common attributes Minimalists (7.2) hid from the overview in the shopping website and the unique attributes some participants surfaced to the overview in the hotel booking website. We highlight participants who only surfaced attributes blue and participants who often hid attributes red. Attributes are sorted by the number of participants that interacted with them. AI-generated attributes are sorted separately.

Hoarders and Minimalists. Prior research has identified tendencies when managing digital information along the spectrum of two extremes: hoarding and minimalism [71]. We found that this behavior also emerged when users were given the ability to surface or hide attributes in our probes (Fig. 22):

1. Hoarders (P2, P4, P10): Participants who did not hide any attributes from the overview.
2. Minimalists (P1, P3, P6, P7): Participants who deliberately hid attributes from the overview.

P2 (Hoarder) did not hide any attributes in the overview, including ones they surfaced manually from the detail view or by using AI. They reasoned, “hiding it has a bigger inertia for me to overcome because if I want it back, I would feel like I wasted more effort to remove it and bring it back again.” They further explained that this is their natural tendency, as they also clutter their screen with many tabs. On the contrary, P3 (Minimalist), who hid five default attributes from the shopping website and one default attribute on the hotel booking website, admitted to the higher effort, but also advocated for the reward, stating, “it probably takes longer to set it up the way I want it, but also then it’s easier to look through because I’m not getting irrelevant stuff.”

Hoarders may be more resilient to clutter in the overview but less resilient to repeatedly surfacing information from the detail view. For them, an interface that packs most information in the overview, allowing them to sift through it themselves, would be preferable. In contrast, Minimalists value the limited space in the overview and prefer to “keep the space clean.” For them, a blank slate in the overview, allowing them to choose which attributes to display, would be more preferable. For those who fall between these extremes, a setting to control a maximum number of attributes shown in the overview, automatically hiding less important ones when more are surfaced, may also be preferable.

Utilizing the overview composition and layout to structure tasks. Among the 10/12 participants that used multiple overviews in shopping probe, five of them (P1, P2, P6, P7, P12) utilized two overviews: one for each monitor type (work and home). For instance, P12 opened only one overview at a time when finding monitors of each type to offload important information for later, saying, “It gives me some mental space where I can preserve it.” They also appreciated how they could open both spaces at once when they needed to compare between the two overviews. P6 did the same, but additionally split their budget into each overview, allocating \$150 for the work monitor and \$100 for the home monitor.

The behaviors of these participants hint at a pattern of utilizing the overview composition and layout to manage sub-tasks. To support this usage pattern, future implementations of malleable overview-detail interfaces can provide support for nested or linked overviews, created either manually or with AI. This can allow users to break complex tasks into more manageable, hierarchical sub-tasks or linearly order overviews to queue sequentially dependent tasks. This aligns with literature that take the task-centric approach to provide hierarchically organized browser tabs [14], dedicated desktop spaces [24], and dynamically linked email threads [10].

7.3 RQ3. What additional customizations did users seek that were not supported by our design probes?

Choose how detail views are invoked (P6). Although we implemented our probes to open detail views by clicking, P6 suggested the probes to allow users to also open details by hovering. They further described how they would also want a “less maximized” view of the details, suggesting the desire to invoke multiple detail views with *different sets of attributes* in different ways.

Attributes about all items in the overview (P6, P10, P12). Participants wanted details across all items in the overview in addition to ones provided for each item. For instance, P6 wanted to calculate the total price of all their bookmarked hotels, while P10 wanted the “best price” based on multiple attributes they surfaced. Indeed, if the detail view can have attributes about an item, the overview can have attributes about the entire collection [73]. Future implementations of malleable overview-detail interfaces should explore how this can be supported.

8 Discussion and Future Work

Our exploration of design probes and observations from our user study reveal opportunities to further investigate malleable overview-detail interfaces. We explore future directions for supporting additional customization aspects, studying prolonged use, implementing them in real-world websites, and enabling user-defined abstractions in other contexts. Finally, we discuss how our design pattern for enhancing end-user customizability can extend beyond the overview-detail pattern.

8.1 Supporting More Variations for Malleable Overview-Detail Interfaces

In this research, we focused primarily on the utility of the malleability of overview-detail interfaces as a whole, rather than the utility of each individual variation. As a result, we did not implement all possible variations within the design space. However, implementing the full range of variations—and potentially more outside of our design space—could provide a more comprehensive understanding of the value each design choice offers. Further probes could also be implemented in contexts beyond shopping and hotel booking to offer insights into which variations and instances users might prefer under different scenarios and gain stronger evidence of various usage patterns. Future work should explore the full range of design variations, including ones beyond our design space, and investigate their long-term utility across a variety of contexts.

8.2 Bringing Malleable Overview-Detail Interfaces to the Real World

Our user study showed promising signs of the utility and value of malleable overview-detail interfaces in meeting diverse user needs. Our work aims to serve as a concrete guide for implementing our overview-detail interface customizations across many real websites and applications. However, the decision to implement them ultimately rests with the developers—from individual providers to large corporations. Increasing customizability requires additional development, which demands significant time, increases codebase complexity, and may even require re-architecting the application entirely. Therefore, the critical question becomes: *How can we practically bring malleable overview-detail interfaces to the real world?*

We have seen many systems take the approach of foregoing the development cost entirely by creating an end-user tool as a browser extension [33, 38, 77], and this can indeed be similarly done for overview-detail interfaces by scraping active websites to identify overviews, detail views, and their associated attributes. In-fact, Sifter [26]’s DOM detection algorithm even identifies a variation of overview-detail interfaces. Our design pattern approach could perhaps expand upon such detection algorithms to identify more variations. However, these “outside-in” approaches often face challenges related to inconsistency and usability issues, making it difficult to implement them effectively.

Another approach is to reduce the cost for developers with a developer toolkit, such as a web framework, that provides many variations and customizations out of the box. This framework can package together an architecture involving GraphQL [20] to dynamically fetch attributes from the back-end and a custom UI component library that provides sets of layouts and compositions to use as defaults and for end-users to customize. While a toolkit may be less viable for existing websites, we aim to significantly reduce the cost for developers building new ones.

8.3 User-Defined Abstractions for All Interfaces

Diverse forms of abstractions are implemented in various information systems and digital artifacts to assist users in understanding, operating with, and communicating with complex knowledge. For instance, scientific data, mechanical systems, and code are often abstracted into various visualizations and diagrams, which involve the use of ellipses, blocks and arrows, and graphs [19, 23, 70]. Recent advancements in AI have also enabled the development of generating different levels of abstractions of text [27, 65, 66] and different summarizations of academic papers [1]. We have also explored dynamic abstractions in the context of overview-detail interfaces, enabling end-users to modify a custom set of attributes.

Across these various systems and contexts, there is a clear need for end-users to define abstractions that suit their unique needs. However, the mechanisms that allow users to easily abstract information extend much beyond modifying data attributes supported in our malleable overview-detail interfaces. For example, to create a results table from a data analysis program, users must copy the output and manually enter it into a separate formatting application, placing data into the right cells. Similarly, to share calendar availabilities to a colleague, users typically send a bullet-point list of available dates and times instead of using a scheduling tool. In

contrast, users who can define their own abstractions within the interface can generate tables directly from program outputs without manual entry and abstract away private event details from their calendar and highlight open time slots by changing the color mappings of events. Enabling such interactions can further foster collaboration, allowing multiple colleagues to transform table and calendar interfaces to converge on a shared abstraction. We see an opportunity to explore how users can go beyond modifying data attributes to convivially creating and tailoring their own abstractions involving semantics of text and data through AI and visual encodings with UI style. Our goal is to study user-defined abstractions as a phenomenon across these diverse contexts, seeking ways to better support them through a unified interaction methodology.

8.4 Malleable Interfaces: Beyond the Overview-Detail Design Pattern

This research focuses on making overview-detail interfaces malleable. While the overview-detail pattern is widely implemented across many interface systems, it is not the only design pattern in use. Research has explored the challenges and opportunities for end-users to customize various interface elements, including dashboards [61], item collections in lists and spreadsheets [26, 41], and menu lists [52]. We have also investigated how different representations of the overview interface can be customized in relation to the overview-detail interface. We believe that exploring other types of design patterns and their design spaces can offer new opportunities for making more interfaces malleable. This approach of investigating the variations and nuances of each design pattern may help us move closer to developing a general malleable interface.

9 Conclusion

This research explored malleable overview-detail interfaces, ones that end-users can customize to address individual needs. We conducted a content analysis of existing overview-detail interfaces in the wild and constructed a design space with three dimensions of variation: content, composition, and layout. Based on these dimensions, we developed interaction techniques to support customization, enabling end-users to customize diverse layouts, compositions of views, manipulable attributes across views, and AI-driven interactions to reformat and generate new attributes. Our user study found that participants produced diverse customizations, surfacing key insights into usage patterns in how they customized overview-detail interfaces. We hope our work paves the way for further exploration of malleable overview-detail interfaces in new contexts and inspires the expansion of malleable interfaces to additional design patterns.

Acknowledgments

We would like to thank anonymous reviewers for their insightful reviews to help us improve the paper. This work was supported by NSF under grant IIS-2432644.

References

- [1] 2024. Elicit. <https://elicit.com/>. Accessed April 5, 2024.
- [2] Eytan Adar, David Karger, and Lynn Andrea Stein. 1999. Haystack: per-user information environments. In *Proceedings of the Eighth International Conference on Information and Knowledge Management* (Kansas City, Missouri, USA) (CIKM '99). Association for Computing Machinery, New York, NY, USA, 413–422. doi:10.1145/319950.323231
- [3] Inc. Airtable. 2024. Airtable. <https://www.airtable.com/>. Accessed December 3, 2024.
- [4] Bill Atkinson. 2024. HyperCard. <https://arstechnica.com/gadgets/2019/05/25-years-of-hypercard-the-missing-link-to-the-web/>. Accessed December 3, 2024.
- [5] Eric Z. Ayers and John T. Stasko. 1995. Using Graphic History in Browsing the World Wide Web. In *Proceedings of the Fourth International Conference on World Wide Web* (Boston, Massachusetts, USA) (WWW4). Association for Computing Machinery, New York, NY, USA, 451–459. doi:10.1145/3592626.3592663
- [6] Benjamin Bach, Natalie Kerracher, Kyle Wm Hall, Sheelagh Carpendale, Jessie Kennedy, and Nathalie Henry Riche. 2016. Telling stories about dynamic networks with graph comics. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 3670–3682.
- [7] Patrick Baudisch, Nathaniel Good, and Paul Stewart. 2001. Focus plus context screens: combining display technology with visualization techniques. In *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology* (Orlando, Florida) (UIST '01). Association for Computing Machinery, New York, NY, USA, 31–40. doi:10.1145/502348.502354
- [8] Benjamin B. Bederson. 2000. Fisheye menus. In *Proceedings of the 13th annual ACM symposium on User interface software and technology*. 217–225.
- [9] Benjamin B. Bederson and James D. Hollan. 1994. Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics. In *Proceedings of the 7th Annual ACM Symposium on User Interface Software and Technology* (Marina del Rey, California, USA) (UIST '94). Association for Computing Machinery, New York, NY, USA, 17–26. doi:10.1145/192426.192435
- [10] Victoria Bellotti, Nicolas Ducheneaut, Mark Howard, and Ian Smith. 2003. Taking email to task: the design and evaluation of a task management centered email tool. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 345–352.
- [11] Michael Bolin, Matthew Webber, Philip Rha, Tom Wilson, and Robert C. Miller. 2005. Automation and customization of rendered web pages. In *Proceedings of the 18th Annual ACM Symposium on User Interface Software and Technology* (Seattle, WA, USA) (UIST '05). Association for Computing Machinery, New York, NY, USA, 163–172. doi:10.1145/1095034.1095062
- [12] Stuart K. Card, Jock D. Mackinlay, and Ben Shneiderman (Eds.). 1999. *Readings in information visualization: using vision to think*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- [13] Joseph Chee Chang, Nathan Hahn, and Aniket Kittur. 2020. Mesh: Scaffolding Comparison Tables for Online Decision Making. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '20). Association for Computing Machinery, New York, NY, USA, 391–405. doi:10.1145/3379337.3415865
- [14] Joseph Chee Chang, Yongsung Kim, Victor Miller, Michael Xieyang Liu, Brad A Myers, and Aniket Kittur. 2021. Tabs.do: Task-Centric Browser Tab Management. In *The 34th Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '21). Association for Computing Machinery, New York, NY, USA, 663–676. doi:10.1145/3472749.3474777
- [15] Fanny Chevalier, Pierre Dragicevic, and Christophe Hurter. 2012. Histomages: fully synchronized views for image editing. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (Cambridge, Massachusetts, USA) (UIST '12). Association for Computing Machinery, New York, NY, USA, 281–286. doi:10.1145/2380116.2380152
- [16] Andy Cockburn, Amy Karlson, and Benjamin B. Bederson. 2009. A review of overview+detail, zooming, and focus+context interfaces. *ACM Comput. Surv.* 41, 1, Article 2 (jan 2009), 31 pages. doi:10.1145/1456650.1456652
- [17] Hai Dang, Karim Benharrak, Florian Lehmann, and Daniel Buschek. 2022. Beyond text generation: Supporting writers with continuous automatic text summaries. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–13.
- [18] Inc. FileMaker. 2024. FileMaker Pro. <https://www.claris.com/filemaker/>. Accessed December 6, 2024.
- [19] Morgan M. Fong, Seth Poulsen, and Geoffrey L. Herman. 2021. What's in a Linked List? A Phenomenographic Study of Data Structure Diagrams. *ASEE Annual Conference and Exposition, Conference Proceedings* (26 July 2021). Funding Information: Thanks to Chinny Emeka and Craig Zilles for their help in verifying and clarifying the code book. Thanks also to Max Fowler, Vidushi Ojha, Mohammed Hassan, Abdu Alawini, and Colleen Lewis for reviewing this draft before submission.; 2021 ASEE Virtual Annual Conference, ASEE 2021 ; Conference date: 26-07-2021 Through 29-07-2021.
- [20] GraphQL Foundation. 2024. GraphQL: A Data Query Language. <https://graphql.org/>. Accessed December 10, 2024.

- [21] Giuseppe Ghiani, Fabio Paternò, Lucio Davide Spano, and Giuliano Pintori. 2016. An environment for End-User Development of Web mashups. *International Journal of Human-Computer Studies* 87 (2016), 38–64. doi:10.1016/j.ijhcs.2015.10.008
- [22] Björn Hartmann, Leslie Wu, Kevin Collins, and Scott R. Klemmer. 2007. Programming by a sample: rapidly creating web applications with d.mix. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology* (Newport, Rhode Island, USA) (UIST '07). Association for Computing Machinery, New York, NY, USA, 241–250. doi:10.1145/1294211.1294254
- [23] Devamardeep Hayatpur, Brian Hempel, Kathy Chen, William Duan, Philip Guo, and Haijun Xia. 2024. Taking ASCII Drawings Seriously: How Programmers Diagram Code. In *Conference on Human Factors in Computing Systems (CHI)*. <https://doi.org/10.1145/3544548.3581390>
- [24] D Austin Henderson Jr and Stuart Card. 1986. Rooms: the use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. *ACM Transactions on Graphics (TOG)* 5, 3 (1986), 211–243.
- [25] Kasper Hornbæk and Erik Frøkjær. 2001. Reading of electronic documents: the usability of linear, fisheye, and overview+detail interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Seattle, Washington, USA) (CHI '01). Association for Computing Machinery, New York, NY, USA, 293–300. doi:10.1145/365024.365118
- [26] David F. Huynh, Robert C. Miller, and David R. Karger. 2006. Enabling web browsers to augment web sites' filtering and sorting functionalities. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology* (Montreux, Switzerland) (UIST '06). Association for Computing Machinery, New York, NY, USA, 125–134. doi:10.1145/1166253.1166274
- [27] Peiling Jiang, Jude Rayan, Steven P. Dow, and Haijun Xia. 2023. Graphologue: Exploring Large Language Model Responses with Interactive Diagrams. In *Symposium on User Interface Software and Technology (UIST)*. <https://doi.org/10.1145/3586183.3606737>
- [28] Hyeonsu Kang, Joseph Chee Chang, Yongsung Kim, and Aniket Kittur. 2022. 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* (Bend, OR, USA) (UIST '22). Association for Computing Machinery, New York, NY, USA, Article 94, 15 pages. doi:10.1145/3526113.3545660
- [29] Hyeonsu B Kang, Tongshuang Wu, Joseph Chee Chang, and Aniket Kittur. 2023. Synergi: A Mixed-Initiative System for Scholarly Synthesis and Sensemaking. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23). Association for Computing Machinery, New York, NY, USA, Article 43, 19 pages. doi:10.1145/3586183.3606759
- [30] David R. Karger. 2007. Haystack: Per-User Information Environments Based on Semistructured Data. In *Beyond the Desktop Metaphor: Designing Integrated Digital Work Environments*. The MIT Press. doi:10.7551/mitpress/1584.003.0006 arXiv:https://direct.mit.edu/book/chapter-pdf/2320273/9780262256483_cad.pdf
- [31] David R Karger, Karun Bakshi, David Huynh, Dennis Quan, and Vineet Sinha. 2005. Haystack: A General-Purpose Information Management Tool for End Users Based on Semistructured Data. In *Conference on Innovative Data Systems Research*. <https://api.semanticscholar.org/CorpusID:14697727>
- [32] Alan C. Kay. 1993. The early history of Smalltalk. *SIGPLAN Not.* 28, 3 (March 1993), 69–95. doi:10.1145/155360.155364
- [33] Tae Soo Kim, DaEun Choi, Yoonseo Choi, and Juho Kim. 2022. Stylette: Styling the Web with Natural Language. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 5, 17 pages. doi:10.1145/3491102.3501931
- [34] David Kirsh. 2010. Thinking with external representations. *AI & society* 25 (2010), 441–454.
- [35] Clemens N. Klokrose, James R. Eagan, Siemen Baader, Wendy Mackay, and Michel Beaudouin-Lafon. 2015. Webstrates: Shareable Dynamic Media. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology* (Charlotte, NC, USA) (UIST '15). Association for Computing Machinery, New York, NY, USA, 280–290. doi:10.1145/2807442.2807446
- [36] R. Kosara, S. Miksch, and H. Hauser. 2002. Focus+context taken literally. *IEEE Computer Graphics and Applications* 22, 1 (2002), 22–29. doi:10.1109/38.974515
- [37] Andrew Kuznetsov, Joseph Chee Chang, Nathan Hahn, Napol Rachatasumrit, Bradley Breneisen, Julina Coupland, and Aniket Kittur. 2022. Fuse: In-Situ Sensemaking Support in the Browser. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) (UIST '22). Association for Computing Machinery, New York, NY, USA, Article 34, 15 pages. doi:10.1145/3526113.3545693
- [38] Google Chrome Labs. 2024. Project VisBug. <https://github.com/GoogleChromeLabs/ProjectVisBug/>. Accessed August 29, 2024.
- [39] Clayton Lewis and Robert Mack. 1982. Learning to use a text processing system: Evidence from “thinking aloud” protocols. In *Proceedings of the 1982 Conference on Human Factors in Computing Systems* (Gaithersburg, Maryland, USA) (CHI '82). Association for Computing Machinery, New York, NY, USA, 387–392. doi:10.1145/800049.801817
- [40] James Lin, Jeffrey Wong, Jeffrey Nichols, Allen Cypher, and Tessa A. Lau. 2009. End-user programming of mashups with vegemite. In *Proceedings of the 14th International Conference on Intelligent User Interfaces* (Sanibel Island, Florida, USA) (IUI '09). Association for Computing Machinery, New York, NY, USA, 97–106. doi:10.1145/1502650.1502667
- [41] Geoffrey Litt, Daniel Jackson, Tyler Millis, and Jessica Quaye. 2020. End-user software customization by direct manipulation of tabular data. In *Proceedings of the 2020 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software* (Virtual, USA) (Onward! 2020). Association for Computing Machinery, New York, NY, USA, 18–33. doi:10.1145/3426428.3426914
- [42] Zhicheng Liu, John Thompson, Alan Wilson, Mira Dontcheva, James Delorey, Sam Grigg, Bernard Kerr, and John Stasko. 2018. Data Illustrator: Augmenting Vector Design Tools with Lazy Data Binding for Expressive Visualization Authoring. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3173574.3173697
- [43] Jock D. Mackinlay, George G. Robertson, and Stuart K. Card. 1991. The perspective wall: detail and context smoothly integrated. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New Orleans, Louisiana, USA) (CHI '91). Association for Computing Machinery, New York, NY, USA, 173–176. doi:10.1145/108844.108870
- [44] Microsoft. 2024. Microsoft Access. <https://www.microsoft.com/en-us/microsoft-365/access>. Accessed December 6, 2024.
- [45] Bryan Min, Matthew T Beaudouin-Lafon, Sangho Suh, and Haijun Xia. 2023. Demonstration of Masonview: Content-Driven Viewport Management. In *Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23 Adjunct). Association for Computing Machinery, New York, NY, USA, Article 60, 3 pages. doi:10.1145/3586182.3615827
- [46] Pedro Molina, Santiago Meliá, and Oscar Pastor. 2002. Just-UI: A User Interface Specification Model. 63–74. doi:10.1007/978-94-010-0421-3_5
- [47] Thanh-Diane Nguyen and Jean Vanderdonck. 2012. User interface master detail pattern on Android. In *Proceedings of the 4th ACM SIGCHI Symposium on Engineering Interactive Computing Systems* (Copenhagen, Denmark) (EICCS '12). Association for Computing Machinery, New York, NY, USA, 299–304. doi:10.1145/2305484.2305535
- [48] Donald A. Norman and Stephen W. Draper. 1986. *User Centered System Design; New Perspectives on Human-Computer Interaction*. L. Erlbaum Associates Inc., USA.
- [49] Inc. Notion Labs. 2024. Notion. <https://www.notion.so/>. Accessed December 3, 2024.
- [50] Srishti Palani, Aakanksha Naik, Doug Downey, Amy X. Zhang, Jonathan Bragg, and Joseph Chee Chang. 2023. Relatedly: Scaffolding Literature Reviews with Existing Related Work Sections. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 742, 20 pages. doi:10.1145/3544548.3580841
- [51] Srishti Palani, Yingyi Zhou, Sheldon Zhu, and Steven P. Dow. 2022. InterWeave: Presenting Search Suggestions in Context Scaffolds Information Search and Synthesis. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) (UIST '22). Association for Computing Machinery, New York, NY, USA, Article 93, 16 pages. doi:10.1145/3526113.3545696
- [52] Jungchul Park, Sung H. Han, Yong S. Park, and Youngeok Cho. 2007. Adaptable versus adaptive menus on the desktop: Performance and user satisfaction. *International Journal of Industrial Ergonomics* 37, 8 (2007), 675–684. doi:10.1016/j.ergon.2007.04.006
- [53] Oscar Pastor and Juan Carlos Molina. 2007. *Model-Driven Architecture in Practice: A Software Production Environment Based on Conceptual Modeling*. Springer-Verlag, Berlin, Heidelberg.
- [54] Peter Pirolli and Stuart Card. 1999. Information foraging. *Psychological review* 106, 4 (1999), 643.
- [55] Peter Pirolli and Stuart Card. 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis.
- [56] Antoine Ponsard and Joanna McGrenere. 2016. Anchored Customization: Anchoring Settings to the Application Interface to Afford Customization. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 4154–4165. doi:10.1145/2858036.2858129
- [57] Gonzalo Ramos, Napol Rachatasumrit, Jina Suh, Rachel Ng, and Christopher Meek. 2022. ForSense: Accelerating Online Research Through Sensemaking Integration and Machine Research Support. *ACM Trans. Interact. Intell. Syst.* 12, 4, Article 30 (nov 2022), 23 pages. doi:10.1145/3532853
- [58] Hugo Romat, Emmanuel Pietriga, Nathalie Henry-Riche, Ken Hinckley, and Caroline Appert. 2019. SpaceInk: Making Space for In-Context Annotations. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 871–882. doi:10.1145/3332165.3347934

- [59] Daniel M. Russell, Mark J. Stefik, Peter Piroli, and Stuart K. Card. 1993. The cost structure of sensemaking. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems* (Amsterdam, The Netherlands) (CHI '93). Association for Computing Machinery, New York, NY, USA, 269–276. doi:10.1145/169059.169209
- [60] Houssein Saidi, Marcos Serrano, and Emmanuel Dubois. 2016. Investigating the effects of splitting detailed views in Overview+Detail interfaces. *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services* (2016). <https://api.semanticscholar.org/CorpusID:6423727>
- [61] Alper Sarikaya, Michael Correll, Lyn Bartram, Melanie Tory, and Danyel Fisher. 2019. What Do We Talk About When We Talk About Dashboards? *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 682–692. doi:10.1109/TVCG.2018.2864903
- [62] Ahmed Seffah. 2015. Patterns of HCI Design and HCI Design of Patterns. In *Human-Computer Interaction Series*. <https://api.semanticscholar.org/CorpusID:2163669>
- [63] B. Shneiderman. 1996. The eyes have it: a task by data type taxonomy for information visualizations. In *Proceedings 1996 IEEE Symposium on Visual Languages*. 336–343. doi:10.1109/VL.1996.545307
- [64] Robert Spence and Mark Apperley. 1982. Data base navigation: an office environment for the professional. *Behaviour & Information Technology* 1, 1 (1982), 43–54. doi:10.1080/01449298208914435
- [65] Sangho Suh, Meng Chen, Bryan Min, Toby Jia-Jun Li, and Haijun Xia. 2024. Luminare: Structured Generation and Exploration of Design Space with Large Language Models for Human-AI Co-Creation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 644, 26 pages. doi:10.1145/3613904.3642400
- [66] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. 2023. Sensecape: Enabling Multilevel Exploration and Sensemaking with Large Language Models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23). Association for Computing Machinery, New York, NY, USA, Article 1, 18 pages. doi:10.1145/3586183.3606756
- [67] Kesler Tanner, Naomi Johnson, and James A. Landay. 2019. Poirot: A Web Inspector for Designers. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300758
- [68] Craig S. Tashman and W. Keith Edwards. 2011. LiquidText: a flexible, multitouch environment to support active reading. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (CHI '11). Association for Computing Machinery, New York, NY, USA, 3285–3294. doi:10.1145/1978942.1979430
- [69] Jenifer Tidwell. 1997. A pattern language for human-computer interface design. Available via *DIALOG* (1997).
- [70] Bret Victor. 2011. Up and Down the Ladder of Abstraction. <http://worrydream.com/LadderOfAbstraction/> Accessed September 5, 2024.
- [71] Francesco Vitale, Izabelle Janzen, and Joanna McGrenere. 2018. Hoarding and Minimalism: Tendencies in Digital Data Preservation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3173574.3174161
- [72] Haijun Xia, Bruno Araújo, Tovi Grossman, and Daniel J. Wigdor. 2016. Object-Oriented Drawing. In *Conference on Human Factors in Computing Systems* (CHI).
- [73] Haijun Xia, Bruno Araujo, and Daniel Wigdor. 2017. Collection Objects: Enabling Fluid Formation and Manipulation of Aggregate Selections. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 5592–5604. doi:10.1145/3025453.3025554
- [74] Haijun Xia, Ken Hinckley, Michel Pahud, Xiao Tu, and Bill Buxton. 2017. WritLarge: Ink Unleashed by Unified Scope, Action, & Zoom. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 3227–3240. doi:10.1145/3025453.3025664
- [75] Haijun Xia, Nathalie Henry Riche, Fanny Chevalier, Bruno Rodrigues De Araújo, and Daniel Wigdor. 2018. DataInk: Direct and Creative Data-Oriented Drawing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI 2018, Montreal, QC, Canada, April 21–26, 2018*. ACM, 223. doi:10.1145/3173574.3173797
- [76] Jiajie Zhang and Donald A Norman. 1994. Representations in distributed cognitive tasks. *Cognitive science* 18, 1 (1994), 87–122.
- [77] Xiong Zhang and Philip J. Guo. 2018. Fusion: Opportunistic Web Prototyping with UI Mashups. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (Berlin, Germany) (UIST '18). Association for Computing Machinery, New York, NY, USA, 951–962. doi:10.1145/3242587.3242632

A Appendix

A.1 List of 33 Categories from Semrush

- Advertising and Marketing
- Airlines
- Apparel and Fashion
- Automotive
- Banking
- Beauty and Cosmetics
- Computer Software and Development
- Computer and Video Games
- Distance Learning
- Education
- Entertainment
- Finance
- Food and Beverages
- Gambling
- Healthcare
- Hospitality
- Human Resources
- Information Technology
- Insurance
- Investment
- Market Research
- Music
- Newspapers
- Online Services
- Real Estate
- Restaurants
- Retail
- Science
- Sports
- Telecom
- Transportation and Logistics
- Travel and Tourism
- Wellness

A.2 Prompts Used for Design Probes

The prompts have been slightly revised to be more presentable in the paper (i.e., renamed variables, removed conditionals inside the string, removed `JSON.stringify(...)` syntax, re-formatted syntax).

A.2.1 Finding an Existing Attribute or Generating a New One.

Your job is to identify the closest attribute that exists in the provided item data given the user prompt. This name should be an attribute of that item. If it doesn't exist, generate a new attribute. Do not respond with any other information.

Your response should be very short---you're just returning a single value. Wrap your response with brackets like so: [Answer].

The attribute name should be capitalized with normal spaces if the name includes more than one word.

Context for this item: `${itemAttributes}`

User prompt: Is it 4k?
Your response: [Is 4k]

User prompt: Multiply the item height with the length
Your response: [Item Height x Length]

User prompt: At least how old is this monitor?
Your response: [Approximate Minimum Age]

User prompt: What's the value of this monitor?
Your response: [Value]

User prompt: `${userPrompt}`
Your response:

A.2.2 Generating the Value of the Attribute Name.

Your job is to produce a value of a Attribute given the attribute name and context about the item.

Your response should be very short---you're just returning a single value. Wrap your response with brackets like so: [Answer].

If you cannot find any relevant attributes given the context provided, set the value to Not Specified. But for cases where the prompt is asking for existence, answer Yes or No instead.

Context for this item: `${itemAttributes}`

Examples:

Attribute Name: Is 4k
Attribute Prompt: Is it 4k?
Your response: [False]

Attribute Name: Item Height x Length
Attribute Prompt: Multiply the item height with the length
Your response: [35.6]

Attribute Name: Approximate Minimum Age
Attribute Prompt: At least how old is this monitor?
Your response: [At least 2 Years]

Attribute Name: Value
Attribute Prompt: What's the value of this monitor?
Your response: [Medium Value]

Attribute Name: \${attribute.name}
Attribute Prompt: \${attribute.prompt}
Your response:

A.2.3 Transforming Attribute Values.

Your job is to transform the value of an attribute of an item.

Your response should be very short---you're just returning a single value. Wrap your response with brackets like so: [Answer].

Context for this specific item: \${itemAttributes}

Example:

Attribute to transform: \${exampleAttribute.name}
Original attribute value: \${exampleAttribute.value}
Transform it in this way: \${userPrompt}
Response: \${exampleAttribute.value}

Attribute to transform: \${attribute.name}
Original attribute value: \${attribute.value}
Transform it in this way: \${userPrompt}
Response:

A.3 Design Probe Screenshots

Shopping Interface




Figure 23: Screenshot of the shopping website in a side-by-side layout

Shopping Interface




Figure 24: Screenshot of the shopping website with two overviews, one list view and one grid view.

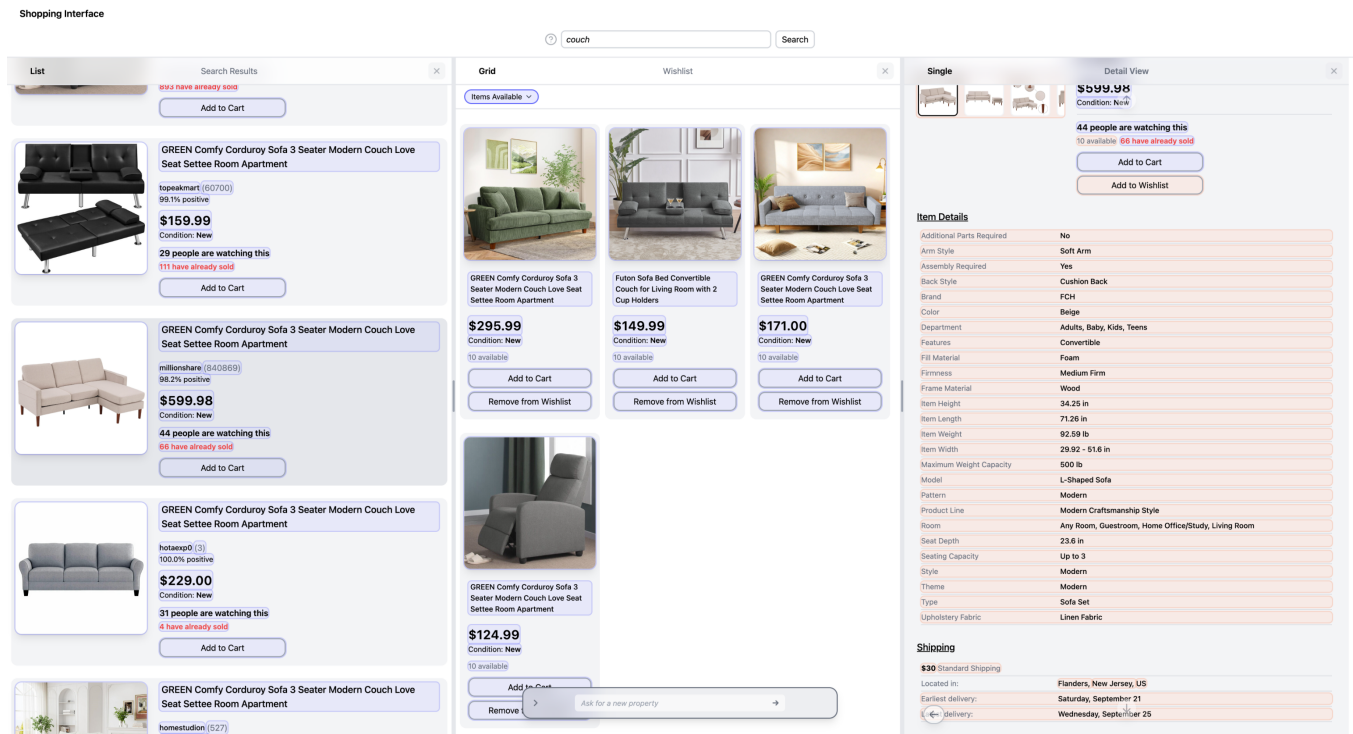


Figure 25: Screenshot of the shopping website with “Attributes Mode” activated. Upon activating “Attributes Mode”, the toolbar provides a text box to prompt AI.

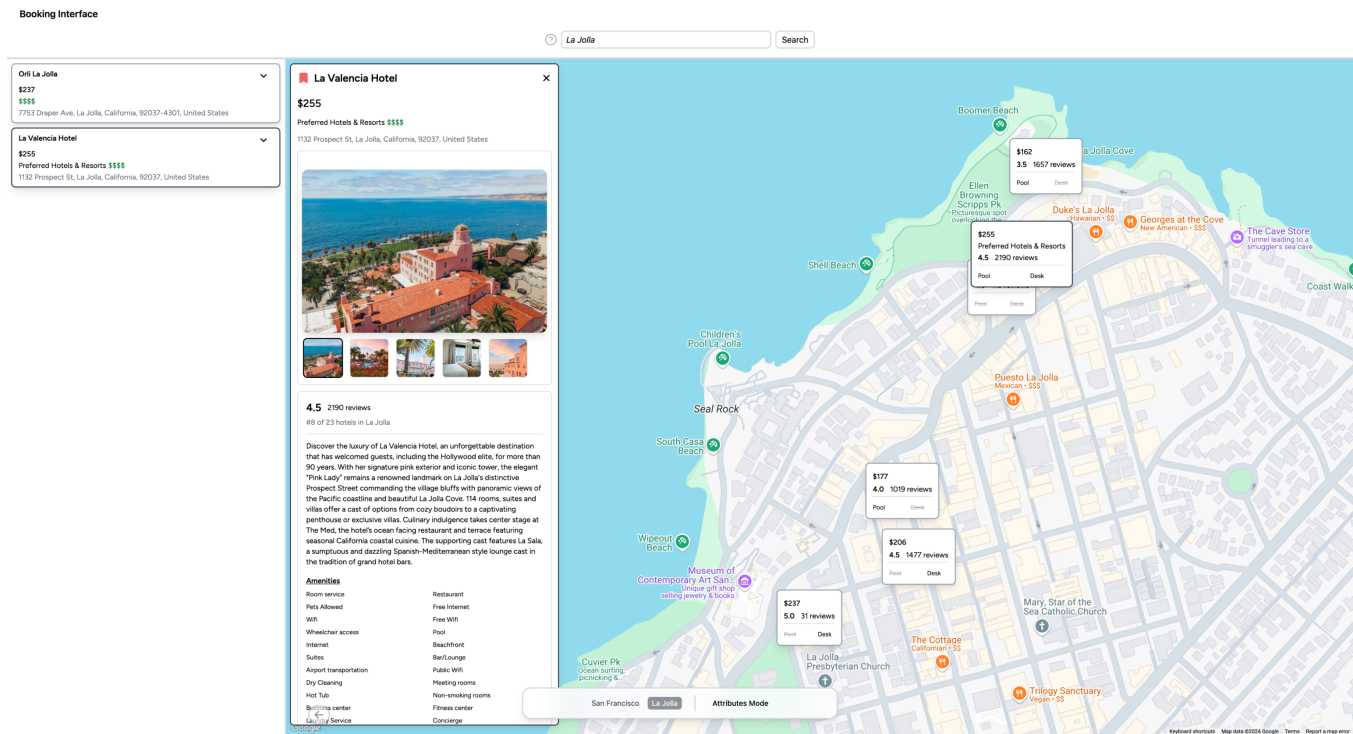


Figure 26: Screenshot of the booking website with different attributes surfaced in the bookmarks list and map view.