

A Design Space for Surfacing Content Recommendations in Visual Analytic Platforms

Zhilan Zhou, Wenyuan Wang, Mengtian Guo, Yue Wang, and David Gotz

Abstract—Recommendation algorithms have been leveraged in various ways within visualization systems to assist users as they perform of a range of information tasks. One common focus for these techniques has been the recommendation of content, rather than visual form, as a means to assist users in the identification of information that is relevant to their task context. A wide variety of techniques have been proposed to address this general problem, with a range of design choices in how these solutions surface relevant information to users. This paper reviews the state-of-the-art in how visualization systems surface recommended content to users during users' visual analysis; introduces a four-dimensional design space for visual content recommendation based on a characterization of prior work; and discusses key observations regarding common patterns and future research opportunities.

Index Terms—Adaptive Visualization, Recommendation, Literature Survey, Design Space

1 INTRODUCTION

Visual analytics as a discipline was born from the recognition that the process of deriving meaningful insights from large volume, high dimensional, heterogeneous, and often conflicting sources of data is fraught with an array of challenges [70]. Much of the research in this area is motivated by complex information problems in which users must work to discover relevant data elements, to view those data from different perspectives, and to progressively connect disparate pieces of information to build an understanding of their analytical target as part of an ongoing sensemaking process [55].

Reflecting on the scale of the challenge in terms of both volume and complexity of information, visual analytics techniques generally aim to combine computational techniques with human analysis activities in a cooperative fashion to exploit the complementary strengths of both computer and human [61]. Moreover, the proliferation of machine learning has led to a wide variety of proposals for how to leverage mixed-initiative and data-driven algorithms during visual analysis [26].

Perhaps not surprisingly, researchers have experimented with many ways in which the combination of algorithmic and human approaches could best be leveraged. This includes, for example, “user-in-the-loop” approaches in which users can provide various forms of guidance to the underlying algorithms (e.g., [12, 69]). It also includes “analytics-in-the-loop” approaches in which model-driven algorithms are designed to provide various forms of assistance to human analysts.

This latter category—using computational approaches to help a human user make analytical progress—can itself come in many forms. One frequent approach is to algorithmically recommend specific *visualizations* (e.g., [33, 49]). In this formulation, the algorithmic problem is to determine the “best” view of a given set of data in a particular task context. Alternatively, and more relevant to the work presented

in this paper, computational approaches have also been proposed to recommend specific *content* to users. In this problem, the goal is to algorithmically help users more effectively identify new information that is relevant to their analytic context.

Visual analytics systems that support this type of content recommendation must include several different components. First, the system must maintain some representation of a user's current analytic context. Second, the system must include some form of recommender which, given the representation of a user's context, evaluates new information to identify which content should be recommended. Finally, recommended content must be surfaced to the user through the user interface of the visual analytics system.

It is this last element—the way in which recommended content is surfaced to users during visual analysis—that is the focus of this paper. By design, content recommendation algorithms are used to identify new content relevant to the user's analytic context which they might not identify on their own. This makes content recommendation a potentially powerful tool in support of the sensemaking process. However, the way in which recommended content is incorporated into a user's analytic workflow can influence the utility of the recommended content itself. The utility is heavily influenced by the design choices of how, when, and where that content is presented to the user.

To better understand the design space for the content recommendation within visual analytics platforms, and motivated by our own work on this topic [83], we conducted a formal survey of the visual analytics literature to identify prior efforts at content recommendation. We searched the literature for adaptive visualization systems that algorithmically identify and surface new content with the aim of helping users discover additional relevant information given their analytical context.

Based on an analysis of the identified articles, a four-dimensional design space is proposed that captures the primary design choices for communicating recommended content during visual analysis. The papers found in the search are then characterized using the proposed design space to identify common patterns, rare design choices, and potential opportunities for future research.

In this way, the research contributions of this paper include: (1) a systematic review of the literature to identify the state-of-the-art in how visualization systems surface recommended content during visual analysis; (2) a design space for visual content recommendation interfaces; and (3) identification of both common approaches and opportunities for future research.

The remainder of this paper is organized as follows. A brief review of related work is provided in Section 2 to better frame the scope of our literature survey. Section 3 then introduces the methodology behind the survey including the search strategy. The design space is presented in Section 4 and used to characterize the papers identified in the literature

- Zhilan Zhou is with the Department of Computer Science at the University of North Carolina at Chapel Hill. E-mail: zzl@email.unc.edu.
- Wenyuan Wang is with the School of Information and Library Science at the University of North Carolina at Chapel Hill. E-mail: vaapad@live.unc.edu.
- Mengtian Guo is with the School of Information and Library Science at the University of North Carolina at Chapel Hill. E-mail: mtguo@email.unc.edu.
- Yue Wang is with the School of Information and Library Science at the University of North Carolina at Chapel Hill. E-mail: wangyue@unc.edu.
- David Gotz is with the School of Information and Library Science and the Department of Computer Science at the University of North Carolina at Chapel Hill. E-mail: gotz@unc.edu.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

search. Finally, before concluding, a discussion in Section 5 identifies key patterns, limitations, and opportunities for future research.

2 RELATED WORK

The literature survey and design space presented in this paper focus on the ways in which content recommendations are surfaced to users during visual analysis. This topic is closely related to two broader areas of research that have received significant attention in recent years. These include both (a) visualization recommendation systems, and (b) more wide-ranging research on recommender systems within the human-computer interaction (HCI) community.

2.1 Visualization Recommendation

Algorithmic solutions that recommend or automatically construct visual representations of data have been studied for several decades (e.g., [57]). Broadly speaking, these types of visualization recommendation systems will automatically configure and suggest one or more visual representations for a given set of data based on either data properties, a user's task context (including, potentially, user preferences), or both. Given these inputs, a visualization recommendation system will typically aim to determine the best visualizations to use with the given dataset in order to help the user achieve some analytical goal [73].

A variety of techniques have been proposed to accomplish this type of visualization recommendation, ranging from rule or pattern-based approaches that infer visualization requirements from user activity (e.g., [30]) to machine learning models trained on large numbers of example visualizations (e.g., [33, 49]).

Despite the variety in both the input and the computational approaches through which visualization recommendations are generated, these systems share a common goal: given an input dataset and perhaps some additional information about the user's task or preferences, recommend one or more views of the input data. These approaches suggest *forms* or *styles* of visualization, but do *not* aim to inject any new *content* into a user's analysis [8]. In sharp contrast, the focus of this paper is on the content recommendation during visualization rather than visualization recommendation as we have just defined it. That is, we focus on papers that (a) describe recommendations of new content that should be brought to the user's attention, and (b) how those recommendations are surfaced to the user. For this reason, visualization recommendations papers (such as those cited in this subsection) are considered outside of the scope for this paper.

2.2 Recommendation Systems in HCI

The HCI research community has an extensive record of impactful research exploring different types of recommendation systems in application scenarios that extend well beyond visualization. Fundamentally, the term *recommendation system* (or *recommender system*) is used to describe any system that makes content suggestions to a user based upon a computed measure of relevance between said content and user's interest [54].

Usually, recommender technology is provided in combination with other information access techniques, such as query-based search [22]. For example, recommendation systems are used widely on websites such as search engines or shopping websites to recommend related documents to read (e.g., [11]) or related products to purchase (e.g., [66]), respectively. They are also widely used on online social media platforms to recommend videos, posts, and other content based on both user's current interactive activity and models of users' historical preference (e.g., [21]).

Many different approaches to solving the content recommendation problem have been proposed including content-based filtering methods, collaborative filtering methods, and hybrid methods. Research on these topics has often focused on optimizing quantitative measures of recommendation accuracy [63]. However, researchers have also recognized the critical importance of user experience and interaction design [41]. This has, for example, led to suggested "design patterns" for user interfaces within the context of recommender systems [22]. The contributions of this paper are similarly focused on content recommendation and the ways in which these recommendations are surfaced

to users. However, unlike past work, this paper aims to specifically study ways in which content recommendations are surfaced within visual analytics platforms.

3 METHODOLOGY

The design space proposed in this paper is derived from an analysis of articles identified through a formal literature survey. This section provides an overview of the survey strategy and the analysis process used to inform the design space. The numbers of articles found at each stage of the search process are also reported.

3.1 Search and Screening

The literature search performed for this paper followed a search strategy and screening process that adheres to the recommended procedures specified in the PRISMA standard (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) [47]. PRISMA defines a multi-stage procedure that has been widely adopted in health and medical fields as a "best practice" for systematic literature reviews. PRISMA also defines a standardized flow diagram to depict the various stages and the number of articles considered, included, and/or excluded at each step.

The PRISMA flow diagram for the literature survey in this paper is shown in Figure 1. As shown in the diagram, articles were identified from two different sources: IEEE Xplore [34] and the ACM Digital Library [1]. A full-text keyword search was performed within both online digital libraries to identify candidate articles. Given the focus of the search, "adaptive visualization(s)" was the phrase used as search terms. The keywords were free-text and not selected from a controlled vocabulary. Searches were conducted in both IEEE Xplore digital library and ACM digital library in December 2021.

This choice of keywords, selected to reflect the focus of this survey on methods of visually surfacing recommended content within visualization-based interfaces, was made after experimentation with various alternatives in both the IEEE and ACM digital libraries, as well as informal exploratory searches made through Google Scholar [29]. Alternatives explored included, for example, "visualization recommendation." However, this phrasing is widely used to describe systems that recommend the form of visualization rather than the content (see Section 2.1). As a result, it returned articles that were out of scope and missed many important papers that deserved to be included. Alternatively, a search for "content recommendation" was excessively broad due to the generic terms used. This search therefore returned an enormous number of articles, many of which were not related to content recommendation at all as defined in this paper (much fewer content recommendations in the context of visualization).

The choice of scoping the search to IEEE and ACM digital libraries reflects the central role these two repositories have in archiving peer-reviewed high-quality articles relevant to visualization and recommendation research. Together, these two resources index a broad range of technical journals and conference proceedings, and they include a number of the leading publications in the disciplines relevant to the scope of this search.

In total, 175 articles were identified via the specified keywords. A screening process was followed to remove duplicates and exclude irrelevant articles by examining titles and abstracts. This process resulted in a total of 92 articles for potential inclusion in the survey.

The choices of keywords and digital libraries used in the search clearly have a significant impact on the set of articles. The impacts of these choices are discussed in more detail along with other limitations of our approach in Section 5.3.

3.2 Eligibility Assessment

The 92 articles that emerged from the screening process were inspected more closely for eligibility via a full-text assessment. This step ensured that papers which appeared relevant from the title and abstract screening were indeed appropriate for inclusion based on the research topic. This further narrowed the number of candidate articles to 46.

Once articles were determined to be relevant in terms of topic, two additional eligibility constraints were applied which excluded 23 additional articles. First, papers were required to present a working prototype to ensure that our survey focused on designs that were used in actual implementations from prior work. This requirement eliminated 12 articles that proposed theoretical concepts or computational algorithms [6, 7, 16, 23, 24, 28, 31, 36, 51, 56, 59, 62], and 8 papers that included results from user studies without prototypes of interactive visualizations [9, 10, 19, 40, 53, 68, 71, 79]. Second, because content recommendation is a dynamic process that occurs while a user interacts with a system, papers that met the prototype requirement were also required to describe user interaction. This requirement resulted in the elimination of 3 additional papers [18, 48, 81].

Finally, we removed 6 papers [2, 20, 60, 76, 78, 82] where the prototypes were deemed not relevant to the scope of the study, were described with insufficient detail to include, or were identified as duplicates that described the same prototypes as other papers in our collection. These steps reduced the included set to 17 papers.

3.3 Additional Search Following Citation Links

Recognizing the limitations of keyword search, the research team enriched the collection of relevant documents by examining all citations from the 17 papers which emerged from the eligibility determination process. Every document cited in the references section for these 17 documents went through the same screening and eligibility steps as the original articles found in the IEEE and ACM digital libraries. This process resulted in 9 additional articles being included for a total of 26 articles in the final collection.

3.4 Reference Management

The open-source reference management software Zotero was leveraged to organize all the references identified in the search and screening process. In addition, spreadsheets were maintained to keep track of the overall review process, decisions, and statistics, as reported in Figure 1.

3.5 Meta-Analysis Process

The final set of 26 publications that were included in the survey was then analyzed by the authors of this paper to extract details regarding the ways in which the described prototypes surfaced recommended content to users. Descriptors of various techniques used in prototype systems in the literature were recorded in a spreadsheet following an open coding process. The descriptors were iteratively refined and grouped to identify the key aspects that could be used to characterize various design choices. New values were added as needed to ensure that all identified prototypes could be accurately described. As the descriptors were refined, previously characterized papers were revisited to ensure they were labeled using the latest set of descriptors. At the end of this process, the descriptors eventually consolidated into the four-dimensional design space described in Section 4.

As the final design space emerged, two authors (ZZ and WW) revisited all articles included in the search and formally characterized each article using terms from the design space. The two authors worked independently and then resolved disagreements through discussion and resolution. In rare circumstances, disagreements led to minor revisions of how dimensions in the design space were described. When this occurred, previously characterized papers were revisited to ensure they were consistently analyzed. The final results of this process are illustrated in Figure 3 and discussed in detail throughout the remainder of this paper.

4 THE DESIGN SPACE

This section introduces a four-dimensional design space which provides a framework for describing how content recommendation capabilities are made visible to users within visualization systems. As described in Section 3, the design space was derived from an analysis of relevant papers identified through a systematic review of the research literature. This section first introduces the dimensions of this design space, and then uses the design space to categorize and discuss prior research identified in the literature review.

4.1 Dimensions of the Design Space

Informed by our analysis of the publications identified in our literature search, we define four key dimensions that help characterize the interface design for content recommendation within visual analytic systems: **Directness**, **Forcefulness**, **Stability**, and **Granularity**. This subsection formally defines each of these dimensions.

We note that we explicitly consider as out of scope issues related to: (a) the identification of what content should be recommended (e.g., similarity measures, individual vs. collaborative approaches, etc.); and (b) how to model user interests and behaviors as the basis for computing recommendations. These out-of-scope aspects are essential parts of recommender systems even when visualization is not employed [35], and they are key aspects to much of the related work referenced in Section 2.2. However, they are beyond the scope of the design space proposed in this paper because they relate to what content is recommended. In contrast, the proposed design space is focused on the methods via which the recommended content is made visible to users.

To help readers understand how these dimensions correspond to practical design decisions, we provide concrete examples using a hypothetical visual analytic platform, *HVAP*. This hypothetical system has a relatively simple interface with just two tabs: the first tab contains visualizations in the form of a scatterplot and a histogram; the second tab contains a text-based list. This is illustrated in Figure 2. *HVAP* is used as an example for explanatory purposes throughout the remainder of this section.

4.1.1 Directness

The *directness* dimension captures the relationship between where a user is interacting with a visualization and the location at which recommended content is surfaced. The directness dimension takes a binary value, which can be either *Direct* or *Indirect*.

A *direct* design places recommended content within the same interface component with which the user is interacting. For example, a direct recommendation for a user interacting with the scatterplot in *HVAP* would appear within the same scatterplot.

In contrast, content surfaced via an *indirect* recommendation would be visible through another interface component. For example, for a user interacting with the *HVAP* scatterplot, indirect recommendations might be surfaced via the histogram (in the same tab but different component) or via the text list (in a different tab).

4.1.2 Forcefulness

The *forcefulness* dimension represents how intrusively the recommendation is surfaced with respect to the user's workflow. Unlike directness, forcefulness exists on a scale with different design choices exhibiting levels of forcefulness. Based on the analysis of the articles identified in our literature, forcefulness is assessed on a five-point scale with 1 being the least forceful and 5 being the most forceful.

In discussing forcefulness, we use the term *context* to refer to the information being displayed within a visual analytics system before any recommended content is surfaced. When the context is displaced completely by a recommendation, the user is compelled to attend to the recommendation before returning to their prior analytical workflow. For example, if a user is interacting with the scatterplot in *HVAP* and a recommendation is displayed such that it obscures the visualizations entirely, this would forcefully interrupt a user's analysis. We rank this type of design as Level 5 on the forcefulness scale.

In contrast, a fully passive design for the recommendation will not alter the context in any way and will only surface information in out-of-view locations for users to optionally attend to at some time in the future. For example, recommended content for the same *HVAP* user interacting with the scatterplot could be displayed within the text list. We rank such a design as Level 1.

Table 1 summarizes the full forcefulness scale using the Level 5 and Level 1 scenarios as the extremes of the scale. As the table shows, two factors contribute to the forcefulness of a recommendation: the degree to which context is displaced, and the visibility of the recommended content.

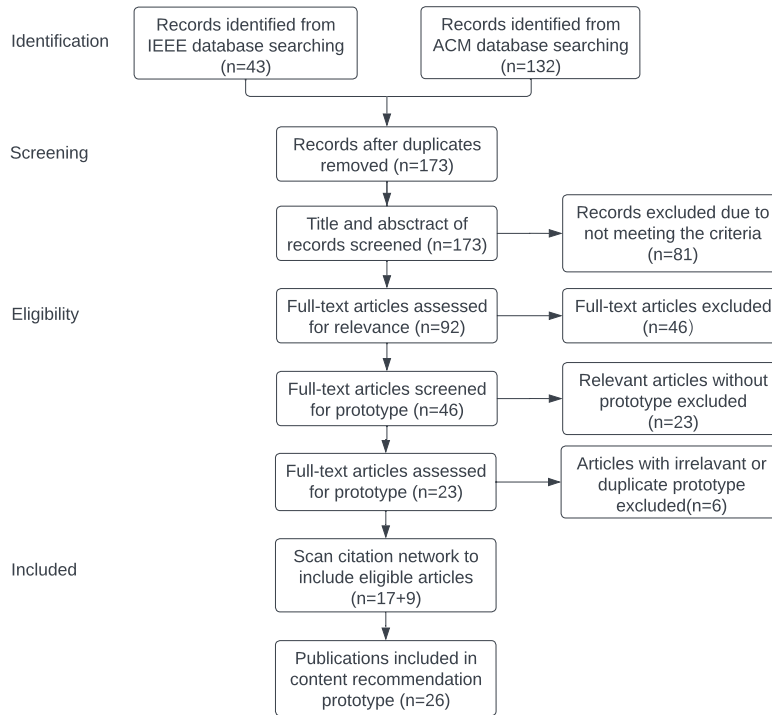


Fig. 1. The PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) flow diagram for our searching and screening process.




Fig. 2. A sketch of HVAP, the Hypothetical Visual Analytics Platform used as an example when defining the design space proposed in this paper. The hypothetical interface includes two tabs: one with visualizations and another with a text-based list.

Example recommendation scenarios within HVAP for each level of forcefulness are provided in Table 2. The examples all assume users are interacting with the HVAP scatterplot at the time of recommendation.

We emphasize that while directness and forcefulness can seem similar, they are conceptually different dimensions. Using HVAP as an example, if a user selects a point in the scatterplot, recommended content can be shown through a tooltip that appears near the hovered point (direct, forcefulness = 4), or a highlight of the five closest nodes and a dimming of others (direct, forcefulness = 3). Alternatively, recommended content can be surfaced in the bar chart component, which is outside of the user’s current focus in the scatterplot (making this alternative indirect). The interface can highlight relevant bars with respect to the selected point (indirect, forcefulness = 3), or replace the original bar chart with a new one (indirect, forcefulness = 5).

4.1.3 Stability

The *stability* dimension describes the temporal dynamics of when the content being recommended is updated within the interface. In this dimension, we only consider updates to the visible portion of a recommendation system. The underlying algorithm that prepares the

Forcefulness Level	Visible Context	Visible Recommendation
5 (Most Forceful)	None	Recommendation only
4	Part	Highlighted over context
3	All	In context, highlighted
2	All	Parallel to the context
1 (Most Passive)	All	None (Hidden)

Table 1. The forcefulness scale has five levels and depends on two factors: the degree to which the context is displaced, and the visibility of the recommended content.

content may update its output on its own schedule that is independent of when the content in the interface is updated, and this algorithmic timing is not the focus of this design space. Design choices within the stability dimension can be classified into three general categories: periodic, event-driven, and on-demand policy.

Periodic updates occur at a fixed time interval. For example, the display of recommended content could be revised every 3 seconds. *Event-driven* updates occur in response to some underlying system event. For example, the display of recommended content could be revised every time a relevant document appears in a data stream, or in response to a resizing of a user interface window. Event-driven designs often result in interfaces that exhibit irregularly-timed changes to the display of recommended content. Both periodic and event-driven designs result in interface updates that are system initiated. These designs result in less-stable interfaces which can change the visual display at times that are not expected and/or inconvenient to a user.

In contrast, *on-demand* systems provide the most stable experience because they only update the display of recommended content in response to explicit user requests (e.g., clicking on a button to update recommendations). This design can result in fewer updates and longer delays between the algorithmic identification of relevant content, but offers more agency to users and reduces the chances that updates will disrupt a user’s workflow.

Forcefulness Level	Elements Visible to User	Recommendation Location
5 (Most Forceful)	Only the recommendation	The visualization tab, obscuring all context
4	Just the scatterplot and the recommendation	Visualization tab, obscuring some context
3	Scatterplot, histogram, and the recommendation	Visualization tab, highlighted within the scatterplot
2	Scatterplot, histogram, and the recommendation	Visualization tab, shown in parallel to context
1 (Most Passive)	Scatterplot and histogram; recommendation not visible	List tab (not visible without additional user interaction)

Table 2. Examples of different levels of content recommendation forcefulness for a user interacting with the scatterplot in HVAP (see Figure 2). The level of forcefulness depends on both (a) what information is visible to users, and (b) the location of the recommendation.

4.1.4 Granularity

The *granularity* dimension describes the unit of content that is recommended to the user. By “unit of content” we refer to the atomic unit of information that is being recommended or visualized (e.g. a document; a semantic concept; a dimension of a dataset; a recommended database entity such as the data for a single house in a real estate application). To characterize this aspect, we focus on the difference between (a) the unit of content that a user is viewing within the context of the recommendation, and (b) the unit of content being recommended. The granularity of recommended content can be coarser, similar, or finer compared to the context.

In practice, recommendations often contain heterogeneous content with mixed granularities. For example, imagine that HVAP is used to analyze food pricing strategies in various countries. In this scenario, imagine the scatterplot is showing the relationship between the wholesale and market prices of various food categories (e.g., fruit, dairy, grains) for a specific country. If the HVAP system recommends information about the prices of oranges, we would consider this a finer granularity recommendation because an orange is a narrower concept than the fruit category. If the system recommends information about the prices of citrus fruits in another country, we would consider this a similar granularity recommendation. If the system instead recommends market research data for both individual fruits as well as the overall fruit category, it would be a mixed granularity recommendation (finer and similar).

4.2 Mapping Literature in the Design Space

Given the design dimensions defined above, we used the design space to characterize each of the 26 papers identified in the literature review. The findings are reported by dimension, with the full results summarized in the table in Figure 3. The results show significant variation across all dimensions, but also some interesting patterns that speak to design trade-offs for visualization interfaces that adopt content recommendation capabilities.

4.2.1 Directness

The *directness* dimension describes the location at which a recommendation is displayed. Designs that surface recommendations within the interface component that is being interacted with by a user adopt a *direct* design, while those that show recommendations in other interface components are considered *indirect*. Often, systems provide recommendations in more than one way. For example, in HVAP, a user interacting with the scatterplot might see recommendations both within the scatterplot (direct) and text list (indirect). This can result in designs that exhibit a combination of both direct and indirect recommendations within different parts on the platform.

Direct Designs. A total of ten papers [4,5,13,27,38,44,45,64,65,80] in our survey incorporate at least one form of direct recommendation. These designs highlight the recommendation content within the interface component that is the focus of user interaction.

A concrete example of this approach can be found in the work of Bae et al. [5]. They developed an extensible multi-application platform that included a document visualization as the primary interface to support document triage. While users examine the visualization to discover details about various documents, the platform will highlight potentially

relevant documents within the visualization to direct users’ attention to partial or whole documents that best match users’ interests as inferred by the system. Because the highlighted recommendations appear within the same visualization that the users are exploring, this is an example of direct recommendation design.

Indirect Designs. A total of 21 papers [3,4,13–15,25,27,32,37,39,42–44,46,50,52,58,72,74,77,80] identified in our review of the literature demonstrated at least one form of indirect recommendation. These systems included design elements which surfaced recommended content through interface components that were outside the user’s current focus of interaction.

For example, Kim et al. [39] describe a system that recommends re-expressions of spatial measurements (e.g. 78 miles) from texts in a way that contextualizes the distances in terms that are more familiar to a user (e.g. 8 times the distance between your home and your company). A user’s primary focus of interaction in this system is the text itself, but the personalized spatial analogies are shown visually using interactive maps placed beside the text. This placement of the recommendation in a visualization beside and distinct from the main area of interaction makes this an example of indirect recommendation.

Most often, indirect recommendations are located within an interface component that is visible and alongside the user’s primary area of interaction. In some rare cases, however, indirect recommendations have been placed on a separate page that requires additional interactions to be revealed. For example, Brusilovsky et al. [14] presented a system providing personalized recommendations for a structured repository of educational examples. Users would typically interact with the examples and browse the repository list on the primary component of the interface. The system then included a secondary display area which visualized a 2D layout of recommended concepts (selected from the pre-defined repository). Users need to perform an additional click to view these recommendations.

Similarly, the educational system proposed by Ahn & Brusilovsky [3] uses a secondary interface to share recommendations, which requires an additional click to access. The system described in that paper adopts a primary visualization interface that uses a Self Organizing Map (SOM) to organize educational resources by topic. Users can click to access a list of recommended topics that are displayed on a secondary interface.

Discussion. Overall, about two-fifths of the papers in our survey adopt a direct design. In contrast, more than four-fifths of the papers adopt an indirect design. This makes indirect designs approximately twice as common. Moreover, half of the direct-design systems also provide indirect recommendations.

Direct designs make recommendations more visible by placing them squarely within a user’s area of attention. This approach, therefore, may give recommendations more visibility even for focused users immersed in complex cognitive tasks. However, this same attribute of direct designs has the potential to be more disruptive to a user’s analytical work. In addition, because direct approaches must be designed to visually integrate within an existing view, there are more constraints on how the recommendations are visually represented.

Indirect designs, because they appear outside of the user’s primary area of focus, may at times be less visible than their direct counterparts. However, such designs usually have ample space and freedom of design for how recommendations are displayed. Moreover, a more forceful

	Mutlu et al. 2016	Roussinov 1999	Wiza et al. 2005	Knittel et al. 2022	Bonada et al. 2016	Karsai et al. 2016	Lei Shi et al. 2009	Silva et al. 2016	Glowacka et al. 2013	Ahn et al. 2015	d'Entremont and Storey 2006	Bae et al. 2010	Silva et al. 2018	Yelizarov and Gamayunov 2014	Lehmann et al. 2010	Brusilovsky and Loboda 2006	Kim et al. 2016	Leuski and Allan 2004	Nazemi et al. 2013	Gobel et al. 2018	Kanai and Hakozaki 2000	Lalle et al. 2021	van Tonder and Wesson 2008	Ahn and Brusilovsky 2011	Brusilovsky et al. 2006	Vartak et al. 2014	Count #	
Directness -- Direct																											10	
Directness -- Indirect																												21
Forcefulness -- 5																												5
Forcefulness -- 4																												6
Forcefulness -- 3																												18
Forcefulness -- 2																												1
Forcefulness -- 1																												2
Stability -- On-Demand																												13
Stability -- Event-driven																												10
Stability -- Periodic																												3
Granularity -- Coarser																												7
Granularity -- Similar																												20
Granularity -- Finer																												11

Fig. 3. A systematic review of the literature found a total of 26 papers as illustrated in Figure 1. Each paper was reviewed by two authors and classified within the proposed four-dimensional design space for the content recommendation within visual analytics platforms. This table summarizes the findings for all papers across the four design space dimensions: directness, forcefulness, stability, and granularity.

design can be used to overcome limitations in visibility. We hypothesize that these reasons explain the greater popularity of indirect designs.

4.2.2 Forcefulness

The **forcefulness** dimension indicates how intrusively the recommendation is displayed, with values ranging from 5 (most forceful) to 1 (least forceful) as defined in Section 4.1. These ratings are based on two factors which influence the forcefulness of a recommendation: (a) the degree to which a recommendation displaces a user’s context, and (b) the visibility of the recommended content. The five levels, derived from an analysis of the papers collected in our literature review, are summarized in Table 1.

We note that there are some cases (e.g. [25]) where systems recommend content in more than one way, at times with different design choices regarding what contextual information remains visible and how the recommended content is displayed. In these cases, each approach to recommendation is judged independently resulting in some papers being assigned to multiple levels of forcefulness. The remainder of this subsection describes which papers fit into each of the five levels and provides concrete examples for each.

Level 5. This level represents the most forceful methods of presenting recommendations, reflecting designs which require a user to pay full attention to the recommended content. A total of five papers from our review [13, 42, 50, 58, 77] fit within the Level 5 category for forcefulness.

Three of these systems are designed to visualize arbitrarily large datasets [50, 58, 77] through a query-based interaction model. The systems analyze user queries to determine a recommended subset using the most appropriate dimensions of the dataset. Each time the user makes a new request, the visualization is fully refreshed with an entirely new view of new recommended subset of data.

In other work, a variation of this approach was adopted which used multiple visualization sub-panels. In this design, some sub-panels were

used to show detailed content recommendation while other sub-panels were responsible for visualizing the context [42]. This design allowed recommendations with Level 5 forcefulness which replaced the overall view with new content in some sub-panels, as well as recommendations at lower levels that preserved some context in other sub-panels.

Level 4. This level represents the second-most forceful design choices in our design space as defined in Section 4.1. At this level, unlike Level 5, a portion of the visible context is maintained as the recommendation is surfaced, and the recommended content is displayed in addition to that remaining context.

A total of six papers [4, 25, 27, 38, 45, 65] from our search reported designs with Level 4 forcefulness. A common theme is that these systems managed the integration of new recommended content by selectively removing prior content that was deemed least valuable to maintain within the display, while at the same time maintaining other prior content that remained important according to some criteria.

This type of design was commonly observed in visualizations that depict hierarchical structures in some way to manage complexity. For example, Karsai et al. [38] present a prototype for visualizing provenance data that is represented using graph-based data structures. The system provides tools for users to interactively view, cluster, and simplify provenance structures. The visualized provenance graphs are updated in response to user actions, and each time new recommended nodes can be added while some existing but less relevant nodes will be removed from view.

Similarly, one of the two designs from d’Entremont and Storey’s work [25] hides irrelevant nodes and replaced them with a statistically-derived representation of those hidden nodes. This is done to restrict the navigation space for a better focus on recommended nodes.

This type of design has also been leveraged in systems that concurrently visualize some representation of user models alongside a primary work area. For example, Ahn et al. [4] propose a design that augments the display of recommended keywords with a representation

of a model of user interests that is constantly updated in response to user interactions with the overall system. As new recommended keywords are identified and incorporated into the view, less relevant keywords are de-prioritized but remain on display. Previous keywords that have become irrelevant are removed. Similarly, Glowacka et al. [27] propose a prototype that will hide less relevant keywords from the view with recommended contents.

Level 3. This level represents systems which maintain the full current context of a user's visualization while integrating recommendations within that context. This was the most frequent design approach observed in our literature review with 18 papers describing recommendations that fit within the Level 3 category [3–5, 13–15, 25, 32, 37, 39, 42–44, 46, 52, 64, 72, 80].

The most frequently observed design within this category, found in five of the included papers [5, 43, 44, 64, 80], was the basic highlighting of recommended content within the existing visual display. As an example, the system proposed by Bae et al. [5] for document triage highlights recommended documents within its visualization of multiple documents, while a details panel showing information about a single document remains unchanged. Similarly, Lehmann et al. [44] highlights recommended concept clusters as users interact with a concept that appear in a visualization of collections of documents.

The second category of recommendation designs in Level 3 takes a similar but opposite approach: reducing the visual salience of less relevant information rather than highlighting relevant content. A total of four papers [25, 32, 43, 80] leverage this approach. For example, one aspect of a system for visualizing tree-based data proposed by d'Entremont and Storey [25] reduces the opacity level of less relevant tree nodes. Similarly, Gobel et al. [32] describe a gaze-based system that reduces the opacity of map legends that are less relevant to a user's interests.

We note that these two approaches (highlighting new recommendations and decreasing the salience of less relevant information) are compatible strategies that can be used simultaneously. For example, this is observed in work from Lalle et al. [43], as well as Yelizarov and Gamayunov [80].

A third common approach, especially with visualizations that adopt a visual form showing networks or clusters where the position is flexible, uses the positioning of elements to emphasize a recommendation. For example, work from Kanai and Hakozaiki [37] as well as Leuski and Allan [46] use positional shifts to move recommended items "close" to the user in 3D visualizations. More recently, Nazemi et al. [52] followed a similar philosophy to visualize document relationships.

Two projects from Ahn and collaborators [3, 4] adopt a more interactive approach in which a spatial positioning of recommended content is arranged nearby user-specified "points of interest," or POIs, based on relevance. The POIs can then be moved interactively to trigger updates from the system.

Meanwhile, Brusilovsky et al.'s [14] ADVISE II system adopts yet another use of position to highlight recommended content. It uses position to place visual representations of newly recommended information close to previously consumed content that it is related to. This is intended to make it easy for users to find the specific recommended content they would find most useful based on what past content they wish to continue examining.

Level 2. This level represents systems which show recommended content in parallel to the main context. Only one paper in the review adopted this more subtle form of recommendation, making it the least common forcefulness level in our survey. This form of recommendation was observed in work by Vartak et al. [74] on SeeDB, a system that allows users to build visual queries to a database. The system provided recommendations to users by adding them to a stack of possible combinations of visualizations.

Level 1. The least forceful level in our design space shows recommendations only through interface components that are not visible to users without additional interaction. This is the most subtle form of visual content recommendation as it requires users to specifically act before recommended content can be consumed. Two papers from our literature review fall into this category [3, 14]. Both systems show

recommendations on a separate panel of the interface which is normally not visible. Users must click to access the recommended content shown on those panels.

Discussion. Our analysis shows that the extremes of the forcefulness scale are relatively rare in practice, with the majority of papers falling within Level 3. Moreover, six of the 14 papers that included more (Levels 4 or 5) or less (Levels 1 and 2) forceful recommendation designs also incorporated Level 3 design elements.

Between the two extremes, most papers adopted more forceful designs. This includes five papers that used the most extreme Level 5 compared to just three that used any designs that fit within Level 1 or 2. This appears to reflect a goal for developers or designers to ensure that users attend to the recommended content. However, especially given the difficulty that systems often have in modeling user information needs to generate precise recommendations, there appears to be an opportunity to more deeply explore the use of less forceful recommendation designs.

4.2.3 Stability

The **stability** dimension characterizes how often the recommendation contents get updated, impacting the perceived stability of the visual interface for the user. Design choices within this dimension can be classified into three general categories as defined in Section 4.1: periodic, event-driven, and on-demand policy. It is possible, in theory, that systems combine multiple types of update policies to govern when visible recommendations should be refreshed. However, our review of the literature found that, in practice, each paper could be classified into just one category.

Periodic. Systems with periodic stability are those that update the display of recommended content following a regular time interval (e.g., every n seconds). In our review, we found three papers [5, 64, 80] that followed this approach.

Prototypes in this category typically select a time interval for recommended content updates that is fixed to an arbitrary number of seconds. The time interval tends to be small (e.g., less than 10 seconds), and is sometimes determined in part using data collected from users during usability pilot studies. For example, Silva et al. [64] updates its recommendations every 2 seconds, while Yelizarov & Gamayunov [80] updates its recommendations every 5 seconds.

Event-driven. This category represents systems that update their recommendations in response to specific events. The triggering events can be either underlying system events or some form of user action (excluding explicit user requests for recommended content). In total, ten of the 26 papers [13, 14, 27, 32, 37, 42–44, 65, 72] surveyed included event-driven updates to recommended content.

The largest group of paper in this category respond indirectly to user interactions [14, 27, 37, 42, 44, 72]. For example, this category includes systems that update recommended content in response to user actions such as a change to visualization settings, selecting or inspecting visualized objects, or moving and resizing items. We describe this type of content recommendation as an "indirect" response to user activity because the user actions are not specifically designed to request updates to recommended content. Even if content recommendation were removed from a system, the interactions in this category would still make sense as part of the system's interaction design. This is in contrast with on-demand designs in the stability dimension which we describe later in this subsection.

Another group of papers in the event-driven category describe prototypes that respond to eye gazes as the triggering event type [13, 32, 43, 65]. One example in this category is recent research by Lalle et al. [43] in a system designed to support text analysis. The system updates recommendations in response to user eye fixations (in this case, a fixation is defined as maintained gaze on the same location for at least 100 milliseconds). The number of fixations that meet this duration requirement are tracked, and recommendation updates are triggered when a sufficient number of fixations have been detected.

The final type of triggering event we observed is data entry [37]. Kanai and Hakozaiki describe an approach to the content recommendation which maintains a model of user preferences to guide the identifi-

cation of related content. The preference model as well as the update of recommended content within the interface are both triggered by a user's data entry actions as they interact with the system.

On-demand. The single largest category in the stability dimension is on-demand. This descriptor applies to systems that only update the recommendations made via the user interface in response to explicit user requests for an update (e.g., by clicking an update button). This pattern was observed in 13 of the 26 papers in our survey [3, 4, 15, 25, 38, 39, 45, 46, 50, 52, 58, 74, 77].

Discussion. Exactly half of the papers in our literature review adopted an on-demand policy to update recommendations. The on-demand approach is most straightforward and predictable for users, allowing them to control exactly when updates to visualization or other interface elements will occur. This approach, therefore, provides users with the most stable user experience. However, on-demand approaches can also result in delays in the communication of recommended content. Event-driven and periodic updates may result in faster updates, but reduce user agency.

It is also important to note the distinction between computational updates which refresh the content that is identified as best for recommendation, and interface updates which refresh the recommended content that is accessible to users via the user interface. Matching the focus of this paper, the characterizations of prior work reported in this section focuses on interface update policies. In practice, computation updates and interface updates can be decoupled and occur on different schedules.

4.2.4 Granularity

The **granularity** dimension describes the relationship between the unit of content being recommended and the unit of content being interacted with by a user at the time of recommendation. As defined in Section 4.1, our design space includes three distinct ordinal values for the granularity dimension: coarser, similar, and finer. Many systems recommended content at multiple levels of granularity, and our review found at least one example of every possible combination of the three levels.

Coarser Granularity. Coarser granularity designs provide recommendations at a higher level of abstraction or via a larger unit of information. A total of 7 papers [3, 14, 38, 44, 45, 58, 77] were found to recommend content at a coarser granularity.

For example, Roussinov [58] describes a prototype that uses an Adaptive Self-Organizing-Map (ASOM) to model user queries as a set of concepts. While users view collections of documents containing individual concepts, the system can recommend groups of concepts and concept aggregations based on those groups.

In another example, Ahn & Brusilovsky [3] proposed a visualization component called *KnowledgeSea* which supports topic-based navigation. It recommends higher-level concepts based on the topic contents and user interests. Along similar lines, Brusilovsky et al. [14] describe an adaptive navigation panel which records a user's progress through information resources and, in response, reveals relevant topic recommendations. These efforts represent two common use cases for recommendation at a coarser granularity: topic-level summarization, and new topic recommendation.

Similar Granularity. Designs that recommend content that has a similar level of granularity as data the user is already viewing within the visual interface are described as having a similar granularity. This was the most common form of recommendation in our review of the literature, with 20 out of 26 papers falling into this category [3–5, 13–15, 25, 27, 32, 37, 39, 42–46, 50, 64, 74, 80].

Consider Bae et al. [5] as one example. The system provides users with visualizations that depict individual documents, and recommendations are also provided at the document level. Similarly, Kim et al. [39] recommends visual representations of analogous distance descriptors (e.g. a map view showing 8 times the distance between your home and your company) for spatial measurements (e.g. 78 miles). The recommended analogies always have a granularity that is similar to the original expressions.

Finer Granularity. Recommendations that suggest content that is of a smaller size fit within the finer granularity category. A total of 11

papers [3, 4, 15, 38, 42, 45, 52, 58, 65, 72, 77] were identified as fitting into the finer granularity category. As one example, Silva et al. [65], describe an approach which builds a hierarchical view of a large data set, and then provides users with recommendations for specific variables or dimensions that yield deeper, lower-level views of the data which move down the hierarchy.

Discussion. Overall, more than three-fourths of the papers recommend contents with similar granularity. This is perhaps unsurprising as it is typically the most natural level of recommendation. More specifically, the data is already represented at the level of granularity being used to display it within the visual interface. Therefore no extra effort is required to obtain data at that same granularity. Contrast this with coarser or finer granularity recommendation approaches. For any system to recommend content at a different (coarser or finer) granularity, additional work is required to retrieve or construct data at the alternative level of detail.

Similarly, it is often easier to visually integrate similar granularity recommendations into a user interface because visual representations for data may already exist. If recommended content is coarser or finer, visual integration of recommended content to an existing context within the interface can itself become challenging.

The one exception where changes in granularity become more natural is in the presence of hierarchical data. The hierarchical structure works naturally with changes in granularity as systems can leverage the parent-child relationships to move between levels of granularity.

However, while similar granularity designs are the dominant approach, it is also valuable to observe that a large minority of systems (10 of 26) provide recommendations at multiple levels of granularity. In fact, at least one example was found for every possible combination of the three levels of granularity.

5 DISCUSSION

The four-dimensional design space proposed in Section 4 contributes a typology for describing strategies in which recommended content is surfaced in a visual analytic system. This brings several key benefits. First, it allows various content display strategies developed for different applications to be concisely expressed using the same vocabulary (i.e., values along each dimension in the design space). This enables researchers to describe and compare specific existing strategies with others. Second, the design space provides a framework for improving an existing strategy. Once an existing strategy is mapped in the design space, design alternatives can be systematically generated by movements within the dimensions of the design space.

This section discusses the main patterns that emerge from the application of this design space to the papers found in our literature review. This section also highlights important limitations in our approach, and identifies opportunities for future work.

5.1 Common Trends Across Dimensions

The results of characterizing the papers found in our literature review using our four-dimensional design space are reported in Section 4.2 and summarized in Figure 3. From these results, we observed a small number of high-level trends that span across dimensions of the design space.

First, we observed a relationship between forcefulness and stability in which designs with higher levels of forcefulness (Level 4 or Level 5) had more stable update policies. More specifically, of the 11 papers forcefulness at Level 4 or 5, seven fell into the on-demand category for stability. This rate for on-demand approaches (7 of 11) is far higher than the rate for on-demand approaches in less forceful recommendation designs (just 6 of 15). This suggests that designers of more intrusive recommendations, which can cause greater disruption to a user's workflow, felt the need to offer a more predictable experience for users.

Similarly, we found an apparent link between systems that offer recommendations are varied granularity with more stable on-demand designs. More specifically, 7 of 10 papers with multiple levels of granularity fell within the on-demand category for stability. We hypothesize that the motivation for this pattern is that the presence of heterogeneous

granularities of recommendation results in a more complex, cognitively demanding information environment, and that this in turn suggests the need for a more stable recommendation workflow.

Returning to forcefulness, we observed that the most common forcefulness category, Level 3, was typically associated with similar granularity of recommendations (16 of 18). These two approaches are highly compatible in that Level 3 forcefulness designs integrate recommendations within the existing visual context. This can most easily be done with recommended content that is of the same type and granularity as the already visualized information. Recommendations that are of non-similar granularity, meanwhile, occur more frequently with either very forceful or less forceful designs. These extremes of the forcefulness spectrum provide designers with more freedom because the display of recommended content becomes disconnected from the existing visual display.

5.2 Content Recommendation as User Guidance

Content recommendation can be viewed as a specific form of user guidance. Ceneda et al. [17] define guidance in visual analytics as a dynamic process that aims to help users make progress in their analyses. Content recommendation is typically employed to provide users with new information that is related to their analysis, and therefore fits within the general definition of a user guidance technique.

Ceneda et al. defined a general model which describes alternative approaches for user guidance at a high level. This model included several dimensions including the type of knowledge gap that an approach is intended to resolve, the inputs and outputs required, and the degree of guidance. Given the more focused scope of our proposed design space (see Section 4.1), the dimensions we proposed are largely independent of those defined by Ceneda et al.

However, there is some alignment between our forcefulness dimension and Ceneda et al.'s *degree* dimension [17]. Ceneda et al. define guidance degree using three characteristic scenarios: orienting, directing, and prescribing. Orienting helps users understand options for navigation without suggesting system-generated recommendations as to which are better options. In our model, recommended content displayed to achieve an orienting function would map to Level 3 forcefulness. In contrast, directing methods communicate system-generated preferences. In our model, this would map to Level 3 or 4 forcefulness depending on how the interface was designed. Meanwhile, guidance that Ceneda et al. describe as prescribing would fit within Level 4 forcefulness in our model. Ceneda et al. also describe fully-automated guidance which in our model would be considered Level 5 forcefulness.

5.3 Limitations

As with all studies based on a systematic review of the literature, the findings in this paper are inherently dependent on the choice of search terms. As outlined in Section 3.1, the focus of this search was difficult to specify in terms of keywords without being overly narrow or overly broad. The final choice of terms was a compromise that yielded relevant articles without also returning result sets in the tens or hundreds of thousands which would make a full-text review logistically impossible. The choice to scope the search to IEEE Xplore the ACM Digital Library also helped focus the scope of the search, but at the expense of eliminating some potentially relevant journals or proceedings. We attempted to mitigate these limitations by adding an extra step to our search process. As described in Section 3.3, we included in our search all papers cited by the papers identified in our original keyword search. This did identify a small number of additional papers from beyond the IEEE and ACM archives. However, we acknowledge that relevant articles may still fall out of our search (e.g., recommendations for next-step actions in analytical workflows [67, 75]). Nevertheless, we believe that the general nature of our proposed design space means that it can also be applied in principle to characterize many approaches described in papers missed by our search.

Another limitation relates to our use of an open coding process described in Section 3.5. Alternative coding approaches could have been followed and might have produced somewhat different results. Relatedly, we recognize that the manual tagging of papers within our

design space includes some subjective judgements that are necessarily made based only on the authors' reading of the published articles. To mitigate the risk of errors, every paper was reviewed by two people who worked independently. However, mischaracterizations are still possible based on the limited information provided in some articles.

5.4 Opportunities for Future Work

There are several possible directions for future research that builds on the results presented in this paper. For example, the relationships between dimensions highlighted in Section 5.1 are based on observations from the analysis of the literature. Experimental studies designed to better understand the potential design tradeoffs between those dimensions would be a valuable next step.

In addition, the results summarized in Figure 3 highlight some areas of the design space are under-explored. Perhaps most interestingly, there is relatively little work exploring more "subtle" aspects of content recommendation. This includes low forcefulness designs, and designs with periodic updates in the stability dimension. These represent opportunities for research exploring potential use cases where these types of designs might be beneficial. For example, it is possible that these subtle design choices could potentially allow for the display of recommended content in a way that supports serendipitous discovery with negligible impact on users' workflows.

Finally, while this paper has focused on the visual display of recommended content, future studies examining design decisions for the underlying models and algorithms that support content recommendation could be valuable. Moreover, results from such a study might help uncover how those models and algorithms can directly influence the eventual visual presentation of recommended content. For example, approaches that leverage hierarchical models might encourage visual designs that surface recommendations at multiple levels of granularity.

6 CONCLUSION

In this paper, we explored the design space for surfacing recommended content within visual analytics platforms. Like visualization recommendations, which aim to recommend the best form of a visualization to users to advance their analytic goals, content recommendation techniques aim to help users discover related information that they might otherwise overlook. More specifically, content recommendation in visual analytics requires technologies which capture some form of model of a user's information needs given their analytic task context, use that model to identify relevant content, and then surface that recommended content within the user interface to communicate those recommendations to users.

Content recommendation can take many forms, and a wide variety of techniques have been proposed to address this general problem. Moreover, the interface for exposing recommended content to users requires a large number of design decisions which can directly impact the usability and utility of the underlying algorithmic recommendation capabilities.

This paper describes our efforts to develop a better understanding of the design space for how recommended content can be surfaced within visual analytic platforms. First, we conducted a systematic review of the relevant literature. We then analyzed the publications identified from that search to develop a multi-dimensional design space for visual content recommendation. We formally defined four distinct dimensions in our design space—directness, forcefulness, stability, and granularity—and we outlined the various design alternative within each of those dimensions. We then characterized each of the papers identified in our literature review using the four dimensional design space to gain insights into the design choices made in prior work. While the detailed analysis shows a number of key patterns and commonly adopted approaches, it also highlights that the "best" design solution is often highly context dependent with each design choice offering its own benefits under certain conditions.

ACKNOWLEDGMENTS

The research reported in this article was supported in part by a grant from the National Science Foundation (#1704018).

REFERENCES

- [1] ACM Digital Library. <https://dl.acm.org/>. Accessed: 2022-03-24.
- [2] J.-W. Ahn and P. Brusilovsky. Envisioning user models for adaptive visualization. In *2008 IEEE Symposium on Visual Analytics Science and Technology*. IEEE, 2008.
- [3] J.-W. Ahn and P. Brusilovsky. Guiding educational resources for iSchool students with topic-based adaptive visualization. In *Proceedings of the 2011 IConference*. ACM, 2011.
- [4] J.-W. Ahn, P. Brusilovsky, and S. Han. Personalized Search: Reconsidering the Value of Open User Models. In *Proceedings of the 20th International Conference on Intelligent User Interfaces, IUI '15*, pp. 202–212. ACM, New York, NY, USA, Mar. 2015.
- [5] S. Bae, D. Kim, K. Meintanis, J. M. Moore, A. Zacchi, F. Shipman, H. Hsieh, and C. C. Marshall. Supporting document triage via annotation-based multi-application visualizations. In *Proceedings of the 10th annual joint conference on Digital libraries, JCDL '10*, pp. 177–186. ACM, New York, NY, USA, June 2010.
- [6] X. Bai, D. White, and D. Sundaram. Context adaptive visualization for effective business intelligence. In *2013 15th IEEE International Conference on Communication Technology*. IEEE, 2013.
- [7] O. Banos, M. B. Amin, W. A. Khan, T. Ali, M. Afzal, B. H. Kang, and S. Lee. Mining minds: an innovative framework for personalized health and wellness support. In *Proceedings of the 9th International Conference on Pervasive Computing Technologies for Healthcare, PervasiveHealth '15*, pp. 1–8. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), Brussels, BEL, May 2015.
- [8] C. Bao, S. Li, S. Flores, M. Correll, and L. Battle. Recommendations for visualization recommendations: Exploring preferences and priorities in public health. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems, 2022*.
- [9] O. Barral, S. LallÉ, A. Iranpour, and C. Conati. Effect of Adaptive Guidance and Visualization Literacy on Gaze Attentive Behaviors and Sequential Patterns on Magazine-Style Narrative Visualizations. *ACM Trans. on Interactive Intelligent Systems*, 11(3-4):28:1–28:46, Aug. 2021.
- [10] O. Barral, S. Lallé, and C. Conati. Understanding the effectiveness of adaptive guidance for narrative visualization: a gaze-based analysis. In *Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI '20*, pp. 1–9. ACM, New York, NY, USA, Mar. 2020.
- [11] J. Beel, B. Gipp, S. Langer, and C. Breitingner. Paper recommender systems: a literature survey. *International Journal on Digital Libraries*, 17(4):305–338, 2016.
- [12] J. Bernard, M. Hutter, M. Zeppelzauer, D. Fellner, and M. Sedlmair. Comparing visual-interactive labeling with active learning: An experimental study. *IEEE transactions on visualization and computer graphics*, 24(1):298–308, 2017.
- [13] S. Bonada, R. Veras, and C. Collins. Personalized Views for Immersive Analytics. In *Proceedings of the 2016 ACM Companion on Interactive Surfaces and Spaces*. ACM, 2016.
- [14] P. Brusilovsky, Jae-Wook Ahn, T. Dumitriu, and M. Yudelson. Adaptive Knowledge-Based Visualization for Accessing Educational Examples. In *Tenth Inter. Conf. on Information Visualisation (IV'06)*. IEEE, 2006.
- [15] P. Brusilovsky and T. D. Loboda. WADEIn II: a case for adaptive explanatory visualization. In *Proceedings of the 11th annual SIGCSE conference on Innovation and technology in computer science education, ITICSE '06*, pp. 48–52. ACM, New York, NY, USA, June 2006.
- [16] D. Burkhardt, K. Nazemi, and E. Ginters. Innovations in Mobility and Logistics: Assistance of Complex Analytical Processes in Visual Trend Analytics. In *61st Inter. Scientific Conf. on Information Technology and Management Science of Riga Technical University (ITMS)*. IEEE, 2020.
- [17] D. Ceneda, T. Gschwandner, T. May, S. Miksch, H.-J. Schulz, M. Streit, and C. Tominski. Characterizing guidance in visual analytics. *IEEE transactions on visualization and computer graphics*, 23(1):111–120, 2016.
- [18] A. Chaudhuri and Han-Wei Shen. A self-adaptive treemap-based technique for visualizing hierarchical data in 3D. In *2009 IEEE Pacific Visualization Symposium*. IEEE, 2009.
- [19] C. Conati and H. Maclaren. Exploring the role of individual differences in information visualization. In *Proceedings of the working conference on Advanced visual interfaces, AVI '08*, pp. 199–206. ACM, New York, NY, USA, May 2008.
- [20] V. Coors. Resource-adaptive interactive 3D maps. In *Proceedings of the 2nd international symposium on Smart graphics, SMARTGRAPH '02*, pp. 140–144. ACM, New York, NY, USA, June 2002.
- [21] P. Covington, J. Adams, and E. Sargin. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*, pp. 191–198, 2016.
- [22] P. Cremonesi, M. Elahi, and F. Garzotto. User interface patterns in recommendation-empowered content intensive multimedia applications. *Multimedia Tools and Applications*, 76(4):5275–5309, 2017. Publisher: Springer Science and Business Media LLC.
- [23] T. Deuschel. On the Influence of Human Factors in Adaptive User Interface Design. In *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization, UMAP '18*, pp. 187–190. ACM, New York, NY, USA, July 2018.
- [24] G. Domik and B. Gutkauf. User modeling for adaptive visualization systems. In *Proceedings IEEE Visualization '94*. IEEE Comput. Soc. Press, 1994.
- [25] T. d'Entremont and M.-A. Storey. Using a degree-of-interest model for adaptive visualizations in Protégé. In *Presentation Abstracts*, p. 34. Cite-seer, 2006.
- [26] A. Endert, W. Ribarsky, C. Turkay, B. W. Wong, I. Nabney, I. D. Blanco, and F. Rossi. The state of the art in integrating machine learning into visual analytics. *Computer Graphics Forum*, 36(8):458–486, 2017.
- [27] D. Glowacka, T. Ruotsalo, K. Konuyshkova, k. Athukorala, S. Kaski, and G. Jacucci. Directing exploratory search: reinforcement learning from user interactions with keywords. In *Proceedings of the 2013 international conference on Intelligent user interfaces, IUI '13*, pp. 117–128. ACM, New York, NY, USA, Mar. 2013.
- [28] M. Golemati, C. Halatsis, C. Vassilakis, A. Katifori, and G. Lepouras. A Context-Based Adaptive Visualization Environment. In *Tenth International Conference on Information Visualisation (IV'06)*. IEEE, 2006.
- [29] Google Scholar. <https://scholar.google.com/>. Accessed: 2022-03-24.
- [30] D. Gotz and Z. Wen. Behavior-driven visualization recommendation. In *Proceedings of the 14th International Conference on Intelligent User Interfaces, IUI '09*, p. 315–324. ACM, New York, NY, USA, 2009.
- [31] M.-L. Guo, M.-H. Wu, and L.-M. Luo. A Real-Time and Multi-Dimension Data Types Aware Visualization Technology in Personalized Learning System. In *Proceedings of the 2017 International Conference on Wireless Communications, Networking and Applications, WCNA 2017*, pp. 122–126. ACM, New York, NY, USA, Oct. 2017.
- [32] F. Göbel, P. Kiefer, I. Giannopoulos, A. T. Duchowski, and M. Raubal. Improving map reading with gaze-adaptive legends. In *Proc. of the ACM Symposium on Eye Tracking Research & Applications*, pp. 1–9, 2018.
- [33] K. Hu, M. A. Bakker, S. Li, T. Kraska, and C. Hidalgo. Vizml: A machine learning approach to visualization recommendation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19*, p. 1–12. ACM, New York, NY, USA, 2019.
- [34] IEEE Xplore. <https://ieeexplore.ieee.org/Xplore/home.jsp>. Accessed: 2022-03-24.
- [35] D. Jannach, M. Zanker, A. Felfernig, and G. Friedrich. *Recommender systems: an introduction*. Cambridge University Press, 2010.
- [36] Jianwei Yu and Bisheng Yang. A framework of spatio-temporal data adaptive visualizations for mobile environment. In *2009 17th International Conference on Geoinformatics*. IEEE, 2009.
- [37] H. Kanai and K. Hakozaiki. A browsing system for a database using visualization of user preferences. In *2000 IEEE Conference on Information Visualization. An International Conference on Computer Visualization and Graphics*, pp. 277–282. IEEE, 2000.
- [38] L. Karsai, A. Fekete, J. Kay, and P. Missier. Clustering provenance facilitating provenance exploration through data abstraction. In *Proceedings of the Workshop on Human-In-the-Loop Data Analytics, HILDA '16*, pp. 1–5. ACM, New York, NY, USA, June 2016.
- [39] Y.-S. Kim, J. Hullman, and M. Agrawala. Generating Personalized Spatial Analogies for Distances and Areas. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, CHI '16*, pp. 38–48. ACM, New York, NY, USA, May 2016.
- [40] Y.-S. Kim, P. Kayongo, M. Grun-de-Mclaughlin, and J. Hullman. Bayesian-Assisted Inference from Visualized Data. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):989–999, 2021. Publisher: Institute of Electrical and Electronics Engineers (IEEE).
- [41] B. P. Knijnenburg, M. C. Willemsen, Z. Gantner, H. Soncu, and C. Newell. Explaining the user experience of recommender systems. *User modeling and user-adapted interaction*, 22(4):441–504, 2012.
- [42] J. Knittel, S. Koch, T. Tang, W. Chen, Y. Wu, S. Liu, and T. Ertl. Real-Time Visual Analysis of High-Volume Social Media Posts. *IEEE Trans. on Visu-*

- alization and Computer Graphics, 28(1):879–889, Jan. 2022. Conference Name: IEEE Transactions on Visualization and Computer Graphics.
- [43] S. Lalle, D. Toker, and C. Conati. Gaze-Driven Adaptive Interventions for Magazine-Style Narrative Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 27(6):2941–2952, 2021. Publisher: Institute of Electrical and Electronics Engineers (IEEE).
- [44] S. Lehmann, U. Schwanecke, and R. Dörner. Interactive visualization for opportunistic exploration of large document collections. *Information Systems*, 35:260–269, Apr. 2010.
- [45] Lei Shi, N. Cao, S. Liu, Weihong Qian, L. Tan, Guodong Wang, Jimeng Sun, and Ching-Yung Lin. HiMap: Adaptive visualization of large-scale online social networks. In *2009 IEEE Pacific Visualization Symposium*. IEEE, 2009.
- [46] A. Leuski and J. Allan. Interactive Information Retrieval Using Clustering and Spatial Proximity. *User Modeling and User-Adapted Interaction*, 14(2-3):259–288, June 2004.
- [47] A. Liberati, D. G. Altman, J. Tetzlaff, C. Mulrow, P. C. Gøtzsche, J. P. Ioannidis, M. Clarke, P. J. Devereaux, J. Kleijnen, and D. Moher. The prisma statement for reporting systematic reviews and meta-analyses of studies that evaluate health care interventions: explanation and elaboration. *Journal of clinical epidemiology*, 62(10):e1–e34, 2009.
- [48] H. N. Miyamura, S. Ohzahata, Y. Shinano, and R. Miyashiro. Adaptive view-dependent tree graph visualization. In *2011 Eighth International Joint Conference on Computer Science and Software Engineering (JCSSE)*. IEEE, 2011.
- [49] D. Moritz, C. Wang, G. L. Nelson, H. Lin, A. M. Smith, B. Howe, and J. Heer. Formalizing visualization design knowledge as constraints: Actionable and extensible models in draco. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):438–448, 2019.
- [50] B. Mutlu, E. Veas, and C. Trattner. VizRec: Recommending Personalized Visualizations. *ACM Transactions on Interactive Intelligent Systems*, 6(4):31:1–31:39, Nov. 2016.
- [51] K. Nazemi, A. Kuijper, M. Hutter, J. Kohlhammer, and D. W. Fellner. Measuring context relevance for adaptive semantics visualizations. In *Proceedings of the 14th International Conference on Knowledge Technologies and Data-driven Business - i-KNOW '14*, pp. 1–8. ACM Press, Graz, Austria, 2014.
- [52] K. Nazemi, R. Retz, J. Bernard, J. Kohlhammer, and D. Fellner. Adaptive Semantic Visualization for Bibliographic Entries. In *International Symposium on Visual Computing*, pp. 13–24, July 2013.
- [53] N. Oscar, S. Mejía, R. Metoyer, and K. Hooker. Towards Personalized Visualization. In *Proceedings of the 2017 Conference on Designing Interactive Systems*. ACM, 2017.
- [54] M. J. Pazzani and D. Billsus. Content-based recommendation systems. In *The adaptive web*, pp. 325–341. Springer, 2007.
- [55] P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis*, vol. 5, pp. 2–4. McLean, VA, USA, 2005.
- [56] P. Pombinho, A. P. Afonso, and M. B. Carmo. Chameleon – A Context Adaptive Visualization Framework for a Mobile Environment. In *2011 15th International Conference on Information Visualisation*. IEEE, 2011.
- [57] S. F. Roth, J. Kolojchick, J. Mattis, and J. Goldstein. Interactive graphic design using automatic presentation knowledge. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '94, p. 112–117. ACM, New York, NY, USA, 1994.
- [58] D. Roussinov. Internet search using adaptive visualization. In *CHI '99 Extended Abstracts on Human Factors in Computing Systems*. ACM Press, 1999.
- [59] T. Ruotsalo, G. Jacucci, P. Myllymäki, and S. Kaski. Interactive intent modeling: information discovery beyond search. *Communications of the ACM*, 58(1):86–92, Dec. 2014.
- [60] A. Rusu, C. Clement, and R. Jianu. Adaptive Binary Trees Visualization with Respect to User-Specified Quality Measures. In *Tenth International Conference on Information Visualisation (IV'06)*. IEEE, 2006.
- [61] D. Sacha, A. Stoffel, F. Stoffel, B. C. Kwon, G. Ellis, and D. A. Keim. Knowledge generation model for visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1604–1613, 2014.
- [62] B. Schönhage and A. Eliëns. A flexible architecture for user-adaptable visualization. In *Proceedings of the 1997 workshop on New paradigms in information visualization and manipulation*, NPIV '97, pp. 8–10. ACM, New York, NY, USA, Nov. 1997.
- [63] G. Shani and A. Gunawardana. Evaluating recommendation systems. In *Recommender systems handbook*, pp. 257–297. Springer, 2011.
- [64] N. Silva, T. Schreck, E. Veas, V. Sabol, E. Eggeling, and D. W. Fellner. Leveraging eye-gaze and time-series features to predict user interests and build a recommendation model for visual analysis. In *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications*, ETRA '18, pp. 1–9. ACM, New York, NY, USA, June 2018.
- [65] N. Silva, L. Shao, T. Schreck, E. Eggeling, and D. Fellner. Visual Exploration of Hierarchical Data Using Degree-of-Interest Controlled by Eye-Tracking. In *Proceedings of the 9th Forum Media Technology*, 2016.
- [66] B. Smith and G. Linden. Two decades of recommender systems at amazon.com. *IEEE Internet Computing*, 21(3):12–18, 2017.
- [67] A. Srinivasan and V. Setlur. Snowy: Recommending utterances for conversational visual analysis. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, pp. 864–880, 2021.
- [68] B. Steichen, G. Carenini, and C. Conati. User-adaptive information visualization: using eye gaze data to infer visualization tasks and user cognitive abilities. In *Proceedings of the 2013 international conference on Intelligent user interfaces*, IUI '13, pp. 317–328. ACM, New York, NY, USA, Mar. 2013.
- [69] C. D. Stolper, A. Perer, and D. Gotz. Progressive visual analytics: User-driven visual exploration of in-progress analytics. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1653–1662, 2014.
- [70] J. Thomas and K. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. National Visualization and Analytics Ctr, 2005.
- [71] D. Toker, C. Conati, B. Steichen, and G. Carenini. Individual user characteristics and information visualization: connecting the dots through eye tracking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 295–304. ACM, New York, NY, USA, Apr. 2013.
- [72] B. van Tonder and J. Wesson. Using adaptive interfaces to improve mobile map-based visualisation. In *Proceedings of the 2008 annual research conference of the South African Institute of Computer Scientists and Information Technologists on IT research in developing countries: riding the wave of technology*, SAICSIT '08, pp. 257–266. ACM, New York, NY, USA, Oct. 2008.
- [73] M. Vartak, S. Huang, T. Siddiqui, S. Madden, and A. Parameswaran. Towards Visualization Recommendation Systems. *ACM SIGMOD Record*, 45(4):34–39, May 2017.
- [74] M. Vartak, A. Parameswaran, N. Polyzotis, and S. R. Madden. SEEDB: automatically generating query visualizations. *Proc. VLDB Endow.*, 2014. ISBN: 2150-8097 Publisher: AC.
- [75] X. Wang, F. Cheng, Y. Wang, K. Xu, J. Long, H. Lu, and H. Qu. Interactive data analysis with next-step natural language query recommendation. *arXiv preprint arXiv:2201.04868*, 2022.
- [76] R. Wasinger, J. Wallbank, L. Pizzato, and J. Kay. Scrutable User Models and Personalised Item Recommendation in Mobile Lifestyle Applications.
- [77] W. Wiza, W. Cellary, and K. Walczak. A Method of Holistic 3D Visualization of Arbitrarily Large Datasets. In *Seventh IEEE International Symposium on Multimedia (ISM'05)*. IEEE, 2005.
- [78] W. Wiza, K. Walczak, and W. Cellary. Periscope: a system for adaptive 3D visualization of search results. In *Proceedings of the ninth international conference on 3D Web technology - Web3D '04*, p. 29. ACM Press, Monterey, California, 2004.
- [79] A. Wu, L. Xie, B. Lee, Y. Wang, W. Cui, and H. Qu. Learning to Automate Chart Layout Configurations Using Crowdsourced Paired Comparison. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, number 14, pp. 1–13. ACM, New York, NY, USA, May 2021.
- [80] A. Yelizarov and D. Gamayunov. Adaptive Visualization Interface That Manages User's Cognitive Load Based on Interaction Characteristics. In *Proceedings of the 7th International Symposium on Visual Information Communication and Interaction*. ACM Press, 2014.
- [81] H. Yu, C. Wang, and K.-L. Ma. Parallel hierarchical visualization of large time-varying 3D vector fields. In *Proceedings of the 2007 ACM/IEEE conference on Supercomputing*, SC '07, pp. 1–12. ACM, New York, NY, USA, Nov. 2007.
- [82] Y. Zhang, J. Zhu, W. Li, Y. Hu, and M. Liu. Adaptive Web 3D Visualization for Diverse Terminals. In *2018 26th International Conference on Geoinformatics*. IEEE, 2018.
- [83] Z. Zhou, X. Wen, Y. Wang, and D. Gotz. Modeling and Leveraging Analytic Focus During Exploratory Visual Analysis. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, 2021.