<https://helda.helsinki.fi>

Characterizing the Quality of Insight by Interactions: A Case Study

He, Chen

2021-08-01

He , C , Micallef , L , He , L , Peddinti , G , Aittokallio , T & Jacucci , G 2021 , ' Characterizing the Quality of Insight by Interactions: A Case Study ' , IEEE Transactions on Visualization and Computer Graphics , vol. 27 , no. 8 , pp. 3410-3424 . <https://doi.org/10.1109/TVCG.2020.2977634>

<http://hdl.handle.net/10138/333450>

<https://doi.org/10.1109/TVCG.2020.2977634>

acceptedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Characterizing the Quality of Insight by Interactions: A Case Study

Chen He, Luana Micallef, Liye He, Gopal Peddinti, Tero Aittokallio, and Giulio Jacucci

Abstract—Understanding the quality of insight has become increasingly important with the trend of allowing users to post comments during visual exploration, yet approaches for qualifying insight are rare. This paper presents a case study to investigate the possibility of characterizing the quality of insight via the interactions performed. To do this, we devised the interaction of a visualization tool—MediSyn—for insight generation. MediSyn supports five types of interactions: selecting, connecting, elaborating, exploring, and sharing. We evaluated MediSyn with 14 participants by allowing them to freely explore the data and generate insights. We then extracted seven interaction patterns from their interaction logs and correlated the patterns to four aspects of insight quality. The results show the possibility of qualifying insights via interactions. Among other findings, exploration actions can lead to unexpected insights; the drill-down pattern tends to increase the domain values of insights. A qualitative analysis shows that using domain knowledge to guide exploration can positively affect the domain value of derived insights. We discuss the study’s implications, lessons learned, and future research opportunities.

Index Terms—Insight, interaction, interaction pattern, entity, visualization, insight-based evaluation

1 INTRODUCTION

Insight in visualization denote an advance in knowledge or a piece of information [11]. Gaining insight is considered one of the major purposes of visual data exploration [70]. Hence, understanding the user insight generation process when using visualization tools is an important goal of visualization research [20, 50]. For instance, Yi et al. [70] identified four categories of insight gaining processes through an extensive literature review: providing an overview, adjusting the level of abstraction and/or the range of selection, detecting patterns, and matching the user’s mental model of the data. Visualization can support these processes to foster insight [2].

As insight generation is usually aided by user interaction in visualization, Reda et al. [57] and Guo et al. [28] explored which types of interactions foster or hinder insight via interaction logs and verbal transcripts. In this paper, instead of investigating the relations between interaction types and the *quantity* of insight, we seek to characterize the *quality* of insight by interactions.

The trend of allowing users to post comments during visual exploration (e.g., [32, 44, 64]) makes the quality of generated insights more critical than the quantity. Wang et al. [65] suggested that unrelated or incorrect insights would degrade the benefits of representing insights in a visualization. However, verifying and validating diverse insights is challenging by nature [65]. We suggest that the quality of insight can be relevant to the types of interactions performed. The interactions performed result from user intent [40, 67], whereas user intent may further affect the quality of insight. For instance, with the intent of taking an in-depth look at an issue, a user may retrieve the details of the data for exploration, and then the insight that the user gains may be deeply related to this issue. Thus, the interaction of *retrieving details on demand* is related to the *depth* of the generated insight based on a distinct user intent.

Gotz and Zhou [25] characterized user interaction at four levels of granularity: tasks, sub-tasks, actions, and events. Tasks and sub-tasks are usually domain specific, whereas events, such as mouse clicks, have no meaning without context. In the action level, each action represents a distinct user intent, which makes actions generic across visualization

tools yet semantically rich. Interaction categorization based on actions also serves our purpose of using user intent to relate interaction and insight. Furthermore, ElTayeb and Dou [20] suggested an extra level between actions and sub-tasks that consists of patterns of multiple actions to support the reuse of analysis across different applications.

This work thus explores *which action or pattern relate to which aspect of insight quality* via a case study. First, we redesigned the interaction of MediSyn [29], a visualization tool that synthesizes five datasets of drug-target relations. MediSyn uses the concept of entities (that is, taking drugs and targets as entities) to afford five types of interactions for insight generation: 1) selecting entities of interest, 2) connecting relevant entities, 3) elaborating by retrieving the details of entity relations, 4) exploring other entities, and 5) entity-based insight sharing. We then evaluated MediSyn with 14 participants by asking them to input their tasks, freely explore the data with their tasks in mind, and derive insights by inputting notes. We graded the recorded insights on four aspects, manually extracted seven interaction patterns from the logged interactions, and analyzed the correlations between interactions and insights. The results show the potential to qualify insights by interactions. Among others findings, exploration actions tend to increase the unexpectedness of insights; the drill-down pattern can lead to insights with high domain values, which resonates with a qualitative analysis. The qualitative analysis of user strategies also uncovered that with domain-specific data, using domain knowledge to guide data exploration helps derive insights with high domain values; users tend to explore public insights when they are stuck with data exploration. We discuss the implications of this study, lessons learned, and future research opportunities.

2 RELATED WORK

The traditional task-based evaluation, which measures task time and error, hinders the assessment of the exploratory feature of visualization [45, 49]. Researchers, therefore, proposed insight-based evaluation instead, which investigates the characteristics of insights derived from visualization tools [58, 59]. Choe et al. [13] and Yang et al. [68] studied categories of insights derived from various charts, as well as the appropriate charts to use to derive these insights, to gain knowledge on building visualization tools for insight discovery and communication. Smuc et al. [62] evaluated a visualization by analyzing 1) the quantity of generated insights, 2) insight categories, and 3) relations between insights (i.e., how insights build on one another). They found that the third analysis was more informative on improving the design of the visualization tool.

However, analyzing insight alone is often limited. Looking into the insight generation process as a whole can shed more light on improving visualization design [45]. Mayr et al. [45] compared three evaluation methods: the task-based method, the insight-based method, and

-
- Chen He and Giulio Jacucci are with the Department of Computer Science, University of Helsinki. E-mail: chen.he@helsinki.fi and giulio.jacucci@helsinki.fi.
 - Luana Micallef was with Human-Centered Computing, Department of Computer Science, University of Copenhagen.
 - Liye He and Tero Aittokallio are with the Institute for Molecular Medicine Finland, University of Helsinki. E-mail: liye.he@helsinki.fi and tero.aittokallio@helsinki.fi.
 - Gopal Peddinti is with the VTT Technical Research Center of Finland Oy. E-mail: gopal.peddinti@vtt.fi.

problem-solving strategies. Problem-solving strategies are extracted by analyzing thinking aloud data, eye movement data, and interaction logs. They found that the insight-based approach informs the types of insights that the tool fosters or hinders, whereas analyzing problem-solving strategies helps identify how the tool should be improved.

The key role that interaction plays in identifying the insight generation, sensemaking, and reasoning processes puts the science of interaction under the focus of the visualization community [21, 53, 54]. Pohl et al. [55] studied interaction patterns extracted from logs during user visual exploration, and they found that patterns can be valid across visualizations. To support interaction pattern analysis, researchers have proposed several tools. Liu et al. [41] and SensePath [47] used multiple-linked views to support the pattern analysis of user online activities. Liu et al. coordinated multiple levels of granularity of web clickstream data. SensePath captures and displays user actions in temporal order. Analysts can inspect the web page, the recorded video, and the transcribed information of a selected action in other linked views [47]. Blascheck et al. [6, 7] proposed two visual analytics tools that integrate transcribed thinking aloud data, eye movement data, and interaction logs to assist with the analysis of user studies. These tools support the interactions such as pattern search and comparisons between participants for analysis. Blascheck et al. further used this concept to study the strategies that users employ to discover the interaction techniques available in a visualization [8].

Apart from interaction pattern analysis, other metrics are used to reveal new facets of people’s interaction with visualization tools. Two such metrics of exploration uniqueness and exploration pacing were proposed by Feng et al. [22]. Battle and Heer [4] found that users tend to plan and execute a few steps of their interactions at a time as an exploration pace.

Recorded interaction logs can also be used to infer user tasks [24], to recover reasoning processes [16, 17, 39], to predict task efficiency and personality traits [9], etc. Dabek and Caban [15] built a model that learned from user interactions of solving close-ended tasks and provided suggestions to new users to guide them through the same tasks. Shrinivasan et al. [60] proposed an algorithm that recommends related notes and concepts to users based on their line of analysis.

Previous work has significantly advanced research on evaluating visualization and understanding users by analyzing user insight or interaction. However, the interrelations between insight and interaction are under-explored. For instance, the thinking aloud data acquired by Blascheck et al. were not further analyzed and connected to insight. There are a few exceptions to this, however. Reda et al. [57] proposed a graph that captures transitions between users’ mental and interaction states. Each design variation generates a transition graph, which facilitates the discovery of which design variation or interaction type fosters or hinders insight. However, this method omits the effects of interaction sequences on insight generation. Guo et al. [28] analyzed the correlations between the types or sequences of interactions and the quantity of insight by asking users to complete open-ended tasks using a visual analytics tool. The results showed, among others, that exploration actions promote the generation of facts, whereas filtering actions inhibit it. With a similar purpose of bridging interaction and insight, this work investigates interaction types and patterns in relation to the quality of insight.

3 DESIGN CONSIDERATIONS

We propose to use the notion of entities to design the interaction for insight generation, and we review the quality of insight for investigation.

3.1 Entity-based interaction design

Entities represent any real-world objects or concepts, such as people, places, and topics. To design interaction for data visualization, we found the notion of entities useful in supporting visual representation and manipulation, as stated in Klouche et al. [38]. Entities can yield other related entities to support exploration; they can be organized to support pattern recognition; and they can be shared to assist in collaboration [38]. These concepts have been successfully applied to various visualization tools. For instance, Pivotpaths [19] allows users

to pivot on an entity to view its relevant entities in a graph. Jigsaw [63] coordinates multiple views, and each one shows entity relations from a different perspective to support pattern recognition. Bier et al. [5] proposed an entity workspace that allows users to organize and share their entity graphs and obtain recommendations based on their peers’ entity graphs.

The trend of allowing users to share insight during visual exploration has led to two new design considerations that we identified from previous work. First, the visualization and insight should be linked to support bidirectional exploration [37, 44]. On the one hand, data exploration should be linked to insight exploration [31, 32]. Heer et al.’s research [32] and InsideInsights [44] display related insights when users navigate to a visualization to which those insights refer. This can be opportunistic. Heer and Shneiderman [31] later proposed the concept of data-aware annotations, which indicates applying annotations to different views of the same data. We found that entities can support this concept well. The entities that compose a view can change, whereas insights can be dynamically displayed to match the entities under exploration (e.g., [14]). The other direction is exploring how to use insight to stimulate data exploration. Bier et al. [5] allowed analysts to discover new entities as potential interests from their peers’ entity graphs. ManyEyes [64] and InsideInsights enable restoring the visualization when the user navigates to an insight. This method, however, limits users’ ability to derive insight from relating multiple views.

We propose using entities to mediate insight and visualization exploration. For instance, the entities mentioned in the insight can be used to stimulate the visual exploration of their relations. Entities in the visualization can trigger user exploration of related insights.

The second consideration is that mechanisms should be provided to help users rationalize individual insights. Alspaugh et al. [3] and Madanagopal et al. [43] identified the key role of analytic provenance in helping data analysts recall, reason, and collaborate through expert interviews. But the provenance feature is not well-supported by existing tools [3]. See Nguyen et al. [48] and Ragan et al. [56] for detailed reviews of existing research on visualizing provenance data. Provenance can be supported by enabling users to attach data sources to notes [66, 71], by allowing them to manually create a trail depicting visualization or data flows [18, 34], or by automatically capturing and displaying actions [27, 42, 47, 72] or visualization states [10, 30, 33, 35].

To support insight communication, some research has featured the manual construction of a knowledge graph. A knowledge graph allows users to take recorded visualization states [46], to create notes [61], or to do both [52] as entities and link the entities as causal relations to communicate findings. The manual construction of a knowledge graph may allow users to flexibly map the graphs to their mental models, but this approach requires extra user effort. To provide the rationale of an insight, a trail that leads to the insight and that records semantically meaningful actions as well as their resulted visualization states is attached to the insight in our design.

3.2 Quality of insight

To evaluate insights, some research (e.g., [13, 28, 62]) has classified insights by analyzing the content of collected insights. For instance, Guo et al. [28] classified insights as facts, hypotheses, and generalizations. To assess the quality of insight, however, we found that it was important to look at one insight from multiple perspectives rather than simply categorizing it. North et al. [49] characterized insight as complex, deep, qualitative, unexpected, and relevant. This is for defining insight rather than for quantifying insight for practical use. To support the practical and in-depth analysis of insight quality, Saraiya et al. [58] characterized insight by directness versus unexpectedness, correctness, breadth versus depth, and domain values based on realistic case studies. We use this characterization to qualify insight in this paper, although previous studies that have analyzed these characteristics have not produced promising results. For example, in quantifying the domain value of individual insights, North et al. [51] found that only the “summary” type of insights was more valuable than the others. In Saraiya et al.’s study, most insights were the “breadth” type rather than the “depth” type. We think that this may be due to the lack of interaction types

supported by the visualizations they studied. The visualizations that North et al. evaluated were mostly static, whereas the ones that Saraiya et al. studied supported mostly standard charts, such as scatter plots and histograms, which facilitated pattern recognition, but limited the user intent of drilling down into the details of individual data items. The limited support of having users elaborate on the data may cause a lack of depth in insights. This work would be the first attempt to link multiple types of interactions to these aspects of insight quality.

4 THE ENTITY-BASED VISUALIZATION TOOL

The tool we used synthesizes five publicly available, manually curated drug-target datasets to assist in the research of personalized cancer therapy. *Target* here means point mutations in specific tumors.

4.1 Visualization design

The visual encodings of data variables have been rationalized in previous work [29]. MediSyn uses a matrix-based view to relate drugs, mutations, and tumors. As Fig. 1B shows, the upper rows depict tumors, the columns denote mutations, the lower rows represent drugs, and the table cells depict their relations. The horizontal line of a cell in the upper rows denotes the tumor to which the drug-target effect refers. In the lower rows, we use colored horizontal bars to depict the corresponding effects and their evidence levels. The hues represent effects. For example, the effects of responsiveness and resistance are depicted by green and red hues, respectively. The length of the bars denotes the evidence level of an effect. As validated in the previous study, MediSyn makes knowledge about genomically informed therapy accessible and evaluable to biologists [29]. Based on the previous study [29], we had three updates in this version:

- We use discreet bar segments to ease the previous version’s perceptual difficulty of comparing the lengths of the bars across columns. From one to five segments, the bars represent ascending evidence levels from preclinical assays to guidelines, such as Food and Drug Administration guidelines.
- We extended the number of datasets from two to five. The five datasets result in 427 types of drugs or drug combinations, 827 kinds of mutations, 58 kinds of cancers, and 20,887 total relations. The various datasets contain 19,299, 365, 320, 408, and 495 numbers of relations.
- We use juxtaposed bars to depict values from various datasets to support comparison. The previous design, which used overlaid layers to depict data values from two datasets, is not scalable to the increased number of datasets. In this design, for instance, two juxtaposed bars of the outlined table cell in Fig. 1B depict resistant effects from two datasets with varying evidence levels.

4.2 Entity-based interaction design

To support insight generation, we redesigned MediSyn’s interaction. We present the types of interactions based on Yi et al.’s taxonomy [69]. This taxonomy categorizes interaction based on user intent, which maps to the “action” tier according to Gotz et al. [25] and serves our purpose of relating interaction and insight by user intent. Additionally, we created a new type of action called “share” that indicates the user intent of exploring shared insights. We also discuss how each type of action can support insight generation.

4.2.1 Select entities of interest

Considering drugs, mutations, and tumors as entities allows users to explore the data from various aspects. As Figure 1A shows, users can select drugs, mutations, and tumors from the list to add to the view. Once a user selects an entity, its relevant entities across all datasets will be retrieved and depicted in the view for analysis. The user can deselect an entity by clicking the cross icon next to an entity (Fig. 1).

4.2.2 Connect related entities

Viewing data from various perspectives facilitates insight [53, 70]. MediSyn allows users to center on an entity to see its connections to other entities in two ways. One is to highlight related entities by

mousing over the entity, and the other is to aggregate the related entities through the reordering of rows and columns by clicking on the entity, which we call a “pivot” interaction. To be specific, when a user mouses over a drug, the columns of its relevant mutations and the upper rows of its relevant tumors are highlighted in a semitransparent orange color. When a user clicks on the drug, its relevant mutations gather to the left of the columns, which pushes the irrelevant columns to the right in animated transitions. The relevant tumors gather to the bottom of the upper rows in the same manner (Fig. 1B). The same actions also apply to the other types of entities, mutations and tumors. For the display of mutations, the same mutation can appear in multiple columns because the same mutation can appear in various tumors. When a user clicks on a mutation, it aggregates the columns of that mutation as well as its relevant drugs and tumors.

This allows the user to analyze the data from various perspectives, centering on the drug, tumor, or mutation. For example, users can focus on a drug by highlighting or aggregating relevant mutations and tumors, and they can investigate how the effects change across mutations and tumors. The users can also sequentially highlight two tumors and check how the drugs overlap.

4.2.3 Elaborate by retrieving the details of entity relations

Being able to view the data in multiple levels of granularity allows users to focus on what is important at the current moment and omit distracting information [5]. With MediSyn, users can retrieve details on demand for further inspection. When a user hovers over a colored bar, the corresponding drug, mutation, and tumor are highlighted, which indicates the intent of viewing an entity relation. When users click on a colored bar, they can see its details on the right side of the matrix (Fig. 1C), which includes the natural language description of the relation and the titles and digital object identifiers (DOIs) of the publication sources. Users can click on DOI links to further review the publications in a new window.

4.2.4 Explore other entities

To generate insight, users first need to collect sufficient information [54]. Exploring other entities is an effective way to do so. In MediSyn, exploring other entities is indicated by the user adding an entity to the view while attending to other information. First, in the current view, the user can add an existing entity to the view by mousing over the entity and choosing to “select” the entity through the drop-down menu (Fig. 1B). Afterward, all relevant entities will be added to the view. Second, the user can select the entity to add to the view from an insight display, as articulated next.

4.2.5 Entity-based insight sharing

The user can input an insight by clicking the upper-left paper-and-pen icon (Fig. 1). A text area resembling a sticky paper will appear. Afforded by sticky papers, the user can drag and drop the area freely on the interface and can continue to analyze the data while taking notes. By clicking the paper icon next to the note-taking icon, the user can check publicly available notes as well as one’s own notes, as shown in Fig. 1D. The user can choose to like or dislike notes by clicking the thumbs-up or thumbs-down icon, and view the provenance trail by clicking the “view provenance” button (Fig. 1D).

To support the bidirectional exploration of the insight and visualization we discussed in Section 3.1, entities mentioned in a note are automatically extracted as tags of the note and are displayed below the note itself, as shown in Fig. 1D. The user can explore the entity mentioned in the note by clicking the tag, which results in a select action. On the other hand, visual cues in the visualization imply opportunities to explore related notes without being distracted [5]. If an entity in the view has related notes available, it will be shown in bold (Fig. 1B). As the user hovers over the bold entity, a drop-down menu allows him or her to choose to “view notes” on part D of Fig. 1.

To help users rationalize others’ insights, each insight has a trail attached that records semantically meaningful actions leading to the insight. To construct the visualization of interaction trails, we take a description of an action together with its resulting view as a node and

Fig. 1. The interface of MediSyn consists of (A) the datasets and entities for selection, (B) the visualization, (C) details related to a clicked colored bar, and (D) public and private notes. The pink sticky paper allows users to enter and drag around a note.

connect them sequentially (e.g., Appended Fig. A1). For simplicity of display, the trail will always take the action of changing the numbers of entities in the view as a starting point, such as selecting, de-selecting, and selecting an entity from a note. Sequential select and de-select actions are combined into one node. The interaction of writing the corresponding note completes the trail. In addition, hovering interaction is omitted. See Supplemental Video 1 for a video demonstration of the interaction of MediSyn. Additionally, MediSyn is accessible with a user registration at <https://d4health.hiit.fi>.

4.3 Use case

To illustrate the use of MediSyn, we describe a use case collected from the user study we will present later. To explore approved drugs for lung cancer, the biologist selected the entity “lung” from the left list. The selected entity and its related mutations and drugs were then displayed in the visualization. The biologist viewed the details of some relations by clicking on or hovering over colored relation bars sequentially. The relations she checked included that the lung cancer carrying mutation EGFR(T790M) was resistant to erlotinib or afatinib treatment, the lung cancer carrying BRAF(V600E) was responsive to the drug combination of “dabrafenib + trametinib” based on guidelines, and so forth. She then retrieved all available notes related to the entity “lung” cancer. The first note in the list was “Lung cancer carrying EGFR (T790M) are highly resistant to EGFR inhibitors like afatinib and erlotinib. However, combination of afatinib together with cetuximab or nimotuzumab is effective.”

The biologist continued to explore the details of several relations and highlighted three entities—“dabrafenib + trametinib,” “egfr tkis + mek inhibitors in egfr mutant,” and “egfr tkis”—to see their relations to targets. Next, she started to input the following note: “There currently are only a few approved targeted therapies for lung cancer, such as ceritinib, afatinib and erlotinib, but several more are undergoing clinical trials.” During this inputting, she retrieved all notes related to “ceritinib.” Afterward, she moused over the relation that showed that the lung cancer carrying BRAF(G469A) is resistant to egfr tkis treatment based on a case report. She then explored the existing targets of egfr tkis treatment by clicking on the entity. After clicking on a series of entities of drug combinations, including “egfr tkis + mek inhibitors in egfr mutant,” she continued to write, “Depending on a mutation, different drugs show desired potency. In particular, lung cancer carrying BRAF(G469A) mutation is resistant to egfr tkis treatment but shows response to a combination of egfr tkis and mek inhibitors in preclinical studies.” She then saved and closed the note.

5 USER STUDY

To investigate the research question—which action or pattern relate to which aspect of insight quality—and explore user strategies of deriving insight using MediSyn, we conducted a lab study.

5.1 Participants

In total, we recruited 18 paid participants from a research institute related to personalized cancer therapy. Two were recruited for a pilot study whose research topics were machine prediction of drug-target interactions. The remaining 16 were recruited for the lab study, but two quit in the middle of the experiment, as one was interested in mutation frequencies in cancer, not drug-target interactions, and the other, a junior researcher, stated she was not able to write notes. The remaining 14 participants (mean age: 31.36; SD: 5.11; female: 5) had sufficient knowledge to analyze the data. Among them, eight were doctoral students, four were postdoctoral researchers, one was a master’s student, and the last was a senior researcher. Their specific focus areas differed. For instance, one focused on lung cancer, another specialized in ovarian cancer, two had expertise in leukemia, and one was studying multiple melanoma and leukemia. Others focused on drug discovery or translational medicine research and did not explicitly point out diseases of focus.

5.2 Pilot study

The purpose of the pilot study was to collect public notes and to have pilot participants go through the procedure of the actual study. The two pilot participants followed the same procedure as in the lab study, as we describe next. However, their task differed from that of the lab study. They were encouraged to write as many notes as they could and to post them publicly. If they reached a certain quantity of qualified notes, they could get higher amounts of compensation. A qualified note was considered a useful observation on the data and had to contain at least one entity name. As the participants were instructed, the published notes were anonymous to other users to avoid the influence of user identity on the experiment. The pilot study was conducted in the participants’ own offices with use of their own computers.

We collected 40 public notes from the pilot study. Another domain expert, one of the authors, screened through these notes and selected 27 as public notes for the lab study. The remaining were screened out because they were either simple observations of drug-target relations for the visualization or purely based on domain knowledge and disconnected from the visualized data. Accordingly, we updated the description of a qualified note for the lab study, which we present next.

Action	Application-specific instantiations
Select	Select an entity of interest from the left list.
Connect	Highlight entities related to an entity by mousing over the entity. Aggregate entities related to an entity by clicking on the entity.
Elaborate	Highlight a table cell depicting entity relations and its related entities by mousing over the cell. Retrieve the details of an entity relation by clicking on a colored bar. Retrieve a publication source by clicking a DOI link.
Explore	Retrieve other entities related to an existing entity by mousing over the entity and clicking “select” from a drop-down menu. Explore an entity mentioned in a note by selecting the entity from the note display.
Share	Retrieve all available notes. Retrieve notes related to an entity by mousing over the entity and clicking “view notes” from a drop-down menu. View the provenance view of a note. Like or dislike a note.

Table 1. Actions and their application-specific instantiations.

On average, the public notes contained 3.74 different entities each (SD: 2.26).

5.3 Lab study

The lab study started with a training session. Participants first walked through the functionalities of MediSyn by following a step-by-step interactive tutorial we created using Intro.js [1]. In total, the tutorial consisted of 12 interactive steps. In each step, the participants carried out 1-2 interactions following the instructions and clicked the “next” button to move to the next step until they finished the tutorial. After this process, they were instructed to create accounts in MediSyn and to freely explore the tool until they were satisfied. They were encouraged to ask the experimenter questions during the training session. The training session took around 20 minutes.

In the actual study, the participants were instructed to first input a task by opening a dialog box. Then they could freely explore the interface with the task in mind and write notes. Whenever they changed tasks, they could open the dialog box, input a new task, and continue to explore that new task. We provided some open-ended task examples to help participants start, such as “to investigate ** disease” and “to explore effective drugs for ** tumor.” They each could finish the experiment when they thought they had written at least five qualified notes. The requirements for a qualified note were as follows: i) it was a useful observation, such as one’s findings, hypotheses, or generalizations, on the underlying data based on one’s domain knowledge (a simple observation of a drug and cancer mutation relationship from the visualization, for example, not being a useful note); and ii) it included at least one entity name, such as the drug, tumor type, and/or mutation.

During the experiment, we logged user interactions and recorded the screen and voice of each participant with the participants’ consent. We did not establish time constraints for the study. Participants were only able to post private notes, and their likes and dislikes of notes were visible only to themselves, as they were instructed.

After the experiment, participants filled out a questionnaire regarding their subjective feedback about MediSyn. The entire experiment took around 1-1.5 hours. The lab study was conducted using a Chrome browser in full-screen state on a 27-inch display with an external mouse and keyboard. The display resolution was 2,560 * 1,440 pixels.

5.4 Measures

To investigate the research question, we logged types of interactions, time stamps and entities related to each interaction, and the user inputs of the tasks and notes. Hovering interaction was recorded when data were highlighted for more than three seconds.

We then constructed interaction trails from the logged sequential interaction data. The trails we collected always started with an entity-selection action, either from the list or from the public notes. De-selecting all entities, or sequentially selecting and de-selecting entities to yield a completely different set of selected entities, marked the start of a new trail. In each trail, the last note-writing action indicates the end of the trail, as we are interested in actions that lead to insight.

We also manually extracted interaction patterns from the interaction trails. An interaction pattern here denotes a sequence of actions consisting of at least two types of actions. To extract patterns, we considered not just the sequential relations of actions but also the entities on which the participants operated. For instance, suppose a user were to start from select actions, then retrieve the details of an entity relation, and then connect entities by pivoting on one entity. If entities in the connect action did not appear in the previous elaboration action, then we would have counted two patterns, as select – elaborate and select – connect. Otherwise, we would have counted this as one pattern: select – elaborate – connect. After extracting all patterns from the trails, we took the patterns performed by at least two participants as candidate patterns for further analysis.

5.5 Assessing notes

We collected 86 notes from the study. Removing 27 notes, 25 of which were data requests (notes such as “FAM46C mutation is quite frequent in Myeloma. Can’t find any information in the database.”) and two of which were user interface suggestions, we took the remaining 59 notes as insights. As stated in Section 3.2, following Saraiya et al. [58], we characterized the insight quality by directness versus unexpectedness, correctness, breadth versus depth, and domain value. Two experts, also authors of this paper, first agreed upon detailed grading criteria (Table 2) and then graded the notes on the four aspects from 1 to 5 independently. For instance, directness versus unexpectedness indicates a spectrum (1 to 5) of one insight quality, where grade one indicates direct insights, and grade five denotes unexpected insights. Directness versus unexpectedness was assessed by the relation of an insight to the user-inputted task. Direct insights related to user tasks, whereas unexpected insights were exploratory and went beyond user tasks. Note that a task could have zero or multiple notes. We assigned a note to the last inputted task, as participants were instructed to do.

Correctness indicates whether the insight can be verified or not. If the insight was a hypothesis, we assessed the rationale for the derivation of the hypothesis. Breadth versus depth relates to the scope of the insight. Broad insights refer to multiple entities or entity groups, whereas deep insights relate to a specific aspect of an entity or entity relation. A domain value ranges from low for simple observations to high for pattern recognition and hypothesis derivation.

Kendall’s tau-b association test shows a correlation between two raters ($p < 0.01$). The tau-b values for the characteristics grading of directness versus unexpectedness, correctness, breadth versus depth, and domain value are 0.415, 0.290, 0.423, and 0.245, respectively. Then, for grades differing by more than one point, two raters had a discussion to either change the score or keep their original values. Finally, we took the average of the grades from two raters as the final grades. The distribution of grades of individual characteristics averaged across participants is shown in Fig. 2. The insights we collected tended to be direct rather than unexpected, which we discuss in Section 6.3.

We computed the correlations between each pair of characteristics. Two moderate-to-strong correlations existed between the domain value and correctness of insights ($r = -0.573$, $p = 0.032$), and the domain

Characteristics	Grading criteria
Directness versus Unexpectedness	Give 1-2 if the note is related to the task. Give 4-5 if the note goes beyond the task description. Give 3 if the note is in between.
Correctness	Give 1-2 if the note is verified to be wrong through the visualization or Internet. Give 4-5 if the note is verified to be right through the visualization or Internet. Give 3 if the correctness is difficult to discern.
Breadth versus Depth	Give 1-2 if the note refers to relations of mutation groups, multiple tumor types, and/or drug groups. Give 4-5 if the note refers to one or more specific relations of drugs and mutations. Give 3 if the note is in between.
Domain value	Give 1-2 if the note is a simple observation from the visualization, such as one drug-target relation or a missing value. Give 1-3 if the note complements missing information with domain knowledge. Give 3-5 if the note reveals patterns across the visualization. Give 4-5 if the note derives hypotheses from the visualization.

Table 2. Grading criteria of insight characteristics.

value and breadth versus depth of insights ($r = -0.607$, $p = 0.021$). The correlations were not statistically significant after the Bonferroni p-value adjustment. Considering all the insights we collected were correct, we could skip the first correlation. However, the second correlation hinted that broader insights tended to have higher domain values in our assessment.

Fig. 2. The distribution of grades of individual characteristics averaged across participants.

5.6 Results

In total, we collected 37 trails, leading to the 59 insights (Appended Fig. A2). Some trails were short, such as “select, select, and write a note,” whereas some others could be as long as 60-70 interactions. Some participants wrote multiple notes in one exploration trail, whereas others wrote one note in each trail.

5.6.1 Interaction patterns

Interaction patterns indicate higher-level user intents that can be accomplished by carrying out multiple actions [24]. The following seven patterns resulted from the study. We named the patterns by the implied user intents. Each action in the pattern can contain arbitrary numbers of the action in succession.

- **Drill-down:** select – (connect) – elaborate. The participants selected an entity or entities from the list, explored entity relations in the view, and then elaborated on certain entity relations. The connect action in this pattern could be zero or more. It was also possible for the participants to carry out the combination of connect – elaborate in succession multiple times. This pattern featured drilling down into the data of interest. In total, nine users used this pattern.
- **Scanning:** select – connect. The participants selected entities from the list and then explored entity relations through connect actions. Differently from the previous pattern, in this pattern, participants did not dig into the details of entity relations. Thus, the count of this pattern excludes the count of the previous pattern. Eleven participants’ trails contained this pattern.
- **Sampling:** explore – (connect) – elaborate. Compared with the first pattern, rather than merely selecting entities from the list, in this pattern, the participants added an entity or entities to the view

through exploration actions, then carried out connect and elaborate actions. Similarly to the first pattern, connect actions existed zero or more times, and the combination of connect – elaborate actions could recur in succession. This pattern featured exploring and drilling down into the details of other interesting data. Six participants used this pattern.

- **Probing:** explore – connect. The participants added an entity or entities to the view by selecting from the existing view or the note display and then explored the entity relations through connect actions. Similarly to the scanning pattern, the count of this pattern excludes the previous pattern. This pattern existed in eight participants’ trails.
- **Expanding:** select – elaborate – connect. The participants selected some entities, then elaborated on certain entity relations, and ultimately switched to connect entities in the view. Differently from the first pattern, in this pattern, the participants transitioned from the details of interest to a broader view of data relations. Only three participants used this pattern.
- **Referencing:** select / explore – share. The participants selected or added entities to the view and then explored notes related to certain entities in the view. Seven participants used this pattern.
- **Annotating:** select – write. The participants selected some entities and wrote notes directly without further exploration. Seven participants performed this sequence of actions.

Referring to previous work, which also extracted action patterns [24, 28, 55], we found two common patterns. The drill-down pattern proposed by Gotz and Wen [24] features users’ intent of narrowing down the analytic focus to a targeted subset of items, which corresponds to the pattern of select – (connect) – elaborate in this work and the locating pattern in Guo et al. [28]. The pattern of explore – (connect) – elaborate was also discovered by Pohl et al. [55] and Guo et al., the latter of whom led us to name the sampling pattern. But for an application, specific actions or patterns can dominate. For instance, a visualization that provides an overview of the data first and allows users to filter the data of interest may not result in the pattern starting with select actions. We thus suggest that the patterns we extracted from the study are specific to a type of applications.

5.6.2 Relations between interaction and insight

To investigate the research question, we took the proportion of each action and pattern from each participant for analysis. The insight characteristics were averaged on each participant. As we considered de-select actions “passive” reactions to select actions, we did not count them in the analysis. We omitted too the expanding pattern due to insufficient data, as well as the insight correctness, because it might have been affected by the user interpretation of the visual encodings. In the analysis, we did not provide hypotheses on the relations of the interaction types and insight characteristics, but rather we investigated each pair to find interesting and noteworthy relations. Depending on whether the values were normally distributed or not via the Shapiro-Wilk test,

	Directness versus Unexpectedness
Exploration action (Mean: 0.07, SD: 0.10)	$\tau_b = 0.384$, $p = 0.064$, CI [0.0127, 0.7137]
Explore – (connect) – elaborate (Mean: 0.09, SD: 0.13)	$\tau_b = 0.489$, $p = 0.024$, CI [0.1800, 0.7574]
Explore – connect (Mean: 0.15, SD: 0.18)	$\tau_b = 0.375$, $p = 0.076$, CI [-0.0132, 0.6805]
Select / Explore – share (Mean: 0.11, SD: 0.13)	$\tau_b = 0.326$, $p = 0.126$, CI [-0.0977, 0.7686]
Number of interactions (Mean: 13.57, SD: 11.87)	$\tau_b = 0.363$, $p = 0.079$, CI [-0.0749, 0.8049]
	Breadth versus Depth
Elaboration action (Mean: 0.14, SD: 0.15)	$\tau_b = 0.262$, $p = 0.202$, CI [-0.1707, 0.6110]
Explore – (connect) – elaborate (Mean: 0.09, SD: 0.13)	$\tau_b = 0.462$, $p = 0.033$, CI [0.0679, 0.7923]
	Domain value
Select – (connect) – elaborate (Mean: 0.20, SD: 0.24)	$\tau_b = 0.365$, $p = 0.081$, CI [-0.1056, 0.7754]
Number of interactions (Mean: 13.57, SD: 11.87)	$\tau_b = 0.256$, $p = 0.207$, CI [-0.1950, 0.6387]

Table 3. The report of the assessment includes Kendall’s tau-b, bootstrapped 95% confidence intervals (CI), and p-values. CIs were computed using the percentile method with 2,000 bootstrap replicates. The p-values should be compared with the Bonferroni-corrected value of 0.003 due to multiple comparisons.

Fig. 3. Scatter plots show the relations between (a) exploration actions and unexpectedness of insights; (b) pattern of explore – connect and unexpectedness of insights; (c) elaboration actions and depth of insights; (d) pattern of select – (connect) – elaborate and domain values of insights. Participants are differentiated by colors.

we calculated Kendall’s tau-b or Pearson’s correlation to explore their relations. Except for the actions of select and connect, all other actions and patterns were positively skewed. We report all moderate to strong relations in Table 3. The p-values should be compared with the Bonferroni-corrected value of 0.003 due to multiple comparisons.

Exploring other interesting entities can lead to unexpected insights. Among the 14 participants, nine used exploration actions leading to an insight. Only two users explored entities from the note display. Kendall’s tau-b test shows a correlation between exploration actions and the unexpectedness of insights ($\tau_b = 0.384$, $p = 0.064$, Fig. 3(a)). In addition, two other correlations exist between the unexpectedness of insights and the sampling pattern ($\tau_b = 0.489$, $p = 0.024$), as well as the probing pattern ($\tau_b = 0.375$, $p = 0.076$, Fig. 3(b)).

Ten participants checked other notes through the visualization. Among them, three participants checked the provenance views of notes. We expect that the action of exploring other notes may lead to unexpected insights, but no correlation is found ($\tau_b = 0.148$, $p = 0.471$). However, a moderate correlation exists between the referencing pattern and the unexpectedness of insights ($\tau_b = 0.326$, $p = 0.126$).

Moreover, a correlation exists between the number of total interactions and the unexpectedness of insights ($\tau_b = 0.363$, $p = 0.079$). It is not so surprising that the more interactions the user carries out, the stronger the tendency is for the insights of that user to go beyond the initial task.

Select actions have a moderate but not statistically significant correlation to the directness of insights. Select actions were used the most (Mean: 0.41, SD: 0.30). We expect that select actions might relate to the directness of insights because with a clear intent in mind, the user might directly select the entities of interest to look up the answer, and consequentially, the answer that he or she gets may be directly related to the task. We calculated Pearson’s r between select actions and the directness versus unexpectedness of insights after removing the outlier in the characteristic of directness versus unexpectedness shown in Fig. 2. A moderate correlation exists between select actions and the directness of insights ($r(11) = -0.266$), but the result is not statistically

significant ($p = 0.379$). Additionally, the annotating pattern has a small to moderate correlation with the directness of insights, and the result is not statistically significant either ($\tau_b = -0.203$, $p = 0.345$).

Ten participants used elaboration actions that led to insights. Most elaboration actions involved retrieving the details of data cells. Only one participant further checked the publication page. Kendall’s tau-b shows the existence of a tendency for elaboration actions to lead to deeper insights ($\tau_b = 0.262$, $p = 0.202$, Fig. 3(c)). In addition, the sampling pattern shows a positive correlation with the depth of insights ($\tau_b = 0.462$, $p = 0.033$), whereas the drill-down pattern shows no relation to the depth of insights ($\tau_b = -0.035$).

Connect actions were used by 13 participants, the second most commonly used type after select actions in deriving insight (Mean: 0.32, SD: 0.17). Its relation to the insights is neutral regarding the characteristic of breadth versus depth ($r(12) = -0.047$, $p = 0.873$).

Regarding the domain values of insights, a positive correlation exists for the drill-down pattern ($\tau_b = 0.365$, $p = 0.081$, Fig. 3(d)). This indicates that *the action of drilling down into the details of interest tends to increase the domain values of the derived insights*. Interestingly, but not surprisingly, the annotating pattern has a small to moderate negative correlation to the domain values of insights, although the correlation is not statistically significant ($\tau_b = -0.219$, $p = 0.315$). In addition, a tendency is evident whereby the larger the number of interactions performed, the higher the domain values the insights have ($\tau_b = 0.256$, $p = 0.207$).

5.6.3 Qualitative analysis of user strategies

To understand the strategies participants took to gain insights with high domain values, we selected two groups of participants whose average domain values of insights were above 3.5 (Group H) and below 2.5 (Group L), respectively, and analyzed the insights generated by each group. Group H consists of participants 12, 9, and 3 (Appended Fig. A2), who took 25.82 (SD:24.91) actions and 8 min 36 sec (SD: 5 min 29 sec) on average to derive an insight. Participants 10, 8, and 7 comprise Group L, and they performed 14.63 (SD 12.24) actions and spent 4 min

One said, “I wanted to ask the same thing in a public note.” Another said that she used notes as a reference. One participant expressed that it was great to see other notes and how others arrived at the conclusions. Regarding user feedback on the provenance view, five users said it helped to explain the logic behind the notes. One stated that provenance helps if the selected entities are different from the current user’s. That is, the provenance view helped with understanding the scope of an insight. We elaborate on two cases from Fig. 4 to show how participants leveraged public notes to arrive at their own discoveries.

Case a: The participant input the task “Breast cancer, the recurrent mutations and effective therapeutics.” After some explorations including two actions of retrieving public notes, he said he was confused with what kind of discovery to write. Then he viewed the provenance of two public notes shown in the red rectangle a, Fig. 4a. Subsequently, he wrote a hypothesis in a similar pattern to the second public note.

Case b: The participant input the task “investigate Myeloma.” Then he selected the entity “Myeloma,” which resulted in five drugs and six mutations in the view. He chose to view notes related to “Myeloma,” which resulted in two public notes shown in the red rectangle b, Fig. 4b. He clicked the thumbs-up of the first note, the same note as the first public note in case a, and opened its provenance view. He then wrote “very limited findings on myeloma,” which we did not consider to be an insight. Then the participant input another task as “Investigate the similarities in mutations and drug responses between Acute myeloid leukemia and Myeloma.” He selected another entity: “Acute myeloid leukemia,” and wrote a note relevant to the first public note he explored.

The two cases imply that when participants were stuck with data exploration, they would then refer to public notes. In the two cases, public notes helped them to come up with an insight and a new task, respectively.

5.6.4 Other subjective feedback

Three users reported that they gained new knowledge concerning drug sensitivity relationships. Two participants found only the information that they had already known. One participant felt no need for such a tool; the relations between drugs and tumors were enough. A lot of the feedback involved requests for more information to be included in the tool. Five participants expressed that more drug-target data were needed. Three participants needed more information on drugs and/or mutations, including mutation frequency in various tumors (two participants) and drug classifications (two participants). Three participants wanted sources of experiments, such as cell lines, mouse models, or patients, along with publications.

5.6.5 Learning curve

To investigate whether usability issues exist which hinder user exploration and insight generation, we assessed the learning curve of MediSyn. Participants tended to take more time and a larger number of actions before writing the first note versus the following notes. The median amount of time that participants spent generating five notes are 460.5, 228.5, 377.5, 181, and 270 seconds, whereas the median steps are 16, 4.5, 8, 3, and 4 steps. Friedman’s analysis of variance shows effects among the order of notes considering the time taken ($\chi^2(4) = 12.29$, $p = 0.015$), or the actions taken ($\chi^2(4) = 12.06$, $p = 0.017$). However, the differences are not statistically significant after applying Benjamini-Hochberg corrected post hoc tests. Two participants stated that the tool was difficult to use in the beginning, but they became familiar with it after some time. Five participants expressed that it was easy to view the relations between various drugs and targets with MediSyn, resonating with the fact that most insights are correct from the gradings (Fig. 2).

Combining objective and subjective feedback, we conclude that the visual encodings of the data are easy to understand. After the user gained acquaintance with the interaction, which happened around writing the first note, MediSyn had no effects on the efficiency of writing notes.

6 LESSONS LEARNED AND FUTURE OPPORTUNITIES

Though the number of participants is not large in the study, results still support the possibility of qualifying insights by interactions. Exploration actions tend to increase the unexpectedness of insights, whereas the drill-down pattern can lead to insights with high domain values, confirmed by a qualitative analysis as well. However, results also show, using domain-specific data, that the involvement of domain knowledge during the exploration can affect the quality of generated insights. Also, considering that a case study from a single domain may limit the generalizability of the findings, we propose that further research with generic types of interactive visualizations can alleviate the effects of domain knowledge on insights quality and has the potential to positively impact this area. Moreover, personality traits have been shown to influence the number of insights generated by Green and Fisher [26]. How personality traits affect the quality of insight is an interesting question for future exploration. We discuss our reflections on this study next.

6.1 Characterizing insights

We used the four metrics from Saraiya et al. [58] in this study to qualify insight, although the quality of insight can not be limited to these metrics, such as the one Chris North proposed as simple versus complex [49]. This study demonstrated that other useful characteristics can also be correlated to various actions through user intent. The interaction logs and insights we collected can be reused. However, assessing the quality of insight can be laborious. To make the assessment process efficient, we found it important to have detailed grading criteria, along with some examples to ensure that the raters were on the same page.

Specific to the visualization we studied, we found that participants frequently spot that specific drug combinations are more effective to certain targets than the individual drugs. This type of insight could be an opportunity for insight automation, which suggests this type of pattern as mixed-initiative analysis [36].

6.2 Interacting with entities

Abstracting various real-world objects or concepts to entities, on the one hand, simplifies the way of designing interaction. For instance, we design the interaction with one or more entities, rather than designing for drugs, mutations, and tumors separately. On the other hand, similarly to the concept of object-oriented programming, the interaction can be flexibly extended to various numbers of entities, and entity affordances can be applied to any real-world data with similar relations. Thus, abstracting drug-target data to entities implies that MediSyn is not limited to this specific type of data, but rather can be applied to any real-world datasets which represent similar relations.

From the interaction trails (Appended Fig. A1), we recognized that the same actions are prone to appear in chunks, especially for the select and de-select, connect, and elaborate actions. Reflecting back on participants’ suggestions of classifying drugs, we suggest that instances of entities can be flexibly grouped based on some criteria to avoid manipulating one by one and instead do so by groups. In our case, mutations can be grouped by genes, and the drugs by approved drug, investigational drug, etc., as one participant proposed. In this way, the users can select entities by groups and look at their relations on a group level to facilitate pattern recognition, as also discussed in Section 5.6.3.

6.3 Insight-based evaluation

We have some considerations on evaluating whether a user’s open-ended exploration goes beyond his or her original goals, such as supporting serendipitous discovery. In this study, we asked participants to input tasks whenever they came up with new ones. This method may not adequately capture the directness versus unexpectedness of insights. Some participants tried to stick with the task they had input and wrote notes, whereas some others wrote new tasks inspired by exploration. For example, in case b of Section 5.6.3, the participant was inspired by a note and input a new task, from which we expected the insight generated to be originally recognized as “going beyond the original task.” This caused most insights to be graded as direct (Fig. 2).

On the other hand, input tasks in the middle of data exploration might interrupt the flow of one’s thinking. Saraiya et al. [58] asked

participants to identify all tasks before exploration. However, we ponder that during exploration, when participants become familiar with the tool, they may come up with new tasks from their knowledge databases. Guo et al. provided an open-ended task to all participants as they studied a generic visualization. This may bring two other challenges. First, there is the question of how open-ended the task should be to allow serendipitous discovery. Second, as proposed by North et al. [51] and Gomez et al. [23], participants need to have enough motivation to explore unfamiliar data and generate insight. We need to address these challenges when studying a generic visualization.

But with familiar data, in this study, we found that “curiosity” is another issue to entice users to go beyond their specific area and generate insight. We observed that some participants tried to focus on their specialized cancer types and did not explore other diseases, so they did not find new information with MediSyn. This phenomenon may also be attributed to the controlled study. An insight-based longitudinal study [59] in one’s natural working environment may reveal other interesting results, which would be worth future exploration.

7 IMPLICATIONS TO KNOWLEDGE-ASSISTED VISUALIZATION

Integrating domain knowledge in visualization, such as the user settings of visualization parameters, to support data exploration, also known as knowledge-assisted visualization [12], is trending but facing some challenges. This study sheds some light on addressing the challenges in the following ways.

First, this study suggests two possible ways to alleviate the difficulty of evaluating shared insights. Compared with visualizations for general content, domain-specific visualizations have limited numbers of domain experts spreading all over the world. Noises in shared user insights, such as comments without domain values, can be more difficult to discover and can conversely cause more serious consequences. The results of this study show, on the one hand, that visualizing insight provenance can help to explain the logic behind the insight. On the other hand, it is promising to qualify insights through the analysis of their interactions. As an example from the study, the drill-down pattern is shown to be positively correlated to the domain values of insights, whereas the pattern of select – write tended to be negatively correlated to it. This implies that the visualization tool can possibly predict the level of the domain values of insights by analyzing their interaction patterns. The study also showed that one insight characteristic can be related to multiple interaction types or patterns (Table 3). Thus, multiple correlations can be evaluated, which may amount to increased confidence in insight characterization. However, as stated in Section 5.6.1, interaction patterns can be specific to certain types of visualizations. Supporting insight characterization in various applications first requires the intervention of experts to learn the relations between insights and application-specific actions, which invites future exploration.

Moreover, qualifying insights by interactions can assist in personalized insight recommendation. As the study shows that elaboration actions tend to relate to the depth of insights, MediSyn can recommend in-depth insights to users who elaborate on similar entities, as they are potentially looking for this type of insights. Insight recommendation can be helpful when users encounter limitations with data exploration, as this analysis shows (Section 5.6.3).

8 CONCLUSIONS

This paper presents a case study that investigates *which interaction type or pattern relate to which aspect of insight quality*. First, we used the concept of entities to devise the interaction of a visualization tool—MediSyn—for insight generation. MediSyn synthesizes five drug-target datasets and supports five types of interactions: 1) selecting entities of interest, 2) connecting relevant entities, 3) elaborating by retrieving the details of entity relations, 4) exploring other entities, and 5) entity-based insight sharing. The action of entity-based insight sharing includes 1) using entities to link the visualization and public notes to support the bidirectional exploration of the two components; 2) allowing users to view insight provenance, which records semantically meaningful actions and their resulted visualization states to help rationalize public notes.

To investigate the research question and user strategies of deriving insight, we conducted a study with 14 participants. During the study, they were instructed to input a task, then freely explore the data with their task in mind and generate insights. We then had two raters grade four aspects, i.e., directness versus unexpectedness, correctness, breadth versus depth, and domain value, of the recorded insights. Apart from the five types of interactions that MediSyn supports, we extracted seven interaction patterns from logged interactions to analyze their relations to aspects of insight quality. Results showed the potential to qualify insights via interactions. Among other findings, exploration actions can lead to unexpected insights; the drill-down pattern positively correlates to the domain values of insights, which is also confirmed by a qualitative analysis. The qualitative analysis of user strategies also reveals that using domain knowledge to guide exploration can increase the domain value of derived insights. To alleviate the effects of domain knowledge on insight quality, we propose that future work on a generic type of visualization can further this research, though we should address user motivation issues by asking them to explore unfamiliar datasets and derive insight. Moreover, how user personality traits relate to the quality of insight is an interesting research question for future exploration.

These results also imply that future research on analyzing interaction types and patterns arriving at an insight can help assess the quality of insight for knowledge-assisted visualization. Opportunities exist for visualization to recommend insights to users who are looking for insights of a certain quality inferred by analyzing their interactions.

9 SUPPLEMENTAL MATERIAL

Supplemental video 1 is a video that provides a demonstration of the five interaction types that MediSyn supports and an overview of the user study. (MP4 67,409 kb)

ACKNOWLEDGMENTS

This work is supported by the Academy of Finland (grants no. 295504, and 286440).

REFERENCES

- [1] Intro.js. <https://introjs.com/>, 2019.
- [2] I. Alexandre. Promoting insight: A case study of how to incorporate interaction in existing data visualizations. In *20th International Conference Information Visualisation, IV 2016, Lisbon, Portugal, July 19-22, 2016*, pp. 203–208, 2016. doi: 10.1109/IV.2016.15
- [3] S. Alspaugh, N. Zokaei, A. Liu, C. Jin, and M. A. Hearst. Futzing and moseying: Interviews with professional data analysts on exploration practices. *IEEE Trans. Vis. Comput. Graph.*, 25(1):22–31, 2019. doi: 10.1109/TVCG.2018.2865040
- [4] L. Battle and J. Heer. Characterizing exploratory visual analysis: A literature review and evaluation of analytic provenance in tableau. *Comput. Graph. Forum*, 38(3):145–159, 2019. doi: 10.1111/cgf.13678
- [5] E. A. Bier, S. K. Card, and J. W. Bodnar. Entity-based collaboration tools for intelligence analysis. In *2008 IEEE Symposium on Visual Analytics Science and Technology*, pp. 99–106, Oct 2008. doi: 10.1109/VAST.2008.4677362
- [6] T. Blascheck, F. Beck, S. Baltes, T. Ertl, and D. Weiskopf. Visual analysis and coding of data-rich user behavior. In *2016 IEEE Conference on Visual Analytics Science and Technology, VAST 2016, Baltimore, MD, USA, October 23-28, 2016*, pp. 141–150, 2016. doi: 10.1109/VAST.2016.7883520
- [7] T. Blascheck, M. John, K. Kurzhals, S. Koch, and T. Ertl. VA²: A visual analytics approach for evaluating visual analytics applications. *IEEE Trans. Vis. Comput. Graph.*, 22(1):61–70, 2016. doi: 10.1109/TVCG.2015.2467871
- [8] T. Blascheck, L. M. Vermeulen, J. Vermeulen, C. Perin, W. Willett, T. Ertl, and S. Carpendale. Exploration strategies for discovery of interactivity in visualizations. *IEEE Trans. Vis. Comput. Graph.*, 25(2):1407–1420, 2019. doi: 10.1109/TVCG.2018.2802520
- [9] E. T. Brown, A. Ottley, H. Zhao, Q. Lin, R. Souvenir, A. Endert, and R. Chang. Finding waldo: Learning about users from their interactions. *IEEE Trans. Vis. Comput. Graph.*, 20(12):1663–1672, 2014. doi: 10.1109/TVCG.2014.2346575
- [10] S. P. Callahan, J. Freire, E. Santos, C. E. Scheidegger, C. T. Silva, and H. T. Vo. Vistrails: Visualization meets data management. In *Proceedings*

- of the 2006 ACM SIGMOD International Conference on Management of Data, SIGMOD '06, pp. 745–747. ACM, New York, NY, USA, 2006. doi: 10.1145/1142473.1142574
- [11] R. Chang, C. Ziemkiewicz, T. M. Green, and W. Ribarsky. Defining insight for visual analytics. *IEEE Computer Graphics and Applications*, 29(2):14–17, March 2009. doi: 10.1109/MCG.2009.22
- [12] M. Chen, D. Ebert, H. Hagen, R. S. Laramée, R. van Liere, K. Ma, W. Ribarsky, G. Scheuermann, and D. Silver. Data, information, and knowledge in visualization. *IEEE Computer Graphics and Applications*, 29(1):12–19, Jan 2009. doi: 10.1109/MCG.2009.6
- [13] E. K. Choe, B. Lee, and m. c. schraefel. Characterizing visualization insights from quantified selfers' personal data presentations. *IEEE Computer Graphics and Applications*, 35(4):28–37, 2015. doi: 10.1109/MCG.2015.51
- [14] Z. Cui, S. K. Badam, M. A. Yalçın, and N. Elmqvist. Datasite: Proactive visual data exploration with computation of insight-based recommendations. *Information Visualization*, 18(2), 2019. doi: 10.1177/1473871618806555
- [15] F. Dabek and J. J. Caban. A grammar-based approach for modeling user interactions and generating suggestions during the data exploration process. *IEEE Trans. Vis. Comput. Graph.*, 23(1):41–50, 2017. doi: 10.1109/TVCG.2016.2598471
- [16] W. Dou, D. H. Jeong, F. Stukes, W. Ribarsky, H. R. Lipford, and R. Chang. Recovering reasoning processes from user interactions. *IEEE Computer Graphics and Applications*, 29(3):52–61, 2009. doi: 10.1109/MCG.2009.49
- [17] W. Dou, W. Ribarsky, and R. Chang. Capturing reasoning process through user interaction. *Proc. IEEE EuroVAST*, 2, 2010.
- [18] C. Dunne, N. Henry Riche, B. Lee, R. Metoyer, and G. Robertson. Graph-trail: Analyzing large multivariate, heterogeneous networks while supporting exploration history. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, pp. 1663–1672. ACM, New York, NY, USA, 2012. doi: 10.1145/2207676.2208293
- [19] M. Drk, N. H. Riche, G. Ramos, and S. Dumais. Pivotpaths: Strolling through faceted information spaces. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2709–2718, Dec 2012. doi: 10.1109/TVCG.2012.252
- [20] O. ElTayeb and W. Dou. A survey on interaction log analysis for evaluating exploratory visualizations. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization*, BELIV '16, pp. 62–69. ACM, New York, NY, USA, 2016. doi: 10.1145/2993901.2993912
- [21] A. Endert, R. Chang, C. North, and M. X. Zhou. Semantic interaction: Coupling cognition and computation through usable interactive analytics. *IEEE Computer Graphics and Applications*, 35(4):94–99, 2015. doi: 10.1109/MCG.2015.91
- [22] M. Feng, E. M. Peck, and L. Harrison. Patterns and pace: Quantifying diverse exploration behavior with visualizations on the web. *IEEE Trans. Vis. Comput. Graph.*, 25(1):501–511, 2019. doi: 10.1109/TVCG.2018.2865117
- [23] S. R. Gomez, H. Guo, C. Ziemkiewicz, and D. H. Laidlaw. An insight-and task-based methodology for evaluating spatiotemporal visual analytics. In *Visual Analytics Science and Technology (VAST), 2014 IEEE Conference on*, pp. 63–72. IEEE, 2014.
- [24] D. Gotz and Z. Wen. Behavior-driven visualization recommendation. In *Proceedings of the 14th International Conference on Intelligent User Interfaces, IUI 2009, Sanibel Island, Florida, USA, February 8-11, 2009*, pp. 315–324, 2009. doi: 10.1145/1502650.1502695
- [25] D. Gotz and M. X. Zhou. Characterizing users' visual analytic activity for insight provenance. *Information Visualization*, 8(1):42–55, 2009. doi: 10.1057/ivs.2008.31
- [26] T. M. Green and B. D. Fisher. Towards the personal equation of interaction: The impact of personality factors on visual analytics interface interaction. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology, IEEE VAST 2010, Salt Lake City, Utah, USA, 24-29 October 2010, part of VisWeek 2010*, pp. 203–210, 2010. doi: 10.1109/VAST.2010.5653587
- [27] D. P. Groth and K. Streefkerk. Provenance and annotation for visual exploration systems. *IEEE Trans. Vis. Comput. Graph.*, 12(6):1500–1510, 2006. doi: 10.1109/TVCG.2006.101
- [28] H. Guo, S. R. Gomez, C. Ziemkiewicz, and D. H. Laidlaw. A case study using visualization interaction logs and insight metrics to understand how analysts arrive at insights. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):51–60, Jan 2016. doi: 10.1109/TVCG.2015.2467613
- [29] C. He, L. Micallef, Z.-u.-R. Tanoli, S. Kaski, T. Aittokallio, and G. Jacucci. Medisyn: uncertainty-aware visualization of multiple biomedical datasets to support drug treatment selection. *BMC Bioinformatics*, 18(10):393, Sep 2017. doi: 10.1186/s12859-017-1785-7
- [30] J. Heer, J. Mackinlay, C. Stolte, and M. Agrawala. Graphical histories for visualization: Supporting analysis, communication, and evaluation. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1189–1196, Nov 2008. doi: 10.1109/TVCG.2008.137
- [31] J. Heer and B. Shneiderman. Interactive dynamics for visual analysis. *Queue*, 10(2):30:30–30:55, Feb. 2012. doi: 10.1145/2133416.2146416
- [32] J. Heer, F. B. Viégas, and M. Wattenberg. Voyagers and voyeurs: Supporting asynchronous collaborative visualization. *Commun. ACM*, 52(1):87–97, Jan. 2009. doi: 10.1145/1435417.1435439
- [33] T. J. Jankun-Kelly, K. Ma, and M. Gertz. A model and framework for visualization exploration. *IEEE Trans. Vis. Comput. Graph.*, 13(2):357–369, 2007. doi: 10.1109/TVCG.2007.28
- [34] W. Javed and N. Elmqvist. Explates: Spatializing interactive analysis to scaffold visual exploration. *Comput. Graph. Forum*, 32(3):441–450, 2013. doi: 10.1111/cgf.12131
- [35] N. Kadivar, V. Y. Chen, D. Dunsmuir, E. Lee, C. Z. Qian, J. Dill, C. D. Shaw, and R. F. Woodbury. Capturing and supporting the analysis process. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology, IEEE VAST 2009, Atlantic City, New Jersey, USA, 11-16 October 2009, part of VisWeek 2009*, pp. 131–138, 2009. doi: 10.1109/VAST.2009.5333020
- [36] E. Kandogan and U. Engelke. Towards a unified representation of insight in human-in-the-loop analytics: A user study. In *Proceedings of the Workshop on Human-In-the-Loop Data Analytics, HILDA'18*, pp. 3:1–3:7. ACM, New York, NY, USA, 2018. doi: 10.1145/3209900.3209912
- [37] Y. Kim, N. H. Riche, B. Lee, M. Brehmer, M. Pahud, K. Hinckley, and J. Hullman. Inking your insights: Investigating digital externalization behaviors during data analysis. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces, ISS 2019, Daejeon, South Korea, November 10-13, 2019*, pp. 255–267, 2019. doi: 10.1145/3343055.3359714
- [38] K. Klouche, T. Ruotsalo, and G. Jacucci. From hyperlinks to hypercues: Entity-based affordances for fluid information exploration. In *Proceedings of the 2018 Designing Interactive Systems Conference, DIS '18*, pp. 401–411. ACM, New York, NY, USA, 2018. doi: 10.1145/3196709.3196775
- [39] H. R. Lipford, F. Stukes, W. Dou, M. E. Hawkins, and R. Chang. Helping users recall their reasoning process. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology, IEEE VAST 2010, Salt Lake City, Utah, USA, 24-29 October 2010, part of VisWeek 2010*, pp. 187–194, 2010. doi: 10.1109/VAST.2010.5653598
- [40] Z. Liu and J. T. Stasko. Mental models, visual reasoning and interaction in information visualization: A top-down perspective. *IEEE Trans. Vis. Comput. Graph.*, 16(6):999–1008, 2010. doi: 10.1109/TVCG.2010.177
- [41] Z. Liu, Y. Wang, M. Dontcheva, M. Hoffman, S. Walker, and A. Wilson. Patterns and sequences: Interactive exploration of clickstreams to understand common visitor paths. *IEEE Trans. Vis. Comput. Graph.*, 23(1):321–330, 2017. doi: 10.1109/TVCG.2016.2598797
- [42] J. Lu, Z. Wen, S. Pan, and J. Lai. Analytic trails: Supporting provenance, collaboration, and reuse for visual data analysis by business users. In P. Campos, N. Graham, J. Jorge, N. Nunes, P. Palanque, and M. Winckler, eds., *Human-Computer Interaction – INTERACT 2011*, pp. 256–273. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- [43] K. Madanagopal, E. D. Ragan, and P. Benjamin. Analytic provenance in practice: The role of provenance in real-world visualization and data analysis environments. *IEEE Computer Graphics and Applications*, pp. 1–1, 2019. doi: 10.1109/MCG.2019.2933419
- [44] A. Mathisen, T. Horak, C. N. Klokmoose, K. Grønbaek, and N. Elmqvist. Insideinsights: Integrating data-driven reporting in collaborative visual analytics. *Comput. Graph. Forum*, 38(3):649–661, 2019. doi: 10.1111/cgf.13717
- [45] E. Mayr, M. Smuc, and H. Risku. Many roads lead to rome: Mapping users' problem-solving strategies. *Information Visualization*, 10(3):232–247, 2011. doi: 10.1177/1473871611415987
- [46] P. H. Nguyen, K. Xu, A. Bardill, B. Salman, K. Herd, and B. L. W. Wong. Sensemap: Supporting browser-based online sensemaking through analytic provenance. In *2016 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 91–100, Oct 2016. doi: 10.1109/VAST.2016.7883515

- [47] P. H. Nguyen, K. Xu, A. Wheat, B. L. W. Wong, S. Attfield, and B. Fields. Sensepath: Understanding the sensemaking process through analytic provenance. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):41–50, Jan 2016. doi: 10.1109/TVCG.2015.2467611
- [48] P. H. Nguyen, K. Xu, and B. Wong. A survey of analytic provenance. *Middlesex University*, 2014.
- [49] C. North. Toward measuring visualization insight. *IEEE Computer Graphics and Applications*, 26(3):6–9, May 2006. doi: 10.1109/MCG.2006.70
- [50] C. North, R. Chang, A. Endert, W. Dou, R. May, B. Pike, and G. Fink. Analytic provenance: Process + interaction + insight. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '11, pp. 33–36. ACM, New York, NY, USA, 2011. doi: 10.1145/1979742.1979570
- [51] C. North, P. Saraiya, and K. Duca. A comparison of benchmark task and insight evaluation methods for information visualization. *Information Visualization*, 10(3):162–181, 2011. doi: 10.1177/1473871611415989
- [52] W. Pike, J. Bruce, B. Baddeley, D. Best, L. Franklin, R. May, D. Rice, R. Riensche, and K. Younkin. The scalable reasoning system: Lightweight visualization for distributed analytics. *Information Visualization*, 8(1):71–84, 2009. doi: 10.1057/ivs.2008.33
- [53] W. A. Pike, J. Stasko, R. Chang, and T. A. O’Connell. The science of interaction. *Information Visualization*, 8(4):263–274, 2009. doi: 10.1057/ivs.2009.22
- [54] 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.
- [55] M. Pohl, S. Wiltner, S. Miksch, W. Aigner, and A. Rind. Analysing interactivity in information visualisation. *KI*, 26(2):151–159, 2012. doi: 10.1007/s13218-012-0167-6
- [56] E. D. Ragan, A. Endert, J. Sanyal, and J. Chen. Characterizing provenance in visualization and data analysis: an organizational framework of provenance types and purposes. *IEEE transactions on visualization and computer graphics*, 22(1):31–40, 2016.
- [57] K. Reda, A. E. Johnson, M. E. Papka, and J. Leigh. Modeling and evaluating user behavior in exploratory visual analysis. *Information Visualization*, 15(4):325–339, 2016. doi: 10.1177/1473871616638546
- [58] P. Saraiya, C. North, and K. Duca. An insight-based methodology for evaluating bioinformatics visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 11(4):443–456, July 2005. doi: 10.1109/TVCG.2005.53
- [59] P. Saraiya, C. North, V. Lam, and K. A. Duca. An insight-based longitudinal study of visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 12(6):1511–1522, Nov 2006. doi: 10.1109/TVCG.2006.85
- [60] Y. B. Shrinivasan, D. Gotzy, and J. Lu. Connecting the dots in visual analysis. In *Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on*, pp. 123–130. IEEE, 2009.
- [61] Y. B. Shrinivasan and J. J. van Wijk. Supporting the analytical reasoning process in information visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pp. 1237–1246. ACM, New York, NY, USA, 2008. doi: 10.1145/1357054.1357247
- [62] M. Smuc, E. Mayr, T. Lammarsch, W. Aigner, S. Miksch, and G. J. rtner. To score or not to score? tripling insights for participatory design. *IEEE Computer Graphics and Applications*, 29(3):29–38, 2009. doi: 10.1109/MCG.2009.53
- [63] J. Stasko, C. Grg, and Z. Liu. Jigsaw: Supporting investigative analysis through interactive visualization. *Information Visualization*, 7(2):118–132, 2008. doi: 10.1057/palgrave.ivs.9500180
- [64] F. B. Viegas, M. Wattenberg, F. Van Ham, J. Kriss, and M. McKeon. Manyeyes: a site for visualization at internet scale. *IEEE transactions on visualization and computer graphics*, 13(6):1121–1128, 2007.
- [65] X. Wang, D. H. Jeong, W. Dou, S.-w. Lee, W. Ribarsky, and R. Chang. Defining and applying knowledge conversion processes to a visual analytics system. *Computers & Graphics*, 33(5):616–623, 2009.
- [66] W. Willett, S. Ginosar, A. Steinitz, B. Hartmann, and M. Agrawala. Identifying redundancy and exposing provenance in crowdsourced data analysis. *IEEE Trans. Vis. Comput. Graph.*, 19(12):2198–2206, 2013. doi: 10.1109/TVCG.2013.164
- [67] K. Xu, S. Attfield, T. J. Jankun-Kelly, A. Wheat, P. H. Nguyen, and N. Selvaraj. Analytic provenance for sensemaking: A research agenda. *IEEE Computer Graphics and Applications*, 35(3):56–64, May 2015. doi: 10.1109/MCG.2015.50
- [68] H. Yang, Y. Li, and M. X. Zhou. Understand users’ comprehension and preferences for composing information visualizations. *ACM Trans. Comput.-Hum. Interact.*, 21(1):6:1–6:30, 2014. doi: 10.1145/2541288
- [69] J. S. Yi, Y. a. Kang, and J. Stasko. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1224–1231, Nov.-Dec. 2008. doi: 10.1109/TVCG.2007.70515
- [70] J. S. Yi, Y.-a. Kang, J. T. Stasko, and J. A. Jacko. Understanding and characterizing insights: How do people gain insights using information visualization? In *Proceedings of the 2008 Workshop on BEyond Time and Errors: Novel evaluation Methods for Information Visualization*, BELIV '08, pp. 4:1–4:6. ACM, New York, NY, USA, 2008. doi: 10.1145/1377966.1377971
- [71] J. Zhao, M. Glueck, S. Breslav, F. Chevalier, and A. Khan. Annotation graphs: A graph-based visualization for meta-analysis of data based on user-authored annotations. *IEEE Trans. Vis. Comput. Graph.*, 23(1):261–270, 2017. doi: 10.1109/TVCG.2016.2598543
- [72] J. Zhao, M. Glueck, P. Isenberg, F. Chevalier, and A. Khan. Supporting handoff in asynchronous collaborative sensemaking using knowledge-transfer graphs. *IEEE Trans. Vis. Comput. Graph.*, 24(1):340–350, 2018. doi: 10.1109/TVCG.2017.2745279

Chen He is a PhD candidate in the Department of Computer Science at University of Helsinki. She received her B.Eng. in Digital Media Technology, Zhejiang University, and M.Sc in System Information Sciences, Tohoku University. Her research interests include the design and evaluation of interactive visualization for insight generation and sharing.

Luana Micallef (1986-2019) was an assistant professor at Human-Centered Computing, University of Copenhagen. She received her PhD in Computer Science from the University of Kent. Her research interests included information visualization, visual analytics, and human-computer interaction, particularly set visualization and design/layout optimization. She is remembered fondly by everyone of us.

Liye He is a PhD candidate at the Institute of Molecular Medicine Finland (FIMM), University of Helsinki. He received his M.Sc in Computer Science from the University of Helsinki. His research interests include computational biology, machine learning models for identifying synergistic drug combinations, and cancer research.

Gopal Peddinti received his M. Sc in Mathematics (University of Hyderabad, India), M. Sc in Computer Science (Indian Institute of Technology Kharagpur), and PhD in Computational Systems Biology (Aalto University). He is a senior scientist in computational biology at VTT Technical Research Center of Finland Ltd. His research interests include the integration and analysis of heterogeneous biological data.

Tero Aittokallio received his PhD degree in Applied Mathematics from the University of Turku. He is an EMBL Group Leader at FIMM, where his computational systems medicine group focuses on network-centric and machine learning-based approaches to modeling and predicting complex relationships between genetic dependencies and medical phenotypes, such as susceptibility to diseases and responses to treatments.

He is also a professor of statistics and applied mathematics at the University of Turku, with research interests in biostatistics, data mining and predictive modeling in biomedical applications.

Prof. Dr. Giulio Jacucci is a professor at the Department of Computer Science at the University of Helsinki. His research field and competencies are in human-computer interaction in particular: multi-modal interaction, interactive intent modeling, mixed reality providing innovation for applications in information discovery, and e-health and wellbeing. He is the co-founder of Etsimo Healthcare Ltd., offering a platform leveraging artificial intelligence on top of health

data, making it possible for health care providers to instantly offer their customers preventive health care. He co-authored two patents in the area of information seeking and on modular screens.

Fig. A2. Thirty-seven interaction trails in 37 rows. Rows are aggregated by participants and attached with participants' indices. The rows of each participant are ordered chronologically.