



REGULAR PAPER

Chen He · Luana Micallef · Barış Serim · Tung Vuong · Tuukka Ruotsalo · Giulio Jacucci

## Interactive visual facets to support fluid exploratory search

Received: 17 January 2022 / Revised: 22 May 2022 / Accepted: 4 July 2022 / Published online: 30 July 2022  
© The Author(s) 2022

**Abstract** Exploratory search starts with ill-defined goals and involves browsing, learning, and formulating new targets for search. To fluidly support such dynamic search behaviours, we focus on devising *interactive visual facets (IVF)*, visualising information facets to support user comprehension and control of the information space. We reviewed existing faceted search interfaces and derived two design requirements (DR) that have not been fully addressed to support fluid interactions in exploratory search. We then exemplified the requirements through devising an IVF tool, which coordinates a linear and a categorical facet representing the distribution and summarisation of items, respectively, and providing context for faceted exploration (DR1). To support rapid transitions between search criteria (DR2), the tool introduces a novel design concept of using facets to select items without filtering the item space. Particularly, we propose a filter-swipe technique that enables users to drag a categorical facet value sequentially over linear facet bars to view the items in the intersection of the two facets along with the categorical facet dynamically summarising the items in the intersection. Three applications with various datasets demonstrate how the features support information discovery with ease. A comparison with a baseline system provided evidence that the task performance of the IVF tool was comparable to the typical query search interface. Another study of 11 participants with realistic email search tasks shows that dynamic suggestions through the timeline navigation can help discover useful suggestions for search; the novel design concept was favoured over using facet values as filters. Based on these practices, we derive IVF design implications for fluid, exploratory searches.

**Keywords** Exploratory search · Information facets · Fluid interaction · Visual data exploration

---

C. He (✉) · B. Serim · T. Vuong · T. Ruotsalo · G. Jacucci  
Department of Computer Science, University of Helsinki, Helsinki, Finland  
E-mail: chen.he@helsinki.fi

B. Serim  
E-mail: baris.serim@helsinki.fi

T. Vuong  
E-mail: vuong@cs.helsinki.fi

T. Ruotsalo  
E-mail: tuukka.ruotsalo@helsinki.fi

G. Jacucci  
E-mail: giulio.jacucci@helsinki.fi

L. Micallef  
Department of Computer Science, University of Copenhagen, Copenhagen, Denmark

## 1 Introduction

Information search is an essential activity that we carry out on a daily basis. Modern browsers have, therefore, seamlessly integrated search functionalities by enabling users to simply input search queries in the top address bar and retrieve a ranked list of relevant links. However, unlike simple lookup searches, users often search with ambiguous information needs or uncertain goals (Athukorala et al. 2016). For instance, when looking for related literature, researchers often start with vague, general queries; when searching for an email, users typically do not have enough memory clues to do a simple lookup search. In such situations, users dynamically modify their search directions as they gain further knowledge of the search space (White and Roth 2009; Wildemuth and Freund 2012). Prior work termed this type of search as *exploratory search* (Marchionini 2006; Ruotsalo et al. 2018) to indicate its learning, investigative nature.

With exploratory search, a query box plus a ranked list of results provide limited support to users' evolving information needs. Advanced visual interfaces are required to support the iterative and incremental process fluidly. Wildemuth and Freund (2012) identified a set of characteristics of exploratory search tasks based on a literature review. Exploratory search tasks focus on learning and investigative search goals; they are open-ended, dynamic, and motivated by ill-defined problems (White and Roth 2009; Wildemuth and Freund 2012), which implies the dynamism of user search behaviours. How to design visual interfaces to fluidly support exploratory search is our main research question.

Facets, composed of orthogonal sets of categories (Hearst et al. 2002), can organise information into meaningful groups (White and Roth 2009; Wildemuth and Freund 2012). Studies on search interfaces with information facets show that users use facets more frequently for exploratory search tasks (Niu et al. 2019; Niu and Hemminger 2015); information facets are essential to support both the expansion and refinement of search queries (White and Roth 2009; Wildemuth and Freund 2012). Recent work on faceted search put more effort into algorithm development and user evaluation (Niu et al. 2019). For instance, techniques have been developed to generate nested facets from ontologies (Zhang et al. 2020) and to predict facets of interest based on user interactions (Chantamunee et al. 2020). However, interactive visualisation design is also essential in supporting faceted exploration (Mahdi et al. 2020).

From a human-computer interaction point of view, there are many benefits of using facets to support fluid exploratory searches. First, facets can serve as entry points for search, which is particularly beneficial when users start with no specific queries in mind (Niu and Hemminger 2015). Second, facets can provide an overview of the search space, which avoids users getting lost in the search process (Niu and Hemminger 2015; Manioudakis and Tzitzikas 2020). Third, facets can guide the next relevant directions to explore and prevent zero-hit queries (Manioudakis and Tzitzikas 2020). Ruotsalo et al. (2020) found that faceted query suggestion increases engagement and supports recall.

Therefore, our goal is to use information facets to address the challenge of supporting fluid interactions in exploratory search. To do this, we first reviewed prior research on faceted search interfaces, identifying their shortcomings in supporting fluid search behaviours. To address these shortcomings, we devised a faceted search interface, which was evaluated through three use cases and two user studies. This article extends an earlier publication (He et al. 2021) by 1) implementing two use cases from different fields, including serendipitous tweets discovery and oncogene occurrence analysis, to demonstrate the generalisability of the tool and 2) conducting a study with a baseline system to investigate the task performance of the tool.

The contribution of this research is three-fold. First, we distill from related work the concept of *interactive visual facets (IVF)* and two design requirements (DRs) to support fluid exploratory search, i.e. providing contextual information for faceted exploration and using facets to support rapid transitions between search criteria. Second, we exemplify the two DRs by devising an IVF tool. The tool coordinates a linear and a categorical facet and introduces a novel design concept – using facet values to select and inspect items without filtering the item space – to support flexible and dynamic search attempts without losing the exploration context. Particularly, users can drag a categorical facet value sequentially over linear facet bars to view the items in the intersection of the two facet values; meanwhile, the categorical facet changes dynamically to summarise the items in the intersection. Third, to demonstrate the fluidity and usefulness of the features in supporting exploratory search, we present three use cases and two user studies. The three cases illustrate how the tool implements the requirements in three different applications. A comparative study with a baseline system evidenced the task performance of the IVF tool is comparable to a typical query search interface. A second study with realistic email search tasks shows that the feature of using facets to select items without filtering the item space was favoured over using facets as filters. We discuss design

implications based on our practices to guide how to effectively design IVF for fluid exploratory searches in the future.

## 2 Related work

One of the earliest faceted search interfaces, Flamenco (Hearst et al. 2002), listed hierarchically faceted metadata to guide users toward possible choices. Subsequent faceted search interfaces extended on it from two aspects by (1) visualising facets to aid user comprehension of the information space (e.g. Dörk et al. 2008, 2010; Seifert et al. 2014; Yalcin et al. 2017) and (2) taking facets as user-manipulable objects to support information exploration (e.g. Wilson et al. 2006; Seifert et al. 2014; Zhao et al. 2013; Klouche et al. 2015; Dörk et al. 2012b). We review related work from the two aspects.

Apart from listing faceted keywords for exploration, facets are frequently visualised in coordinated views. Each view depicts one type of facet, such as chronological and geographical views of the search space; e.g. see Dörk et al. (2008, 2010); Seifert et al. (2014); Yalcin et al. (2017); Rauch et al. (2015); Qu et al. (2020). Brushing and linking techniques are often used in these cases to support exploration of item distribution in various views. For instance, VisGets (Dörk et al. 2008) coordinates temporal, spatial, and topical views of items to enable users to visually formulate queries from the three aspects in parallel. Likewise, Keshif visualises data properties in coordinated views and uses linked brushing and cross-filtering for exploration (Yalcin et al. 2017).

On the other hand, individual items can be directly drawn upon a facet layer using zooming and panning for item inspection; e.g. see Dörk et al. (2012a); Grierson et al. (2015); Viégas et al. (2006); Mauro et al. (2020). In this case, only 1-2 facets are visualised at one time. Themail (Viégas et al. 2006) displays word lists along a timeline to show the evolution of user conversation with an individual. Faceted selection and filtering are supported for detailed inspection. Likewise, Mauro et al. (2020) arrange items categorised by colour on a map. Items could be filtered by facets using checkboxes, treemaps, or sunburst views.

The above two techniques could also be integrated to facilitate exploration, enabling item filtering through coordinated views and showing filtered items over existing or additional facets (Rauch et al. 2015; Wang et al. 2019; Zeng et al. 2021). Allowing item filtering from multiple categorical facets, such as authors and keywords, VIStory (Zeng et al. 2021) depicts filtered items on a timeline colour-coded by user-selected facet groups.

To facilitate user exploration, existing visualisations generally support Boolean queries (e.g. Seifert et al. 2014; Zhao et al. 2013; Yalcin et al. 2017; Mauro et al. 2020; Zeng et al. 2021) or weight adjustment of facet values (e.g. Klouche et al. 2015; Chang et al. 2019; di Sciascio et al. 2016, 2018; Klouche et al. 2017). For Boolean queries, intra-facet selections filter items through a logical OR operation, whereas inter-facet selections create a logical AND operation for filtering. For instance, PivotSlice (Zhao et al. 2013) enables data slices in a matrix view through dragging and dropping numeric or categorical facet values to two perpendicular axes. The two query axes support flexible Boolean combinations of queries, which results in a logical OR operation on values in the same facet panel of an axis and a logical AND operation over the corresponding panels of the two axes.

Rather than simply using facets to filter results, SearchLens (Chang et al. 2019) allows users to personalise search results by adjusting the importance of facet values. In uRank, users can reorder search results according to their relevancy to weighted keywords (di Sciascio et al. 2016, 2017, 2020b). An improved version of uRank incorporated social aspects where users weigh a hybrid model to rank results (di Sciascio et al. 2018, 2020a).

However, the selection or weight adjustment of facet values often results in drastic layout changes of the item space, inhibiting rapid transitions between search criteria. For instance, PivotPaths visualises facet values in nodes and allows pivot operations to transform node layouts to support, such as single-facet querying and double-facet comparison (Dörk et al. 2012b). Similarly, Fluid Views uses dual layers to position items on a context map, such as a timeline or a topic overview (Dörk et al. 2012a). Semantic zooming allows transitions between item overview and details for inspections.

To better support rapid transitions between search criteria, query previews (Plaisant et al. 1999; Qvarfordt et al. 2013) are used to put users in the current search context and avoid dramatic layout changes before users are certain of their next search steps. Existing faceted search interfaces have incorporated query previews in two ways by showing the number of results related to individual facet values (Hearst et al. 2002; Seifert et al. 2014; Kreutz et al. 2018; Rauch et al. 2015; Zeng et al. 2021) or by highlighting relevant facet

values and result items to an attempted query (Dörk et al. 2008, 2010; Yalcin et al. 2017). For instance, FacetScape enables users to view the changes in the number of results before selecting or deselecting facet values (Seifert et al. 2014). MSpace (Wilson et al. 2006) allows users to arrange facets hierarchically in columns, from left to right, and provides item previews when hovering over facet values to assist in choice-making. This low-cost technique of using mouse-over interactions to trigger query previews has also been widely used in coordinated views. As an example, mousing over a data attribute in Keshif highlights the distribution of the data items in other views (Yalcin et al. 2017).

The low discontinuity between displayed information invoked by query previews facilitates low-cost, rapid transitions between different queries. However, more preview techniques could be devised to leverage the advantage better.

### 3 Definition and DRs

As elaborated in Sect. 2, existing faceted search interfaces expanded on *information facets* in two directions, incorporating visualisation and interaction techniques. Thus, we propose to name information facets as *interactive visual facets (IVF)* for the purpose of devising faceted interfaces to assist searches. We define IVF in exploratory search as *visualising information facets to support user comprehension and control of the information space*. This research aims to devise IVF to support fluid exploratory searches.

Elmqvist et al. (2011) characterised fluidity in information visualisation as smooth interaction, responsive graphics, and conscientious user experience. To support fluid exploratory search, Klouche et al. (2018) proposed to use entity affordances to devise interaction techniques. Entities can yield other relevant entities to facilitate exploration; entities can be organised to support pattern recognition; entities can be shared to assist in collaboration (e.g. Andolina et al. 2018). Kules and Shneiderman (2008) proposed a set of design guidelines for using categorised facets for exploratory search. Shneiderman et al. (1997) proposed a four-phase framework for text searches that focused on the iterations of query refinement. However, exploratory searches usually go beyond search queries and involve other types of interactions, such as navigation and browsing. The Flamenco project suggested integrating a direct search method and browsing to enable smooth transitions from one search direction to the next, without users feeling lost or stuck (Hearst et al. 2002).

To summarise, fluidity in exploratory search concerns the beginning, middle, and end of a search. As discussed in Sect. 1, facets can easily address the search entry points issue and avoid dead ends. Thus, we attempted to devise IVF during the search to provide users with a fluid experience. To this end, we derived two DRs that have not been fully addressed according to the analysis in Sect. 2 to drive the design of a fluid, exploratory search interface:

- *DR1: Provide contextual information for faceted exploration* Contextual information can avoid users feeling disconnected from or lost in their search experience (Hearst et al. 2002; Dörk et al. 2012b). Context conveys the time and space of users' exploration. Time-related context informs users of the exploration process. For instance, which items have/have not been retrieved as results. As Jankun-Kelly et al. (2007) stated, a visualisation that does not communicate the states of user visits could be inefficient in assisting data exploration. A review of prior work in Sect. 2 suggests that this time aspect of the requirement was often neglected. Space-related context informs users about the current search space. As stated in Sect. 1, visualising facets is beneficial in providing an overview of the search space. For example, to guide publication exploration, the facets of related authors and keywords shown in PivotPath depict context of current search queries (Dörk et al. 2012b). Interactions between facets, such as brushing and linking among facet views, can reveal additional relationships in the search space. This requirement emphasises that incorporate both aspects of context (time and space) is necessary to provide a fluid exploration experience.
- *DR2: Use facets to support rapid transitions between search criteria* During an exploratory search, users would modify information needs with the accumulation of knowledge (Wildemuth and Freund 2012); thus, user search actions can be tentative. A fluid interface should provide easy transitions between user search attempts with low cognitive or mechanical costs. Query previews can serve this purpose, but existing applications are limited to previewing the number of relevant results to facets or highlighting relevant facet values under hover (Sect. 2). Advanced query preview techniques could be devised to

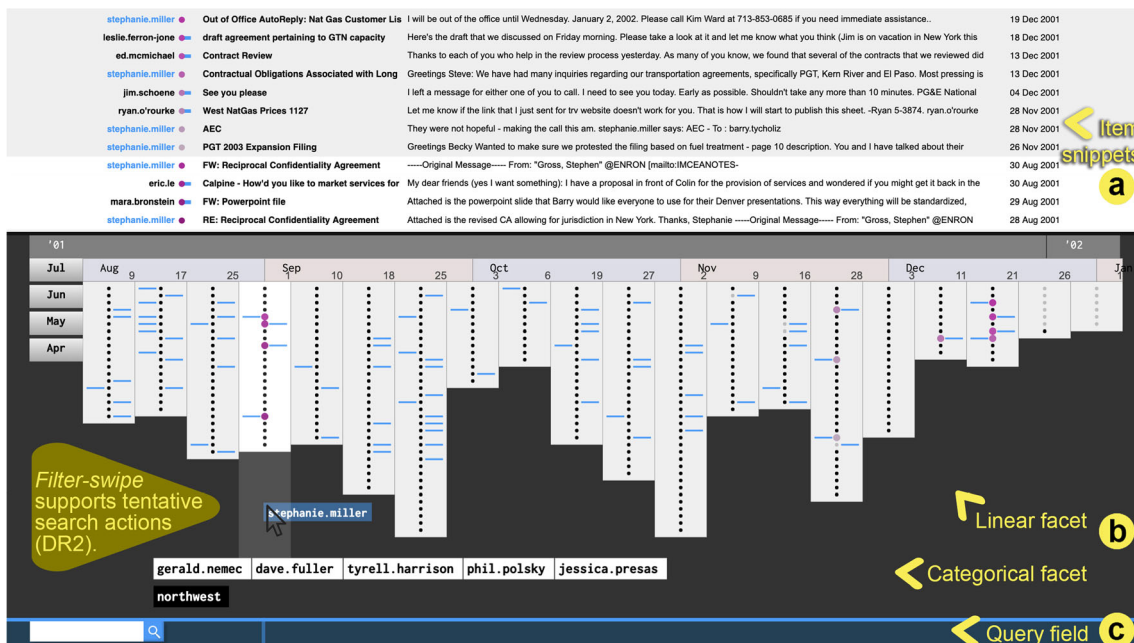
address this requirement better. Next, we present a case study devising IVF to address these two requirements.

#### 4 The IVF tool

We present a design of IVF addressing the two DRs to support fluid, exploratory search. Figure 1 shows an overview of the tool implemented with email data. The interface comprises an item snippet view, a facet space, and a query field. This spatial separation provides *multiple entry points for search*, including item detail exploration, faceted navigation, and query searches. The item snippet view shows snippets of selected items as a stack – newly selected items replace older ones to support skimming (Kules and Shneiderman 2008). In the case of emails, the snippet view contains the sender, title, the beginning of the email, and the date. Each row has a coloured dot that represents its freshness through saturation. More saturated colours indicate more recently selected items, which corresponds to the colour of the dots in the linear facet (Fig. 2b). Clicking on a snippet item opens a new window that shows its details.

The facet space coordinates a linear facet (time in this case) and a categorical facet, such as contact and keyword suggestions. The linear facet area shows the distribution of items over time. A queue of dots in each bar represents a list of items in a specified time period. The range of time for each bar is adjusted automatically according to the number of items over that period, avoiding item overflow. The linear facet view is scalable to display varied numbers of items from the cohort, as it arranges linear spans into pages once it cannot hold all items on one page.

The categorical facet, in the case of emails, consists of keyword and contact suggestions. We extract keywords within the body and subject fields of emails by computing the term frequency-inverse document frequency (TF-IDF) (Ramos 2003) and selecting the highest scoring keywords. The keywords are generated by combining words within the same noun phrases detected using dependency parsing (The Stanford Natural Language Processing Group 2014). Contacts are extracted from the sender, recipient, and carbon-copy field.



**Fig. 1** Overview of the IVF tool visualising 400 emails: **a** an email snippet display, **b** time as a linear facet and contact and keyword suggestions as categorical facets, and **c** a query field. A contact, “stephanie.miller”, is dragged onto a linear facet bar, so items in the intersection of the two facet values are selected and shown in the snippet view and the categorical facets suggest entities relating to the highlighted items. In this case, highlighted items are emails in this bar that were sent by or carbon-copied to this contact. The snippet display highlights the relevant emails with a white background colour. The categorical facet area suggests contacts and keywords dedicated to these items. If users find the categorical facet value relevant, they can drop it into the query field to filter the items; otherwise, they can just release it to end this search attempt. This action is called filter-swipe and supports tentative search actions (DR2). Data from the publicly available Enron email corpus (Cohen 2015) are shown

The query field allows for typed queries and dragging a categorical facet value as a query filter. A typed query is broken down into individual words if it is composed of more than one word. Multiple queries are joined by AND relations, i.e. only text items containing all of the queried words qualify to be displayed. Users can click on a queried word shown next to the input box to remove the query. When a query is selected or deselected, the tool will update the facets to describe a set of items satisfying the existing queries. In this way, the facets *guide users toward the next possible searches* and make sure user-dragged queries will not result in an empty result set, i.e. *avoid dead ends*.

#### 4.1 Faceted exploration within context

Initially, the most recent items in the timeline view are selected by default. The number of selected items depends on the height of the snippet display, avoiding scrolling. Users can select items by clicking on the dots in the linear facet or scrolling to sequentially select items in the linear bar. The selected dots are shown in purple with varying saturations; more saturated colours represent more recently selected items (Fig. 2b), which corresponds to the dots in the snippet rows. For the older items that have been pushed out of the stack of the snippet display, dots in the linear facet will turn grey to indicate they have been viewed (Fig. 2c). Visual encodings of the dots indicate the states of items (i.e. if they are selected, have not or have been selected), and *inform users about the exploration process (time-related context of DRI)*.

Mousing-over can trigger the interactions between the linear and categorical facets to support exploration in space-related context. Initially, the categorical facet, contact and keyword suggestions in the case of emails, show the summarisation of the items in the current span of the linear facet. Hovering over a potential query, which can be a contact or a keyword, highlights the relevant dots in the linear facet and the snippet view with various visual encodings (e.g. Fig. 3a). A line affords a contact relation where the sender and co-recipient lines adhere to the general chronological direction: The left-side line indicates a sender relation, and the right-side line denotes a co-recipient relation (Fig. 2d). A circle around a dot depicts a keyword relation (Fig. 2e). This puts a categorical facet value, such as a topic, *in the context of the linear facet (DRI)*, such as time. Users can see, for instance, how the frequency of the topic varies in time through the topic distribution over a timeline. Hovering over a linear facet bar, indicated by the white background of the bar, converts contact and keyword facets to summarise the items in the bar (e.g. Fig. 3b). The dynamism of the summarisations enables users to *assess the next relevant queries in the context of the overall content (DRI)*. For example, it allows users to detect special topics in a certain time bar among the topics over various time bars. If dots are selected in the highlighted bar (i.e. dots in purple), the items of the selected dots in the snippet view are highlighted with a white background colour. Mousing over the row of a dot is indicated by a bright blue background, as is the corresponding item in the snippet view.

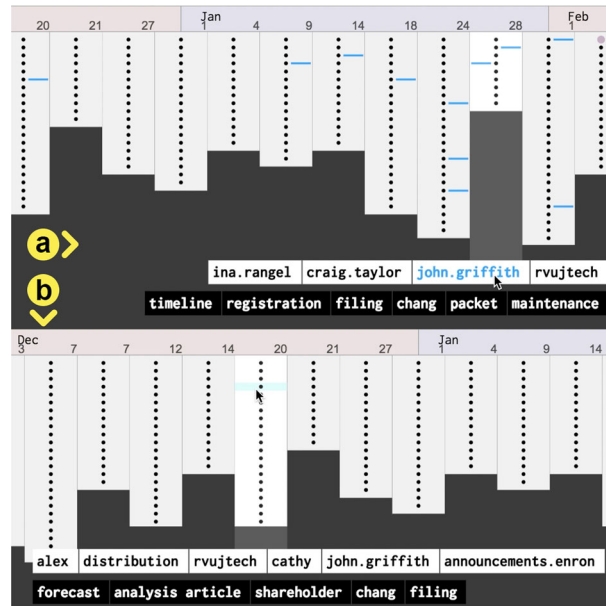
#### 4.2 Rapid transitions between search criteria

To support rapid transitions between search attempts (DR2), we propose to use facet values to select items for inspection without filtering the item space. Users can click on a categorical facet value for quick item selection. Selected items will be colour-coded in the linear facet, i.e. purple dots with varied levels of saturation, and shown in the snippet view. If this evokes too many items to fit into the snippet display, more recent items will be selected. This design allows users to flexibly inspect items of interest from various perspectives while maintaining the present exploration context.

The concept of using facets to select items without filtering the item space is further reflected in the design of a novel filter-swipe technique that enables dynamic query previews by flexibly combining two facet values. Users can drag a categorical facet value over linear facet bars sequentially, and as a result, the categorical facet will dynamically show facet values relating to the items in the intersection of the current linear bar and the dragged facet value, and the snippet view will display selected items in the intersection



**Fig. 2** Various states of the dots inform users about the exploration context (DR1). Each dot represents an item in the collection. **a** Dots that have not been selected yet; **b** dots under selection. The more saturated colour represents the more recently selected items; **c** dots that have been selected and pushed out of the stack of the snippet view; **d** and **e** the blue marks indicate relationships that can be visually overlaid with the other three states. They depict item relations to the currently focused categorical facet value; in the case of emails, **d** affords contact relations where the left-side line indicates a sender relation and the right-side line denotes a co-recipient relation; **e** denotes keyword relations



**Fig. 3** Interactions between linear and categorical facets supporting space-related context of DR1. **a** Mousing over a categorical facet value highlights the distribution of relevant items in the linear facet with various visual encodings (Fig. 2d, e); **b** Mousing over a linear facet bar transforms the categorical facet to summarise the items in the bar

(Fig. 1). If users find the dragged facet value relevant, they can put it in the query field to filter the item space. Otherwise, they can just release the object to quit this tentative search action and grab another categorical facet value to filter-swipe. The mouse interactions of the filter-swipe technique also trigger the mouse-over and click effects on the categorical facet value. That is, the relevant dots in the linear facet are highlighted with various visual encodings and the most recent items are selected when users proceed to drag the facet value. Overall, the filter-swipe technique enables flexible combinations of a linear and a categorical facet value for item inspection and features *rapid transitions between search attempts (DR2)*. See a video demonstration of the IVF tool at <https://youtu.be/v0tUAXPjqfg>.

## 5 Use cases

To demonstrate how the features of the tool address the DRs, we implemented the tool in three different applications (Table 1). The first case visualises emails for email finding, the second case uses tweets for serendipitous discovery and the third case involves a study of acute myeloid leukemia (AML) to recognise mutation co-occurrence patterns across patient ages. The analysis focuses on how the features of the tool fluidly support the critical aspects of exploratory search (i.e. learning, investigation, and forming new targets for search) (White and Roth 2009; Wildemuth and Freund 2012).

### 5.1 Email finding

Email, as a daily communication method, requires easy and efficient management. Looking for an email from a large collection of emails is a common yet burdensome task. To demonstrate how the features of the tool support email finding, we loaded 2,499 emails from the Enron corpus (Cohen 2015) into the IVF tool.

**Table 1** An overview of use cases

Case	Linear facet	Categorical facet	No. of items	Data source
Email finding	Time	Sender & co-recipient, keyword	2499	The Enron corpus (Cohen 2015)
Serendipitous tweet discovery	Time	Username, keyword	1226	Aweiand (Aweiand 2014)
Recognition of age-related oncogene co-occurrences	Age	Mutated gene	198	cBioPortal (Gao et al. 2013)

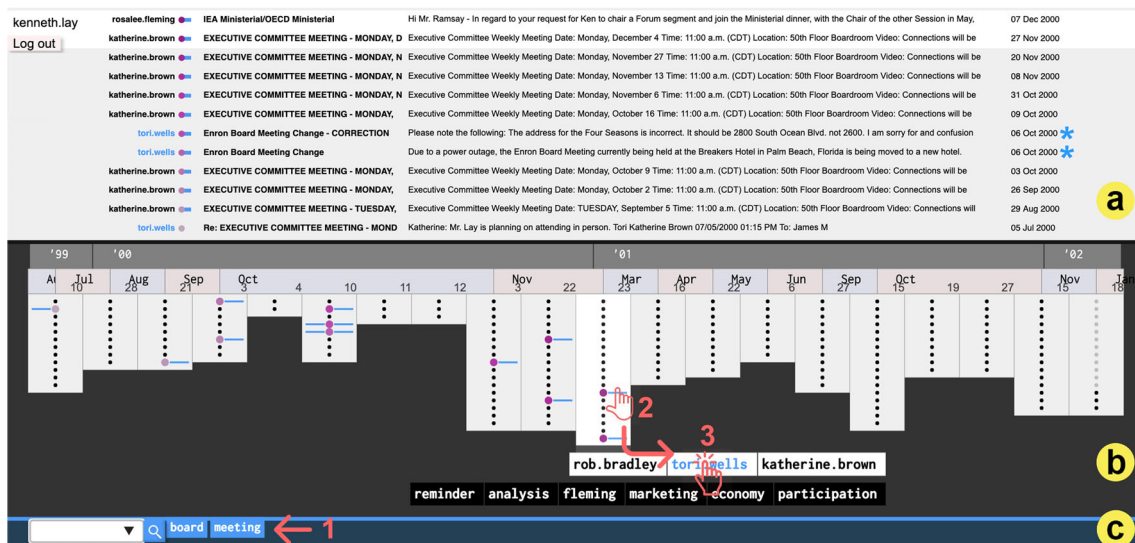
See Sect. 4 for a detailed explanation of how the tool functions. We discuss how the IVF tool could help email finding through the following scenario:

Kenneth opens the IVF tool with all his emails and wants to find an email about the address of a board meeting happened around early 2001 as he remembers. He types “board meeting” in the query box to filter the emails. The resulting emails range from late 1999 to early 2002 (Fig. 4). He hovers his mouse over the columns around early 2001 and a suggested contact “tori.wells” reminds him of this person’s appearance in the meeting. Then he clicked on the contact to quickly select relevant emails. The emails with the address information of the board meeting appeared in the snippet view, which turn out to come in October 2000 (emails marked with \* in Fig. 4). In this case, navigating through the linear facet helps the discovery of a familiar contact and further leads to the right email using quick item selection.

## 5.2 Serendipitous tweet discovery

Every day, countless messages are posted to Twitter. One challenge related to using this service is how to avoid missing interesting posts amid the noise of user-generated content. To address this, we incorporated 1,226 tweets retrieved from Aweiand (2014) and ranging from October 15, 2011, to October 20, 2011, into the tool (Fig. 5a). The snippet view lists the selected tweets including usernames, snippets of the tweets, and dates. The linear facet shows the distribution of tweets over the timeline, whereas the categorical facet provides an overview of relevant users and keywords extracted using the same technique employed to analyse email data, i.e. TF-IDF. We discuss how the tool’s features can be used to support rapid skimming and serendipitous tweet discovery through a scenario.

Ann has not logged into Twitter for five days. After logging in on the sixth day using the IVF tool, Ann gets 1,226 unread tweets. She turns over the pages of the timeline, browsing through the suggested keywords, such as “Android”, “keyboard”, and “upgrade”, all of which seem to be related to technology. These are not so interesting to Ann, so she hovers her mouse over each timeline bar to look for something interesting in smaller collections. Some keywords pop up through the sequential hovering actions, such as “Mandarin”. She then puts the mouse over the keyword “Mandarin” and two relevant items in two timeline bars are highlighted. She then drags the keyword over the two timeline bars sequentially to check the corresponding tweets in the snippet view (Fig. 5a). From one tweet, she learns that a friend is learning Mandarin to prepare for his upcoming trip to Hong Kong, whereas the other tweet seems related to iOS5, an operating system used for mobile devices. She then releases the keyword and continues sequentially

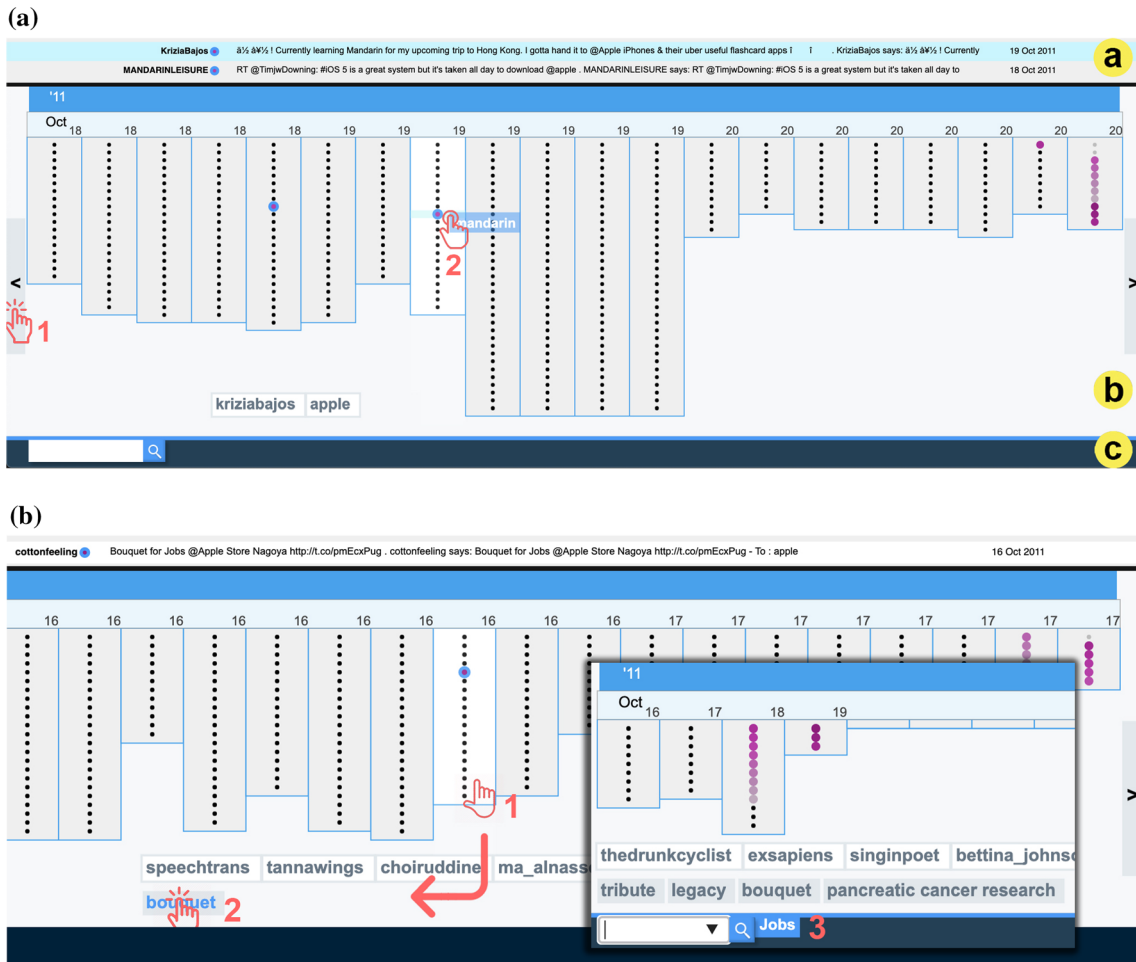


**Fig. 4** The IVF tool loaded with 2499 emails for email finding: **a** a snippet view with each row showing one selected email represented by the sender, the title, the beginning of the content, and the date; **b** a facet space with time as a linear facet and contacts and keywords as categorical facets; **c** a query field that supports typed queries as well as dragging categorical facet values as queries. After filtering the emails by queries of “board” and “meeting” (1), Kenneth clicks on the contact “tori.wells” (3) accessed through hovering over the first column of the year 2001 (2) for quick item selection. Selected emails are shown in reverse chronological order in the snippet view and represented by purple dots. Emails marked with \* are the target emails



hovering over the timeline bars. Another keyword, “bouquet”, catches her attention. She clicks on the keyword and sees that the selected tweet in the snippet view says, “bouquet for Jobs at \*\*\* apple store.” (Fig. 5b) She remembers this person just passed away because of some kind of disease. She then inputs his name, “Jobs”, in the query box to check all messages related to him, which results in 32 tweets in the view (Fig. 5b). The updated keyword suggestions include “pancreatic cancer research”. As a result, she realises he passed away due to pancreatic cancer. She then scrolls through the dots in the timeline to sequentially select and rapidly skim through their snippets.

During this exploration, the user utilises the context to explore keywords, uses keywords of interest for quick item inspection, and discovers new information. When she hovers over timeline bars, some interesting keywords pop out from the context facilitating learning and investigation (DR1). Quick item selection without filtering the space (DR2) helps her formulate new targets to explore, such as “Jobs”, and acquire further information.



**Fig. 5** A serendipitous tweet discovery scenario. (a) The IVF tool is used here to visualise 1,226 tweets for serendipitous discovery: a) a tweet snippet view in which each row presents a username, a tweet snippet, and a date; b) a facet space that includes a timeline and user and keyword suggestions; c) a query field. To begin with, Ann turns over the pages of the timeline by clicking on the arrow on the left side (1) and browses through the suggested users and keywords. Then Ann drags a keyword of interest, “Mandarin” (2), onto a timeline bar to inspect relevant tweets (filter-swipe). As a result, two tweets containing this keyword are selected and displayed in the snippet view. Their corresponding dots are surrounded with blue circles indicating keyword relations. A bright blue background colour highlights the tweet being hovered over. The categorical facet suggests usernames related to the content in the intersection of the current timeline bar and the dragged keyword, which is the tweet being hovered over in this case. No keyword suggestions in this case based on the algorithm. (b) Continuing browsing, Ann then clicks on another keyword of interest (2), “bouquet”, accessed through mousing over a timeline bar (1) to select relevant tweets. A selected tweet, “bouquet for Jobs at \*\*\* apple store”, reminds her of a recent event for further investigation. After filtering tweets by a typed query, “Jobs” (3), Ann gets what she is looking for through a keyword suggestion, “pancreatic cancer research”

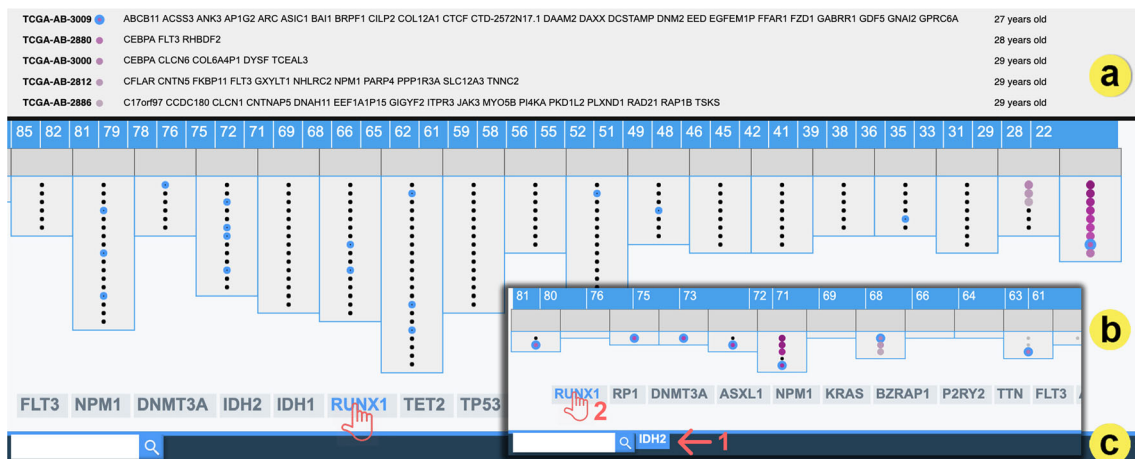
### 5.3 Recognition of age-related oncogene co-occurrences

Cancer is a disease driven by the so-called “driver” mutations that affect cell’s tumorigenic properties. Normally, a cancer cell has multiple mutations. Different mutation combinations can affect the response of a given treatment differently. On the other hand, ageing is associated with the accumulation of mutations and can increase the risk of developing cancer. In this use case, we incorporated a study of AML from cBioPortal (Gao et al. 2013) to investigate the age-related oncogene co-occurrence patterns. The study involved 198 samples from 198 patients with an age range from 18 to 88 years old. On average, each patient had 12.87 mutated genes ( $SD = 7.40$ ).

As Fig. 6 shows, the linear facet depicts the sample distribution across patient ages, whereas the categorical facet displays the most frequent mutated genes across samples ordered from left to right. The more frequent the mutated gene, the greater the likelihood that the gene is an oncogene. The item snippet view lists selected patient data, including patient identity numbers, lists of mutated genes, and ages. We demonstrate the use of the IVF tool to support simple data analysis through the following scenario.

A bioinformatician is exploring oncogene co-occurrence patterns across AML patients of various ages. Initially, the interface suggests a list of the mutated genes occurring most frequently in the cohort (Fig. 6). Hovering over the genes, he sees that some of them are distributed in a biased manner across patient ages. For example, genes *IDH2* and *RUNX1* tend to appear in patients over 55 years old. The bioinformatician knows that *IDH1/2* mutations are often pre-leukemic mutations that require other mutations for the disease to progress, but he does not know their relation to patient ages. Therefore, he conducts a PubMed search to find this information. Interestingly, he finds a paper indicating that *IDH2* mutations are associated with older age through an analysis of 805 adults with AML (age range: 16 to 60 years) (Paschka et al. 2010). Then, he starts to explore mutation co-occurrence patterns of *IDH2*. He drags gene *IDH2* into the query field, which results in a remainder of 20 samples in the view (Fig. 6). The categorical facet shows that the most frequent mutated gene in the current dataset is *RUNX1*, so he performs another online search on these two mutated genes in AML patients. A paper shows that *RUNX1* mutations are indeed associated with older age and are correlated with inferior prognosis, whereas patients with *RUNX1* and *IDH2* co-mutations experience a relatively better outcome (Gaidzik et al. 2016).

In this case, the IVF tool provides simple data analysis functions to arouse deeper investigation and support the discovery of new knowledge. The distribution of categorical facet values over the linear facet provides the context to discover distributional-biased entities (DR1). The categorical facet facilitates drilling down into the dataset, as well as entity co-occurrence pattern discovery, to stimulate in-depth investigation.



**Fig. 6** The IVF tool presents 198 AML patient mutation records for age-related oncogene co-occurrence recognition: **a** a patient record snippet view in which each row represents one patient record with a patient identity number, a snippet of mutated gene information, and an age number; **b** a facet space with age as a linear facet and most frequent mutated genes as a categorical facet; **c** a query field that supports typed queries as well as dragging genes as queries. After loading the data, the bioinformatician mouses over suggested genes and discovers that some genes, such as *RUNX1* and *IDH2* distribute biasedly toward patients over 60 years old. He then drags the mutated gene *IDH2* to the query field for filtering (1). The categorical facet shows the most frequent mutated genes among the current subset in order from left to right. The mutated gene that occurs most frequently with *IDH2* in this dataset is *RUNX1* which is under hover (2). As a result, the dots with blue circles indicate the samples with the co-occurring mutated genes *IDH2* and *RUNX1*

## 6 User studies

We conducted two user studies with email data to understand 1) the search performance of the IVF tool compared with the traditional query search interface and 2) how the features of the tool are used in realistic search scenarios.

### 6.1 Comparative study

We evaluated the search performance of the IVF tool by comparing with a baseline system through a within-subjects design. We asked users to use both systems to find target emails with varied search difficulties. The baseline system mimics the typical query search interface with the linear and categorical facets replaced by a longer list of email snippets while maintaining the same free text search facility as the IVF tool (Fig. 7).

#### 6.1.1 Data

We used and preferred the Enron email corpus (Cohen 2015) over participants' personal email collections to have the same set of questions for the participants and to ensure complete control of factors, such as the amount of information known about the target email and the last time the email was viewed. The Enron email corpus provides an extensive collection of real emails from various users. We opted for a subset of the emails received by two of Enron's important managerial officers with a similar number of emails (2500 and 2142). Each email set was used in only one of the two systems (the IVF tool or baseline).

#### 6.1.2 Participants

We recruited 16 participants through advertising at the university. All participants (nine females and seven males, age range: 22-42, age mean: 25.8) were university students from diverse areas, including Computer Science, Psychology, Cognitive Science, Nursing, Chemistry, and Linguistics. Half of the participants stated that they often use Gmail, while the rest mentioned OS X Mail, Outlook 365, and the university's Webmail. None of the participants had ever heard of the Enron email collection. Each had two movie tickets as compensation for the participation.

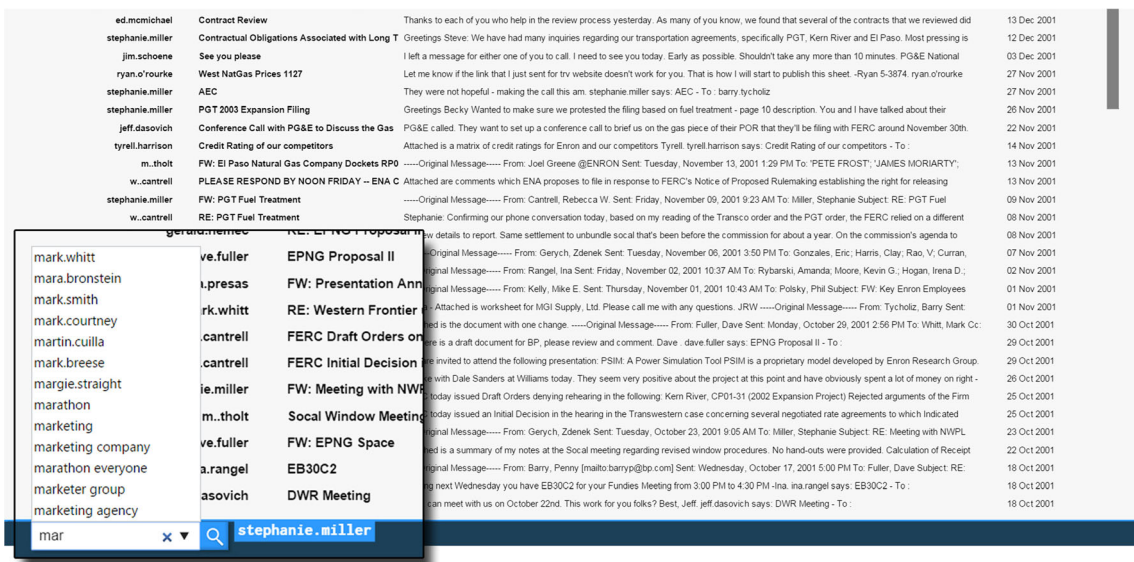


Fig. 7 The baseline system displays 34 emails instead of 12 in the IVF tool and enables continuous scrolling through 400 emails (the same number visualised in the linear facet of the IVF tool). Categorical facet values could be accessed through the query box via auto-complete. Participants could also type in the query box as “dec 2011” to navigate in time

### 6.1.3 Tasks

We evaluated each system on three types of tasks that differed in difficulty levels based on the amount of information provided about the emails to be found (Table 2). T1 was easy as specific unique keywords about the email are known, and a search query with these keywords would return no more than ten emails. T2 and T3 were harder than T1 as a search query with any of the words provided would return (if any) around 30-50 emails, i.e. none of the keywords uniquely identify the target email. All the information required to identify the correct email in all tasks was visible in the email snippets (without the need to open and read the entire email).

In T1 and T2, the month and year when the email was sent were known, but in T3, a 2-month range was provided. The sender was known in T2, but in T3, two possible senders were provided. We indicated whether the mentioned sender was the first or last name of the person or whether it was the name as it appeared in the contacts list. Based on the tasks we collected from our second study, which comes later in this section, we thought these tasks cover realistic email finding tasks while ensuring that the participants were provided enough information to find the correct email.

With the within-subjects design, we counterbalanced the order of both the two sets of emails and the two systems, so there were four possible orderings. The order of the tasks was fixed from easy to hard. Each task level for each dataset contained two questions, the order of which was randomised.

### 6.1.4 Procedure

For each system, the procedure consisted of a training, a practice, and an actual task session. Pilot studies were conducted to ensure the viability of the procedure. The training session comprised: (1) a live demonstration of the system and its search facilities; (2) three questions from easy to hard, over a prepared dataset with 530 emails in the 'inbox' directory of Enron's vice president, Barry Tycholiz, which the participants had to complete. After training, the participants read about the Enron employee they had to impersonate and then carried out three practice questions in the order of easy to hard. During the training and practice sessions, participants could ask questions to the experimenter to resolve any confusion. The actual task contained six questions as mentioned earlier (two for each task level; easy to hard). We set the time limit for each question as 5 min. In total, the experiment took an hour on average.

The experiment for both systems was run on a 3.40 GHz quad-core PC with 16 GB RAM using a 21" monitor with 1920 \* 1200 resolution. For the IVF tool, 12 emails were visible in the list of email snippets, and 400 emails were visible in the timeline visualisation in one instance. For the baseline, 34 emails were visible in the list of email snippets in one instance, with the possibility to see 400 emails through scrolling (Fig. 7).

### 6.1.5 Hypotheses


To evaluate task performance, we measured the time taken to find the email requested by each task question and whether the correct email was successfully found or not (success was 0 if the email was not found within 5 min). Considering the interaction the IVF tool provided was less familiar to the participants than the typical query interface, we hypothesised that *for the three task types, the task performance of using the IVF tool is comparable to the baseline system.*

### 6.1.6 Results

For each participant, we computed the mean success and time for each task type and interface. Figure 8 shows the box plots of success and time per interface and task. As the distributions of the time and success data were not normal, we used non-parametric tests for the analysis. Wilcoxon Signed-rank tests showed the

**Table 2** Example tasks in three difficulty levels based on the amount of information provided about the email to be found

Task	Difficulty	Example
T1	Easy	Find the email received in September 2001 about the cycling club meeting.
T2	Moderate	Find the first email sent by someone whose first name is Terrie in November 2001 about an Enron reward.
T3	Hard	Find the email sent by david.morris OR someone whose first name is Sherri, around April 2001 and May 2001, asking a colleague a possible meeting place and time.



**Fig. 8** Success (left) and time (right) per task type

**Table 3** Effect sizes and  $p$  values of the interfaces' effects on task performance regarding three task types and how different task types affected task performance with the two interfaces

	Interface			Task	
	T1	T2	T3	IVF	Baseline
Success	$r = 0.35$ , $p = 0.35$	$r = 0.11$ , $p = 0.61$	$r = 0.33$ , $p = 0.18$	$W = 0.34$ , $p = 0.005$	$W = 0.26$ , $p = 0.016$
Time	$r = 0.19$ , $p = 0.45$	$r = 0.24$ , $p = 0.35$	$r = 0.01$ , $p = 0.98$	$W = 0.85$ , $p < 0.001$	$W = 0.78$ , $p < 0.001$

two interfaces did not exhibit statistically significant differences in task performance regarding the three task types, whereas Friedman tests indicated that the three task types impacted task performance significantly for both interfaces (Table 3).

Additionally, we investigated the learning curve of the IVF tool by analysing whether the IVF tool significantly improved task performance on T3 relating to T1, compared with the baseline. To do this, we calculated the time and success differences between the two interfaces on T1 and T3 and performed Wilcoxon Signed-rank tests between the differences. No significant result was found, which evidenced that the IVF tool did not have a steep learning curve and was easy to use.

Thus, we confirm the hypothesis and conclude that the IVF tools appeared to be easy to learn and have a comparable performance compared with the query search interface.

## 6.2 Exploratory study

To understand *how people use the IVF tool's features in practical exploratory search scenarios*, we conducted a user study involving realistic email search tasks. Locating a specific email in a personal email box with limited memory cues is difficult. Research shows that memories may be organised by episodes, such as the location and relative time of an event (Elsweiler et al. 2008). We pondered that user interaction with the tool's facets could assist users' memories and widen search opportunities. In this study, we investigated how users interacted with the coordinated facets to locate a specific email by analysing user interaction data and users' perception of how helpful facets were for email finding through questionnaires.

### 6.2.1 Participants

We recruited 11 participants from two large universities (6 graduate students, 4 postdoctoral researchers, and 1 administrative staff). A pre-screening questionnaire was administered to ensure that the participants performed email searches at least three times a week and that their email accounts contained more than 300 emails. Four participated the study with their personal email accounts (which they also used for work and study related correspondence) and seven participated with their institutional accounts. All participants were fluent in English and most of their email correspondence was in English. Each participant received two movie tickets as a reward for the completion of the study.

### 6.2.2 Procedure

To address privacy issues related to personal data, we adopted a diary-keeping method from Elsweiler et al. (2008) to collect email search tasks. To increase the difficulty of the tasks, we carried out the email search experiment 90 days after the diary study was completed.

**Diary study** The participants were asked to keep an online diary for 30 days to record all instances in which they had to find an email. We defined the scope of finding as all types of search actions, including

typed queries. In total, we collected 127 diary entries (4–18 entries from each participant). We excluded 26 of these entries because (1) they were repeated entries; (2) they did not refer to a certain target email, but rather described cases in which the participant wanted to make sure that there were no emails with specific information; or (3) the target email was in the sent box. This left us with 101 entries, which we then used as email search tasks related to the participant’s inboxes.

**Lab experiment** The study started with a training session in which each participant completed 4 training tasks that involved finding emails in an unfamiliar inbox, an employee account from the Enron corpus (Cohen 2015), using various combinations of available information, such as sender, topic, co-recipient, and date. Then, we extracted emails within a recent 2-year span from participants’ inboxes to the IVF tool. Participants were asked to find the emails mentioned in their diaries using a desktop computer with a 24-inch monitor. The tasks were introduced in random order and in the form of task cards. Each task card included a diary entry the participant had previously written and a questionnaire. There were no time constraints for executing the tasks. The participants indicated task completion either by opening an email and confirming it as the correct email (success) or by clicking on the “Give up” button on the screen (failure). Participants could skip a task if the task definition written in the diary was too vague or if the specific email was not included in the participant’s inbox anymore. An experimenter was present during the study to make sure the procedure was followed and answer any technical questions.

### 6.2.3 Data collection

During the search tasks, we logged all user interactions and the state of the interface at those moments including filtering criteria, suggestions, and the number and distribution of emails in the timeline interface. To maintain privacy, our logs did not include any textual content from participants’ accounts.


To help clarify the log data gathered during each session, we administered a questionnaire at the end of each task. In the questionnaire, the participants were asked whether they found relevant contacts or keywords from the suggestions and whether the suggestions supplemented any of the information that was missing at the beginning of the search session.

### 6.2.4 Results

In total, we obtained 73 session logs. For seven of these tasks, participants remarked that they did not remember which emails their diary entries referred to. Of the remaining 66 tasks, 58 (88%) were completed successfully, and 8 (12%) were unsuccessful.

Of the 66 tasks, 23 tasks were query-only, i.e. they did not include any facet use. Thirty tasks were mixed sessions in which participants used queries in combination with facets. Most of these sessions (20/30) started with queries to reduce the number of emails in the display before participants proceeded to use facets. Thirteen tasks did not involve any use of typed queries. A majority of the tasks in this group (7/13) started with a timeline navigation action, such as clicking on a month, and then relied on suggestions or the timeline to select emails. Participant-wise, some were more oriented toward submitting typed queries, whereas others relied on timeline navigation and suggestions, but the difference in using query or facet was not statistically significant among the participants (Fig. 9). Search strategies also varied among tasks that belonged to the same participant, i.e. participants adapted their strategies to different tasks.

The timeline was used in 39 (59%) sessions (Table 4). In 17 (44%) of these sessions, specific timeline navigation actions, such as pointing to months and directly selecting dots, led to the selection of the correct email. Selecting emails through timeline navigation was rarely used at the tasks’ initial stages, but typically followed query actions or month-navigation actions. Suggestions were used in the majority of search



**Fig. 9** Some participants (C, F, K) were more oriented toward typed queries, whereas other participants (J, E) relied more on timeline navigation and search suggestions. A Friedman’s test showed the difference in using query or facets were not statistically significant among the participants ( $\chi^2(2) = 4.05, p = 0.13$ )

**Table 4** The usage of facets among the 66 search tasks

	Timeline (linear facet)	Suggestions (categorical facet)
Use	39 (59%) tasks	38 (58%) tasks
Success	17 (26%) tasks	23 (35%) tasks

sessions, 38 (58%) out of the 66 search sessions. In 23 (61%) of these sessions, suggestions led to the selection of the correct email either directly or when used as queries.

*Suggestions were generally used following data-specification actions, such as querying or navigating to a month (35/38).* A possible explanation is the lack of relevant suggestions at the start of a search session. The log data indeed shows that relevant suggestions (regarding the emails marked as correct by the users) were not encountered as often at the initial stages of the search compared to later stages, i.e. after participants made data-specification actions.

*The timeline can provide a context to support suggestion discovery.* Out of the 72 instances of suggestion usage, 32 (44%) of them were accessed through the timeline, i.e. participants accessed suggestions when hovering overtime periods. This reveals that the timeline can provide a context for suggestion exploration.

*Suggestions were used more frequently for selection than filtering.* Among 72 instances of suggestion use, suggestions were dragged to the query area for filtering only 11 (15%) times (Table 5). A Wilcoxon signed-rank test showed using suggestions for selection was statistically significantly more frequent than for filtering among the participants with a  $p$  value = 0.037 and an effect size  $r = 0.645$ . In other words, participants generally preferred to keep their current context to inspect items rather than decreasing the size of the item space. Additionally, we collected 13 instances of filter-swipe usage, a small portion (18%) of the 72 total instances of suggestion usage.

*The most commonly used suggestions were contacts.* Among the 72 instances in which suggestions were used, the overwhelming majority, 63 (88%), were contact suggestions (Table 5). A Wilcoxon signed rank test showed contacts were statistically significantly more often used than keywords by the participants with a  $p$  value = 0.006 and an effect size  $r = 0.847$ . Of the 11 instances in which suggestions were used as filters, all were contact suggestions. The subjective evaluation of the suggestions through the questionnaire (Table 6) shows that the contact suggestions were found relevant for 35 search tasks (53%) and supplemented missing information in 13 tasks (20%). For keyword suggestions, the respective figures were 18 (27%) and 11 (17%). A Wilcoxon signed-rank test showed participants tended to find contacts more relevant for email finding than keywords with a  $p$  value = 0.054 and an effect size  $r = 0.585$ , which echoed the earlier findings on email search queries that most queries referred to people and especially to senders (Harvey and Elsweiler 2012). However, in actual utility, participants perceived contacts and keywords equally supplemented missing information.

**Table 5** The usage of suggestions among the 72 instances

Feature	Direct selection	Drag filtering	Filter-swipe
Instances	48 (67%)	11 (15%)	13 (18%)
Type	Contact	Keyword	Contact vs. Keyword
Instances	63 (88%)	9 (13%)	$r = 0.847, p = 0.006$

**Table 6** Questionnaire results

	Found relevant	Supplemented missing information
Sender	30 tasks (45%)	8 tasks (12%)
Co-recipient	11 tasks (17%)	6 tasks (9%)
Keyword	18 tasks (27%)	11 tasks (17%)
Contact vs. Keyword	$r = 0.585, p = 0.054$	$r = 0.027, p = 0.943$

Contacts were found to be relevant and supplemented missing information more frequently than keyword suggestions. Note that the cases in which the suggestions supplemented missing information are a subset of the cases in which they were found relevant.

In summary, the 66 search cases demonstrate that the IVF tool facilitates email finding. A majority (65%) involved the use of facets to guide searches. Dynamic query suggestions through the timeline navigation could help discover relevant suggestions in the context (44% of suggestion usage, DR1) in which contact suggestions were more often used than keyword suggestions (effect size = 0.847,  $p = 0.006$ ). The design of using facet values to select items without filtering the item space (DR2) was favoured over using facet values as queries to filter the item space among the participants with an effect size = 0.645 and  $p = 0.037$ .

## 7 Discussion

The IVF tool, which coordinates a linear facet, a categorical facet, and result items, exemplifies the two DRs, enabling contextual information and rapid transitions between search criteria, to support fluid exploratory searches. Three use cases and two user studies revealed the usability and usefulness of the tool in supporting exploratory search with various datasets and for various exploration purposes. The cases show that, in the context of the categorical facet, users can discover keywords of interest for quick item selection and inspection; in the context of the linear facet, users can explore the distribution of categorical facet values, which stimulates deeper investigation. The tool exhibited comparable performance in search to a query interface. The novel design concept of using facet values to select items for inspection without filtering the item space was favoured over using facets as filters according to the exploratory study. Based on the findings, we discuss design implications to address the DRs further and inspire the design of fluid exploratory searches using IVF.

### 7.1 Design implications

#### 7.1.1 *Semantic zooming of the linear facet*

To improve user exploration of the contextual information, we could devise the tool to support semantic zooming of the linear facet, such as zooming in or out to visualise bars over the span of days, weeks, and months; users can aggregate the dots into semantically meaningful collections for pattern recognition. This allows the facet space to hold more items in one view and increase the flexibility of user exploration. In the patient records user case, because of the insufficient quantities of samples in various age bars, it is not possible to compare oncogene co-occurrence patterns among various ages. With semantic zooming, users can, for instance, gather samples into larger collections under larger age spans to support recognition of age-related oncogene co-occurrence patterns.

#### 7.1.2 *Rapid skimming of item snippets*

The filter-swipe technique requires users to depress the mouse button all the time during item inspection. In the comparative study, the participants had to carefully read the contents of the email to see if they matched the task description, and had to hold the mouse clicked over a long period of time, which can require extra effort on the part of users. Also reflecting on the Twitter use case, which required users to rapidly skim through item snippets, We suggest that the snippet view be further improved to support item skimming. For instance, relevant items can be made more salient by highlighting the words in the snippets that appear in the categorical facet. Supporting rapid skimming can increase the efficiency and precision of user assessment of the relevance of selected items. Meanwhile, it can also improve the usability of the filter-swipe technique.

#### 7.1.3 *Visual encoding of the item selection order*

To enable quick item selection without filtering the item space, the snippet view functions as a stack in which the most recent selections replace older ones. Although this permits flexible switching between selection modes, such as using a suggestion or directly selecting an item from the timeline, it also occasionally makes it harder for participants to recognise the emails that are related to their most recent selections. Currently, we encode item selection order using dot saturation, which is less efficient for visual perceptions compared to position encodings according to Mackinlay's ranking, in which visual variables affect how accurately humans perform the corresponding perceptual task for various data types (Mackinlay



1986). A simple solution would be to fix the position of the order of selected items. For example, using a rolling mechanism to always push the newest selected items to the top and move older items toward the bottom can help improve users' sense of item selection.

## 7.2 Limitations

### 7.2.1 Scalability

The IVF tool's linear facet can display scalable numbers of items shown in the use cases as it arranges linear spans into pages. As discussed earlier, semantic zooming of the linear facet that aggregates the dots into semantically meaningful glyphs can be incorporated to allow one page to hold more items. Regarding the scalability of facets, like most faceted visualisations, the tool is limited in terms of the dimensions of the metadata that it can process.

### 7.2.2 User study

**Internal validity** Other factors could affect the results of the studies, including how the features were presented in the training sessions and participants' personal preferences. For instance, some participants might prefer to stick to query search and be reluctant to use new techniques. Additionally, in the second study, participants had different tasks from each other, which might affect feature usage.

**Generalisability** The participants of the two user studies were limited to university affiliates, and the number was not large, which may hinder the generalisability of the results to a larger population. On the other hand, we are interested in whether the findings from the email studies can be generalised to other content such as tweets. More studies are required with other types of content to explore the results' generalisability.

## 8 Conclusion

This paper focuses on devising IVF to support fluid interactions for exploratory search. First, we reviewed existing faceted search interfaces by their interaction and visualisation design of facets. We then proposed the idea of IVF and derived two DRs that have not been well addressed to support fluid, exploratory searches. We exemplified the requirements by devising an IVF tool. The tool comprises an item snippet view, a facet space, and a query field to support multiple entry points for search. The facet space coordinates a linear and a categorical facet to represent the distribution and summarisation of a collection of items and *provide contextual information for faceted exploration (DR1)*. The novel design concept of using facets to select items without filtering the item space introduced by the tool facilitates *rapid transitions between search attempts (DR2)*. On the one hand, users can quick-select and inspect relevant items by clicking on a categorical facet value; on the other hand, users can employ a filter-swipe technique to dynamically preview results under the flexible combination of two facet values.

We then presented three use cases to demonstrate how the IVF tool supports fluid exploratory search in various scenarios. The first case, email finding, illustrated how users could find the contact of interest through the timeline navigation and use the contact form quick item inspection and discovery. The second case with tweets exemplified that dynamic summarisation using categorical facets suggested useful keywords for inspecting items of interest so that users could skim through large collections of text items without missing interesting tweets. The third case, which used AML patient mutation records demonstrated that the linear facet helped expose distributional-biased categorical values to support data analysis and the discovery of new knowledge for further investigation.

We compared the task performance of the IVF tool to the traditional query search interface and found evidence that the tool had comparable performance to the baseline and was easy to learn. We investigated the practical use of the IVF tool through a lab study with realistic email-finding tasks. Results show that, often, dynamic suggestions through the timeline navigation help users discover relevant search suggestions in the context (DR1). Among the participants, the design of facet-based item selection that does not filter the item space was favoured over using facet values as queries to filter data (DR2). The filter-swipe technique is not ideal for detailed item inspection but is useful for item skimming.

Based on these practices, we derived a set of design implications to address the requirements for fluid exploratory searches further. First, semantic zooming of the linear facet can further facilitate the recognition of data distribution. Second, keywords in result items can be highlighted to support rapid skimming. Third, using position alongside saturation to encode the order of results can improve recognition of recent results. Finally, facet scalability remains a challenge for visualisation. In future work, more participants of diverse backgrounds could help quantify the findings; studies on other types of content are needed to generalise the findings.

**Acknowledgements** We thank Heidi Sahlberg from the University of Helsinki for her help with the email user study. This research was supported by the Academy of Finland through the Strategic Research Council project Data Literacy for Responsible Decision-Making [Grant Number 335194].

**Funding** Open Access funding provided by University of Helsinki including Helsinki University Central Hospital.

#### Declarations

**Conflict of interest** The authors declare that they have no competing interests.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

#### References

- Andolina S, Klouche K, Ruotsalo T et al (2018) QueryTogether: enabling entity-centric exploration in multi-device collaborative search. *Inf Process Manag* 54(6):1182–1202. <https://doi.org/10.1016/j.ipm.2018.04.005>
- Athukorala K, Glowacka D, Jacucci G et al (2016) Is exploratory search different? A comparison of information search behavior for exploratory and lookup tasks. *J Assoc Inf Sci Technol* 67(11):2635–2651. <https://doi.org/10.1002/asi.23617>
- Aweiand (2014) Twitter sentiment analyses. <https://github.com/aweiand/TwitterSentiment>. Accessed in December 2020
- Chang JC, Hahn N, Perer A et al (2019) SearchLens: composing and capturing complex user interests for exploratory search. In: *Proceedings of the international conference on intelligent user interfaces*. ACM, pp 498–509
- Chantamunee S, Wong KW, Fung CC (2020) An exploration of user-facet interaction in collaborative-based personalized multiple facet selection. *Knowl Based Syst* 209(106):444. <https://doi.org/10.1016/j.knosys.2020.106444>
- Cohen WW (2015) Enron email dataset. <http://www.cs.cmu.edu/~enron>. Accessed in December 2020
- di Sciascio C, Sabol V, Veas EE (2016) Rank as you go: user-driven exploration of search results. In: *Proceedings of the international conference on intelligent user interfaces*. ACM, pp 118–129. <https://doi.org/10.1145/2856767.2856797>
- di Sciascio C, Sabol V, Veas EE (2017) Supporting exploratory search with a visual user-driven approach. *ACM Trans Interact Intell Syst* 7(4):18:1-18:35. <https://doi.org/10.1145/3009976>
- di Sciascio C, Brusilovsky P, Veas EE (2018) A study on user-controllable social exploratory search. In: *Proceedings of the international conference on intelligent user interfaces*. ACM, pp 353–364. <https://doi.org/10.1145/3172944.3172986>
- di Sciascio C, Brusilovsky P, Trattner C et al (2020) A roadmap to user-controllable social exploratory search. *ACM Trans Interact Intell Syst* 10(1):8:1-8:38. <https://doi.org/10.1145/3241382>
- di Sciascio C, Veas EE, Barria-Pineda J et al (2020b) Understanding the effects of control and transparency in searching as learning. In: *Proceedings of the international conference on intelligent user interfaces*. ACM, pp 498–509. <https://doi.org/10.1145/3377325.3377524>
- Dörk M, Carpendale S, Collins C et al (2008) VisGets: coordinated visualizations for web-based information exploration and discovery. *IEEE Trans Vis Comput Graph* 14(6):1205–1212
- Dörk M, Gruen D, Williamson C et al (2010) A visual backchannel for large-scale events. *IEEE Trans Vis Comput Graph* 16(6):1129–1138
- Dörk M, Carpendale S, Williamson C (2012a) Fluid views: a zoomable search environment. In: *Proceedings of the international working conference on advanced visual interfaces*. ACM, pp 233–240
- Dörk M, Riche NH, Ramos G et al (2012) PivotPaths: strolling through faceted information spaces. *IEEE Trans Vis Comput Graph* 18(12):2709–2718
- Elmqvist N, Vande Moere A, Jetter HC et al (2011) Fluid interaction for information visualization. *Inf Vis* 10(4):327–340. <https://doi.org/10.1177/1473871611413180>
- Elsweiler D, Baillie M, Ruthven I (2008) Exploring memory in email refinding. *ACM Trans Inf Syst (TOIS)* 26(4):21:1-21:36. <https://doi.org/10.1145/1402256.1402260>
- Gaidzik VI, Teleanu V, Papaemmanuil E et al (2016) RUNX1 mutations in acute myeloid leukemia are associated with distinct clinico-pathologic and genetic features. *Leukemia* 30(11):2160–2168. <https://doi.org/10.1038/leu.2016.126>

- Gao J, Aksoy BA, Dogrusoz U et al (2013) Integrative analysis of complex cancer genomics and clinical profiles using the cBioPortal. *Sci Signal* 6(269):1–11. <https://doi.org/10.1126/scisignal.2004088>
- Grierson HJ, Corney JR, Hatcher GD (2015) Using visual representations for the searching and browsing of large, complex, multimedia data sets. *Int J Inf Manag* 35(2):244–252. <https://doi.org/10.1016/j.ijinfomgt.2014.12.003>
- Harvey M, Elswailer D (2012) Exploring query patterns in email search. In: Proceedings of the European conference on advances in information retrieval. Springer, Berlin, pp 25–36. [https://doi.org/10.1007/978-3-642-28997-2\\_3](https://doi.org/10.1007/978-3-642-28997-2_3)
- He C, Micallef L, Serim B et al (2021) Interactive visual facets to support fluid exploratory search. In: The 14th international symposium on visual information communication and interaction. ACM, pp 15:1–15:10. <https://doi.org/10.1145/3481549.3481565>
- Hearst M, Elliott A, English J et al (2002) Finding the flow in web site search. *Commun ACM* 45(9):42–49
- Jankun-Kelly TJ, Ma K, Gertz M (2007) A model and framework for visualization exploration. *IEEE Trans Vis Comput Graph* 13(2):357–369. <https://doi.org/10.1109/TVCG.2007.28>
- Klouche K, Ruotsalo T, Cabral D et al (2015) Designing for exploratory search on touch devices. In: Proceedings of the SIGCHI conference on human factors in computing systems. ACM, pp 4189–4198. <https://doi.org/10.1145/2702123.2702489>
- Klouche K, Ruotsalo T, Micallef L et al (2017) Visual re-ranking for multi-aspect information retrieval. In: Proceedings of the conference on conference human information interaction and retrieval. ACM, pp 57–66. <https://doi.org/10.1145/3020165.3020174>
- Klouche K, Ruotsalo T, Jacucci G (2018) From hyperlinks to hypercues: entity-based affordances for fluid information exploration. In: Proceedings of the conference on designing interactive systems. ACM, pp 401–411. <https://doi.org/10.1145/3196709.3196775>
- Kreutz CK, Boesten P, Witry A et al (2018) FacetSearch: a faceted information search and exploration prototype. In: Proceedings of the conference “Lernen, Wissen, Daten, Analysen”, vol 2191. CEUR-WS.org, pp 215–226
- Kules B, Shneiderman B (2008) Users can change their web search tactics: design guidelines for categorized overviews. *Inf Process Manag* 44(2):463–484
- Mackinlay J (1986) Automating the design of graphical presentations of relational information. *ACM Trans Graph (TOG)* 5(2):110–141
- Mahdi MN, Ahmad AR, Subhi MA et al (2020) Visualization in faceted search engine—a review. In: Proceedings of the conference on big data and analytics. IEEE, pp 84–89
- Manioudakis K, Tzitzikas Y (2020) Faceted search with object ranking and answer size constraints. *ACM Trans Inf Syst* 39(1):9:1-9:33. <https://doi.org/10.1145/3425603>
- Marchionini G (2006) Exploratory search: from finding to understanding. *Commun ACM* 49(4):41–46. <https://doi.org/10.1145/1121949.1121979>
- Mauro N, Ardissono L, Lucenteforte M (2020) Faceted search of heterogeneous geographic information for dynamic map projection. *Inf Process Manag* 57(4):102257. <https://doi.org/10.1016/j.ipm.2020.102257>
- Niu X, Hemminger BM (2015) Analyzing the interaction patterns in a faceted search interface. *J Assoc Inf Sci Technol* 66(5):1030–1047. <https://doi.org/10.1002/asi.23227>
- Niu X, Fan X, Zhang T (2019) Understanding faceted search from data science and human factor perspectives. *ACM Trans Inf Syst* 37(2):14:1-14:27. <https://doi.org/10.1145/3284101>
- Paschka P, Schlenk RF, Gaidzik VI et al (2010) IDH1 and IDH2 mutations are frequent genetic alterations in acute myeloid leukemia and confer adverse prognosis in cytogenetically normal acute myeloid leukemia with NPM1 mutation without FLT3 internal tandem duplication. *J Clin Oncol* 28(22):3636–3643. <https://doi.org/10.1200/JCO.2010.28.3762>
- Plaisant C, Shneiderman B, Doan K et al (1999) Interface and data architecture for query preview in networked information systems. *ACM Trans Inf Syst* 17(3):320–341. <https://doi.org/10.1145/314516.314522>
- Qu D, Lin X, Ren K et al (2020) Airexplorer: visual exploration of air quality data based on time-series querying. *J Vis* 23(6):1129–1145. <https://doi.org/10.1007/s12650-020-00683-6>
- Qvarfordt P, Golovchinsky G, Dunnigan T et al (2013) Looking ahead: query preview in exploratory search. In: Proceedings of the SIGIR conference on research and development in information retrieval. ACM, pp 243–252. <https://doi.org/10.1145/2484028.2484084>
- Ramos J (2003) Using TF-IDF to determine word relevance in document queries. In: Proceedings of the instructional conference on machine learning, vol 242. Citeseer, pp 29–48
- Rauch M, Klieber W, Wozelka R et al (2015) Knowminer search—a multi-visualisation collaborative approach to search result analysis. In: International conference on information visualisation, pp 379–385. <https://doi.org/10.1109/iV.2015.72>
- Ruotsalo T, Peltonen J, Eugster MJA et al (2018) Interactive intent modeling for exploratory search. *ACM Trans Inf Syst* 36(4):44:1-44:46. <https://doi.org/10.1145/3231593>
- Ruotsalo T, Jacucci G, Kaski S (2020) Interactive faceted query suggestion for exploratory search: whole-session effectiveness and interaction engagement. *J Assoc Inf Sci Technol* 71(7):742–756. <https://doi.org/10.1002/asi.24304>
- Seifert C, Jurgovsky J, Granitzer M (2014) FacetScape: a visualization for exploring the search space. In: Proceedings of the international conference on information visualization. IEEE, pp 94–101. <https://doi.org/10.1109/IV.2014.49>
- Shneiderman B, Byrd D, Croft WB (1997) Clarifying search: a user-interface framework for text searches. *D-lib Mag* 3(1):18–20
- The Stanford Natural Language Processing Group (2014) Natural language processing software. <https://nlp.stanford.edu/software>. Accessed in December 2020
- Viégas FB, Golder S, Donath J (2006) Visualizing email content: portraying relationships from conversational histories. In: Proceedings of the SIGCHI conference on human factors in computing systems. ACM, pp 979–988. <https://doi.org/10.1145/1124772.1124919>
- Wang Y, Yu M, Shan G et al (2019) Vispubcompas: a comparative analytical system for visualization publication data. *J Vis* 22(5):941–953. <https://doi.org/10.1007/s12650-019-00585-2>

- White RW, Roth RA (2009) Exploratory search: beyond the query-response paradigm. *Synth Lect Inf Concepts Retrieval Serv* 1(1):1–98
- Wildemuth BM, Freund L (2012) Assigning search tasks designed to elicit exploratory search behaviors. In: *Proceedings of the symposium on human–computer interaction and information retrieval*. ACM, pp 4:1–4:10. <https://doi.org/10.1145/2391224.2391228>
- Wilson M, Russell A, Smith DA et al (2006) mSpace: Improving information access to multimedia domains with multimodal exploratory search. *Commun ACM* 49(4):47–49
- Yalcin MA, Elmqvist N, Bederson BB (2017) Keshif: rapid and expressive tabular data exploration for novices. *IEEE Trans Vis Comput Graph* 24(8):2339–2352. <https://doi.org/10.1109/TVCG.2017.2723393>
- Zeng W, Dong A, Chen X et al (2021) Vistory: interactive storyboard for exploring visual information in scientific publications. *J Vis* 24(1):69–84. <https://doi.org/10.1007/s12650-020-00688-1>
- Zhang G, Tao S, Zeng N et al (2020) Ontologies as nested facet systems for human–data interaction. *Semantic Web* 11(1):79–86. <https://doi.org/10.3233/SW-190378>
- Zhao J, Collins C, Chevalier F et al (2013) Interactive exploration of implicit and explicit relations in faceted datasets. *IEEE Trans Vis Comput Graph* 19(12):2080–2089. <https://doi.org/10.1109/TVCG.2013.167>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.