Extending an Association Map to Handle Large Data Sets

Tamara Babaian, Wendy Lucas, Alina Chircu and Noreen Power

Bentley University, Waltham, MA 02452, USA
`{tbabaian, wlucas, achircu, npower}@bentley.edu`

Abstract. This paper presents the Association Map-Large (AM-L), an interactive visualization of entity associations. AM-L is an extension of a previously reported AM interface that has been enhanced with search and interaction features for supporting larger data sets. We report on a user study with thirty two participants, which assesses user performance and experience with AM-L versus a tabular representation of the same data in the context of an enterprise system. Participants with varying levels of experience were given both simple and complicated tasks to complete with each system. Results indicate greater enjoyment and lower levels of mental effort when using AM-L, as well as less time spent on average when performing tasks. Accuracy results in terms of correctness indicate a learning curve, with overall performance worse with AM-L on the first two simple questions and first complex question, but then as well or better on subsequent questions. Given that the AM-L interface is unlike any with which the users had prior experience, it is not surprising that some exposure to the interface, such as training, would be helpful prior to use. Suggestions from the participants will inform future enhancements to the interface, which will be validated with further studies.

1   Introduction

Interactive data visualizations are becoming more and more common in the news media but have not yet entered the mainstream of decision support systems in the workplace. In the context of enterprise information systems, tables remain the prevalent format for reporting and aiding in decision making processes. Studies show that users perceive visualizations as useful ways to reduce the complexities associated with enterprise information systems [1], yet investigating what kinds of visualizations would work well in practice is an open research area.

In our previous work, we presented an interactive visualization called the Association Map (AM) [2], which was designed to assist enterprise system users in exploring associative relationships between two or three entities. Such exploration occurs in a wide variety of contexts, such as when looking for doctors who treated a patient for a particular problem, or picking an alternative vendor for a material that is used in one or more plants. The AM visualization enables users to easily find connections from a single entity (e.g., view all doctors seen by a specific patient and all health issues for

that patient). It also supports searching for connections between two or three specified entities and then highlighting all connections that are found. The results of a side-by-side evaluation of the AM interface with two table-based interfaces presenting similar association data are reported in [3]. In terms of task completion time and accuracy, user performance with the AM interface was superior to that achieved by the same users utilizing a standard SAP report. In another comparison, in which a customized version of an SAP report was used that showed the data in a simplified, three-column format, the results did not show any significant difference in user performance. Both of these comparisons, however, revealed the users' overwhelming preference for the AM.

The version of the AM used in the studies described above was designed for small data sets, limiting its use in practice. Given the users' performance and enthusiastic responses regarding their experience with the AM visualization [3] and the acknowledged difficulties in working with enterprise systems (see [4-5] for example), it is valuable to further develop, fine tune and evaluate visual interactive interfaces for performing enterprise system tasks.

In this paper, we will present a new version of the AM designed to enable exploration of large data sets. The original AM interface required that all data items be displayed in sufficiently large font for the user to see while fitting on a single screen. The new version provides a way to work with large data sets by employing zooming to highlight selected items, while still keeping the visualization to one screen. To facilitate easy observation of the selected items, the selected items from the middle column are also moved to the center of that column (as shown in Figure 5). Finally, a partial match search feature was added to facilitate searching. Displaying the same set of data in a tabular format would require multiple pages and the use of a scrolling mechanism in order to locate the needed records.

We have conducted a pilot evaluation involving 32 subjects performing tasks with varying levels of complexity using the new version of the AM, called the Association Map-Large (AM-L), versus a simplified, three-column table from SAP. The pilot revealed that, while the users were much more satisfied with AM-L compared to the table format, there was a learning curve associated with this unfamiliar interface. Initially, performance in terms of the correctness of responses with AM-L lagged behind the correctness of responses with the table when performing the first two simple tasks and the first complex task. After completing those initial tasks, however, user performance with AM-L matched or exceeded performance with the SAP table on both simple and complex tasks.

The rest of this paper is organized as follows: after presenting related work and the previous version of AM in Section 2, we describe the functionality of AM-L in Section 3. Section 4 outlines the user study setup for evaluating AM-L and comparing it to a tabular SAP interface. This is followed by a presentation of the findings from the study in Section 5. We then discuss lessons learned and directions for further fine-tuning of the interface in Section 6, followed by our conclusions in Section 7.

2   Related Work

 In the introduction, we have already commented on the studies in enterprise system usability that provided the impetus to our work in developing novel interactive visualizations for common workplace tasks. While interactive visualizations have recently made their way into many online media outlets, there are only a few examples of developed systems employing such visual interfaces (e.g.[6-8]). Research on the efficacy of visual representations and techniques in the context of specific applications is scarce. Dilla and Raschke [9] present a review of existing data visualization applications for financial fraud investigation and detection and outline a research agenda for investigating the key factors in the effectiveness of such visualizations. Our work presented here is, generally-speaking, a step in that direction: although we do not situate the visualization in the context of any particular application, we aim to understand what representations and features contribute to effective use of a visualization for problem solving that is traditionally supported by a variety of tabular representations.

Visualization evaluation is different from HCI usability evaluation, because visualizations are often credited for enabling discoveries and insights into data - notions that are not captured by usability. The issue is further complicated given that these benefits are often achieved after long-term use and exploration and thus are difficult to establish via a simple user study or even an expert evalutaion. The challenges in evaluating visualizations gave rise to a special workshop (BELIV) dedicated to this topic. Saket, Endert, and Stasko [10] present a review of a number of recently introduced visualization evaluations that focus on memorability, engagement, enjoyment and fun. In the work we present here, in which we seek to create visualizations in service of specific user tasks, we combine the traditional metrics of effectiveness with an evaluation of enjoyment [11].

2.1   Background

The previous version of the AM interface, called AM-N, consists of a circular display in which three types of entities are arranged in sorted order as labeled circles on the left and right sides of the diagram and as rectangles along the vertical center (see Figures 1-3). The items of each entity type are shown in different colors. Connections between the items indicate the presence of a relationship. Interactive features of the diagram include:

- A search interface, consisting of three fields for searching on the left, the right, or the center entity item, respectively.
- A point-and-click interface for each item. Clicking on a circle reveals connections between that item and the two other entities in a ternary relationship. Clicking on a rectangle reveals binary relationships between the item in the center and items on the left and on the right.
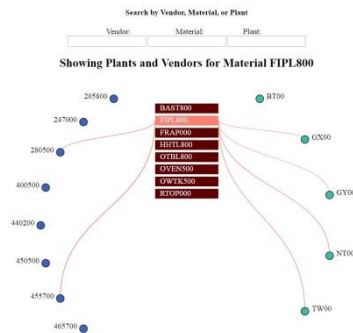
Thus, the AM interface can be used to explore relationships between items.
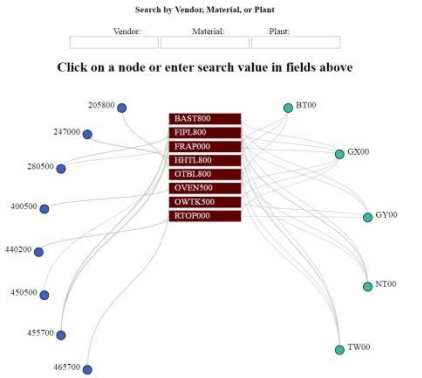
Fig. 1. AM-N interface showing all associations (from [3])
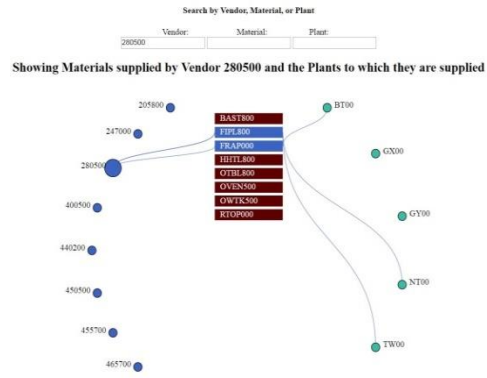
Fig. 2. AM-N interface showing associations for the selected Vendor on the left (from [3]).

Fig. 3. AM-N interface showing associations for the selected Material in the middle (from [3]).

For example, consider an AM depicting associations between Vendors, Materials and Plants, as shown in Figures 1-3. The various states shown in these figures present the following information:

1. When no item is selected, all connections between items are shown in grey (Fig. 1).
2. Selecting a specific Vendor reveals the connections between that Vendor, the Materials it supplies, and the Plants that are being supplied with those Materials by that Vendor. All items and connections are displayed in the same color to indicate this ternary relationship (Fig. 2).
3. Selecting a specific Plant is similar to selecting a Vendor, though a different color is used.
4. Selecting a specific Material reveals connections between that Material, the Vendors that supply it, and the Plants that use it (Fig. 3).
5. Specifying two of three entity keywords in the search interface will result in highlighting the items and connections in the ternary relationship matching the search terms. For instance, a search for Vendor A and Material B will highlight connections between Vendor A, Material B (if, indeed, A supplies B), and those Plants

that receive B from A. Other Materials supplied by A and other Vendors supplying B will not be highlighted by this search.

6. Similarly, a search specifying all three entity keywords results in the highlighting of the connection between the matching entities of three different types, if that connection exists in the data.
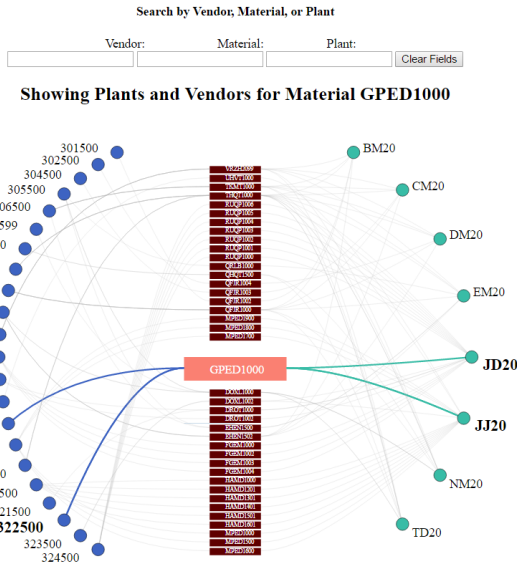

3 AM-L

Following the encouraging results of our initial exploratory user study of the AM [3], we have developed a new version of the interface, AM-L, to improve the search capabilities and accommodate the need for working with a larger set of entity instances. Compared to the version of the AM interface outlined in Section 2.1, AM-L introduces several features motivated by the need to extend AM's applicability to handling larger data sets. In this version, we have modified the presentation and interaction features for the entity displayed in the center (corresponding to Materials in the instance of the AM-L presented here). To handle the situation where the set of materials is so large that the items cannot be represented using a font of a readable size, we added the following three features:

- When a mouse points to the material, the rectangle representing the material is highlighted in a different color and enlarged; furthermore, the highlighted material's label is shown in large font, to the left of the material box (Fig. 4).
- A partial search feature has been added to aid locating an item in a large set. When a user enters a search keyword in any of the search boxes on top of the page, all items of the specified type are highlighted in yellow. A partial search on the center column (Material) highlights the matching items in yellow and also enlarges and moves those items into the center of the diagram (Fig. 5).
- Any time the items in the central column are involved in a selection result, they are enlarged and moved to the center, as shown in Figures 5 and 6.

To improve the readability of the diagram and better reflect the semantics of the data in the visual features, we made the following changes to the way the diagram reacts to selection:

- For the items displayed as circles, the selected circle is enlarged and the item label is highlighted in boldface font. Figure 7 demonstrates the selection of a Vendor.
- When an item is selected from the central column (Material), we use different colors for the connections between the selected Material and Plants versus the connections from the Material to the Vendors to emphasize that the connection lines visualize two different binary relationships in which the selected Material is involved, and not a ternary relationship (Fig. 6).

Fig. 4. AM-L interface with user pointing at an item in the central column highlights the item and shows an enlarged label on the left

Fig. 5. AM-L interface showing results of a partial search for the central column. Materials matching the specified keyword 'HA' are highlighted, and shown in large in the center

Fig. 6. AM-L interface showing results of selection of the central item (Material). The connections to items on left and on right sides are displayed in different colors
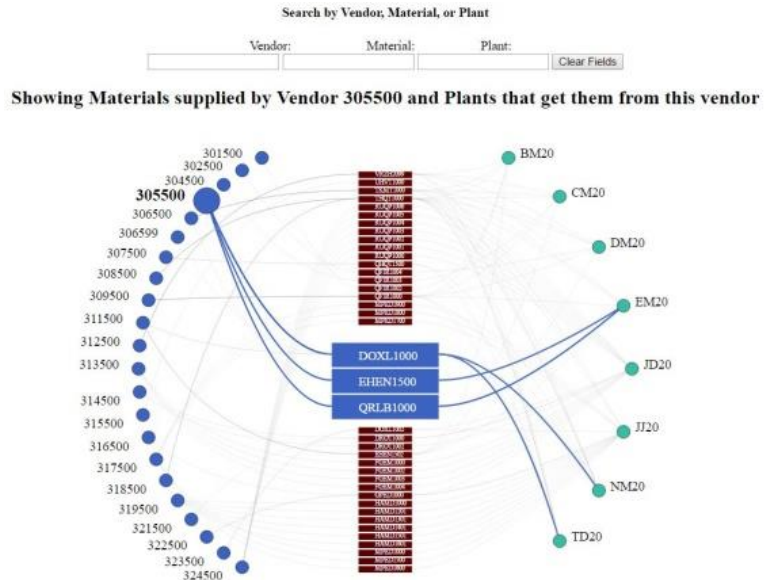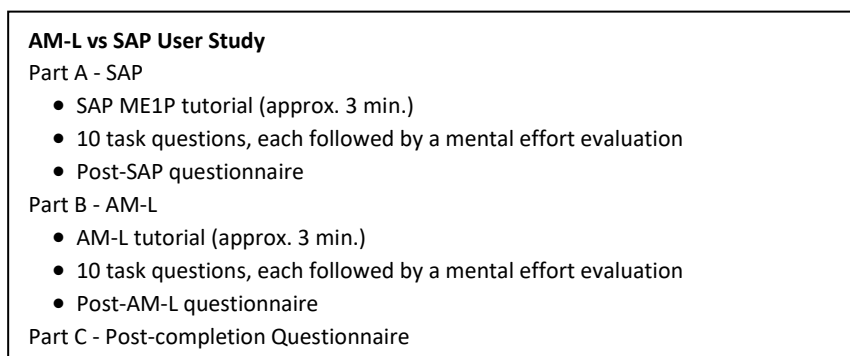

Fig. 7. AM-L interface showing results of selection of a left-side item

4 User Study Set-up

We conducted a user study to collect user feedback regarding the AM-L interface and to compare AM-L with a tabular SAP user interface in terms of user performance, satisfaction, and perceptions of mental effort involved in solving the same problems using the different interfaces. Thirty two students in graduate-level courses in a business university took part in this evaluation; participating in the user study was part of their required coursework. Two participants had over five years experience with SAP, 26 were novices with fewer than two months of experience with SAP, and the other four used it less than 6 months. All of them were using AM-L for the first time.

The study involved three parts (see Figure 8). Parts A (SAP) and B (using AM-L) each consisted of watching a short tutorial video of the appropriate interface, followed by ten task questions, and concluded with a questionnaire on the properties of the interface. After each task question, we asked respondents to evaluate the degree of perceived mental effort required to answer the question. The first six task questions were internally designated as simple; the last four were designated as complex The simple questions required examining associations of only one entity instance with other entities, for example: *List all vendors that supply material(s) to Plant SD00. Enter 'none' if you don't find any such vendors*. Complex questions required review and analysis of associations of two or more entity instances to come up with the answer, for example, *Plants AL00 and SD00 order the same materials. List all vendors that supply a material to one of these two plants, but not to the other. Enter 'none' if you don't find such vendors.* We recorded the participants' answers and the time it took to answer each question. Questions in Parts A and B were identical in terms of structure but had data with different labels.

Participants were randomly split into two groups, with the first group performing Part A before Part B (15 participants), and the second group doing Part B first (17 participants). Both groups (SAP-first and AM-L-first) concluded by filling-out the Post-completion Questionnaire (Part C), which asked users about their experience and perceptions with SAP and AM-L interfaces.

AM-L vs SAP User Study

Part A - SAP
- SAP ME1P tutorial (approx. 3 min.)
- 10 task questions, each followed by a mental effort evaluation
- Post-SAP questionnaire

Part B - AM-L
- AM-L tutorial (approx. 3 min.)
- 10 task questions, each followed by a mental effort evaluation
- Post-AM-L questionnaire

Part C - Post-completion Questionnaire

Fig. 8. Components of the user study

5   Analysis of the Results

In this section we present our analysis of user performance, self-reported mental effort, perceived interface complexity, as well as satisfaction and enjoyment from using the interface (Section 5.1). The Post-completion Questionnaire included open-response answers regarding user preferences and improvement suggestions for AM-L, summarized in Section 5.2.

5.1   Quantitative Findings

We analyzed user performance using the common metrics of accuracy and time on task. The users answered six simple questions, referred to as s1 through s6, and four complex questions, labeled c1 through c4. We describe the accuracy of their answers first.

Accuracy. Figures 9-11 depict the average percentage of correct answers per each question in the SAP-first group, AM-L-first group and overall. Table 1 presents the ranges of accuracy values for each question within each group and for each interface. The accuracy results depicted in Figures 9-11 suggest there might be a learning curve associated with using the AM-L visualization: notice that in both groups, the percentage of people that produced correct answers with AM-L rapidly increases from the first simple question (s1) to the third simple question (s3), thereby surpassing the accuracy with the SAP interface. The accuracy drops however at the first complex question (c1), and then picks up again and is slightly above the corresponding SAP numbers overall, with the exceptions of s1, s2, and c1 (Figure 11).

The average percentage of correctly answered questions overall was 78% for SAP versus 76% for AM-L. If we consider per-question accuracy (see Figure 11) and conjecture that the low accuracy results on questions s1, s2 and c1 are due to a learning curve associated with the use of a completely new interface, then the higher-or-comparable average accuracy with AM-L over SAP on all other questions is encouraging, as low initial accuracy may be overcome with training.

Table 1. Ranges of accuracy of responses per question in each group for each interface

Accuracy (% correct answers per question) ranges		
	SAP	AM-L
SAP-first	30%-100% (Avg 79%)	10%-100% (Avg 75%)
AM-L-first	40%-90% (Avg 76%)	50%-100% (Avg 77%)
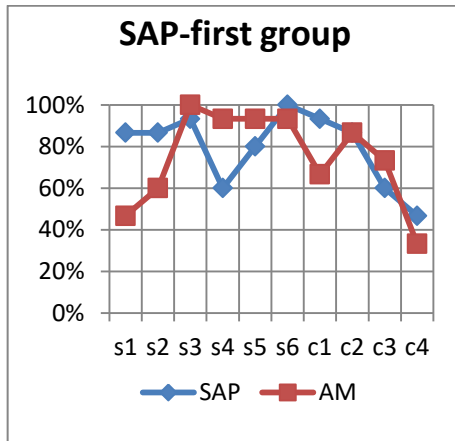
Fig. 9. Percentage correct answers in the SAP-first group
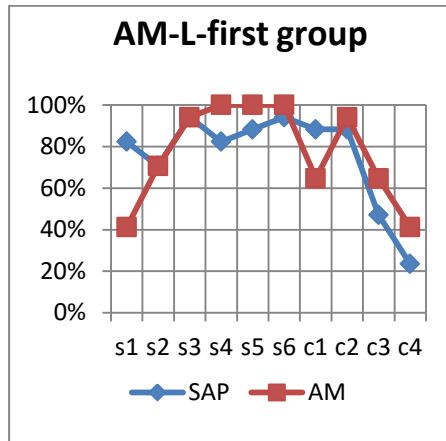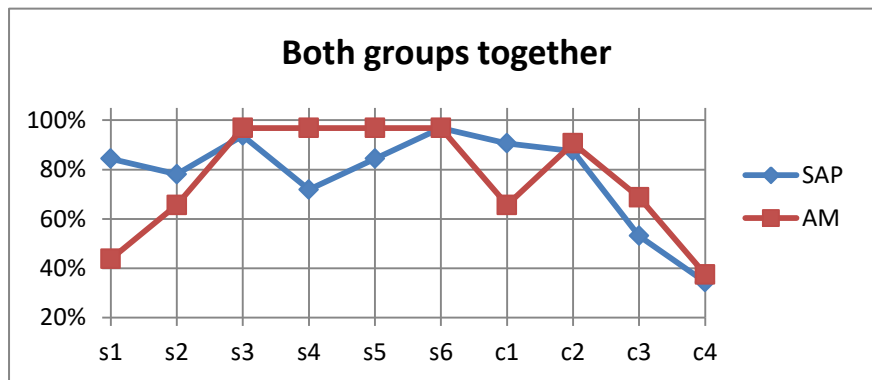

Fig. 10. Percentage of correct answers in the AM-L-first group

Fig. 11. Overall percentage of correct answers

Time-on-task. Time-on-task measures are presented in Figures 12-14. The average time to complete the task is greater for the SAP interface for all questions except question c2 in the SAP-first group of participants. Notably, the steep jump in AM-L time on question c2 (Figure 13, AM-L-first group) is due to one participant, who spent over forty five minutes on a task that took other users in this group less than two minutes on average. There is a similar outlier in the SAP-first group (Figure 12) that explains the jump in the SAP time on question c4: one user took over 20 minutes, while others spent less than four minutes on average.

Overall, the average total time to complete all 10 questions was higher when SAP was used. When considering performance within each group (SAP-first or AM-L-first) on both interfaces, given that the structure of questions is the same, it is reasonable to expect that the time spent with the second interface will be reduced, since the general strategy for answering the question has already been developed while solving, essentially, an identical task using the other interface first. This explains the relatively

closer times between the two interfaces in Figure 13 compared to Figure 12. Overall, users are faster with the AM-L interface, and the time difference between AM-L and SAP is greater when the participants have already solved the problem using SAP. In terms of the average total time per participant, in the SAP-first group, participants spent 27 minutes using SAP and 14 minutes using the AM-L interface. In the AM-first group the averages are 20 and 18 minutes, respectively.

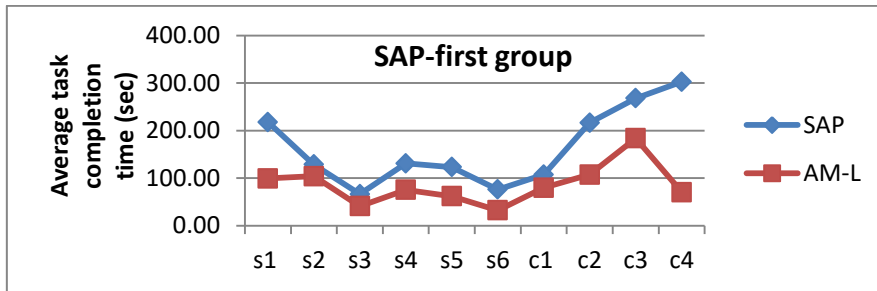**Fig. 12.** Average task completion time per question in SAP-first group
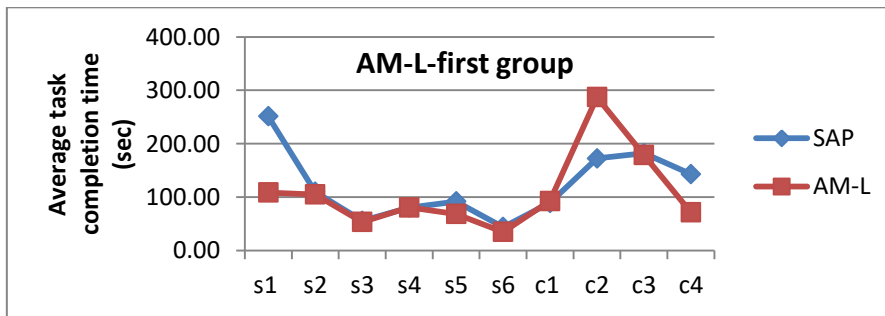

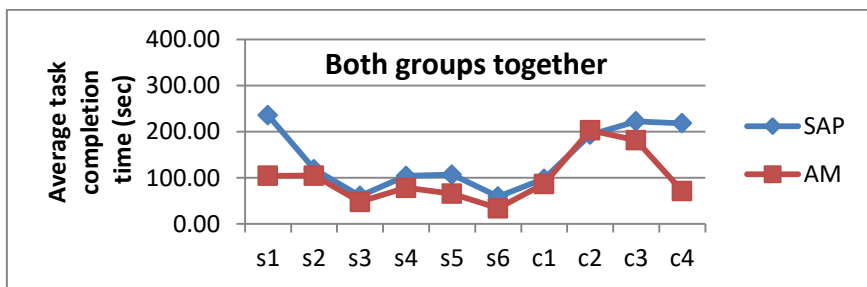**Fig. 13.** Average task completion time per question in AM-L-first group

Fig. 14. Overall average task completion time per question.

Paired sample (repeated measures) t-tests indicated several significant differences between the time-on-task metrics for questions completed with the SAP versus the

AM-L interfaces. In the SAP-first group, most simple questions (s1, s4, s5, and s6), most complex questions (c2, c3, c4), as well as all simple questions taken together (s1-s6), all complex questions taken together (c1-c4), and all questions taken together (s1-s6 and c1-c4) had significantly lower time-on-task when using the AM-L interface (with significance ranging from 0.1 to 0.01). In the AM-L first group, one third of the simple questions (s1 and s6) and one fourth of the complex questions (c4), as well as all simple questions taken together (s1-s6) had significantly lower time-on-task when using the AM-L interface (with significance ranging from 0.1 to 0.01). Taken together, these results indicate that the AM-L interface requires less time than the SAP interface for completing both simple and complex tasks.

Mental effort. To obtain additional insights regarding the benefits of each interface, we also asked respondents to rate the mental effort required to complete each task on a scale from 1 (low) to 9 (high). The average mental effort values per question in both SAP-first and AM-first groups, depicted in Figure 15, indicate lower effort for AM-L. The ranges of average reported values for simple and complex questions within each group and with each interface are presented in Table 2.

Table 2. Ranges of average mental effort values per simple (s1-s6) and complex (c1-c4) questions within each group for each interface.

Self-reported mental effort ranges

	SAP		AM-L	
	s1-s6	c1-c4	s1-s6	c1-c4
SAP-first	2.67 - 3.73	3.13 -5.53	2.27 - 2.80	2.60 - 3.87
AM-L-first	3.06 - 4.82	3.29 - 5	1.88 - 2.76	2.88- 4.35

Paired sample (repeated measures) t-tests indicated significant differences between the mental effort perceptions for the SAP and AM-L interfaces (with significance ranging from 0.01 to 0.001) for several tasks (both simple and complex) in both groups as well. Taken together, these results suggest that the AM-L interface requires less mental effort than the SAP interface for several tasks with varying complexity.
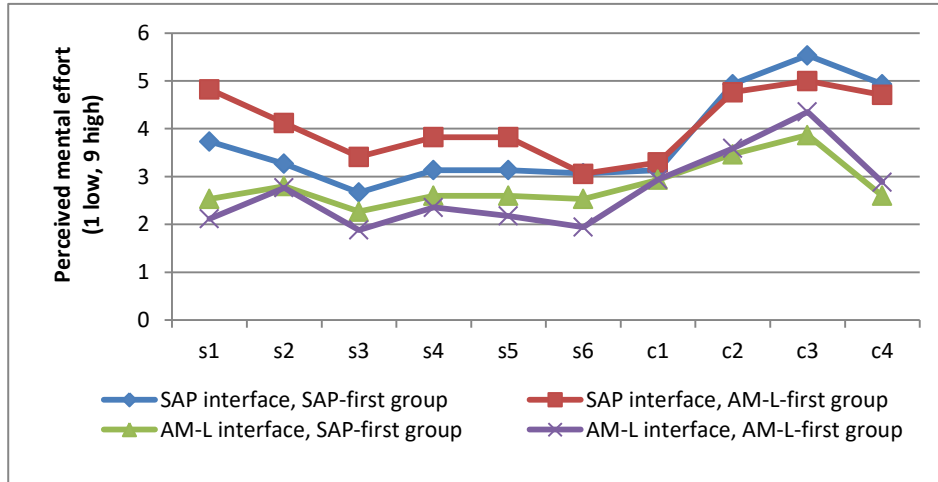
Fig. 15. Perceived mental effort for each interface, by question

Search and Sort. We further analyzed user perceptions of the interface features by asking participants to rate the usefulness of the sorting and search capabilities provided by each interface on a scale from 1 to 7, or indicate that they were not aware of the specific capability. Questionnaire responses show users rating the search features in both SAP and AM-L interfaces highly (see Table 3). According to the SAP logs, none of the students used the sorting feature, which was demonstrated in the tutorial, while the majority (28 out of 32) used the search feature. Search is essential to operating with AM-L; the items for each entity are always presented in a sorted order.

Table 3. Average value of usefulness of search and sort features assigned by users who reported being aware of the feature.

Search and Sort usefulness values (1 lowest - 7 highest)

	SAP		AM-L	
	search	sort	search	sort
SAP-first 15 users	5.8 (2 unaware)	1.3 (7 unaware)	6.4 (1 unaware)	5.4 (1 unaware)
AM-L-first 17 users	5.8 (1 unaware)	2.6 (2 unaware)	5.9 (1 unaware)	5.1 (1 unaware)

Perceived interface complexity and enjoyment. Last, but not least, we analyzed the user perceptions of each interface using perceived interface complexity and perceived enjoyment metrics adapted from previous research studies (see Figure 16 for question details and Figure 17 for metrics summary).

13

Perceived Interface Complexity
(answers on a 7-point Likert scale ranging from 1 strongly disagree to 7 strongly agree)
PIC_1: The interface was complex.
PIC_2: The interface was crowded.
PIC_3: The interface was interactive. *(reverse coded)*
PIC_4: The interface displays lots of information.

Perceived Enjoyment
PE_1: Using the interface was: 1 Unexciting .... 7 Exciting
PE_2: Using the interface was: 1 Dull.. 7 Neat
PE_3: Using the interface was: 1 Not Fun .... 7 Fun
PE_4: Using the interface was: 1 Unappealing .... 7 Appealing
PE_5: Using the interface was: 1 Boring .... 7 Interesting

Fig. 16. Questions on perceived interface complexity (adapted from [12]) and perceived enjoyment (adapted from [11]) regarding the AM-L and SAP interfaces
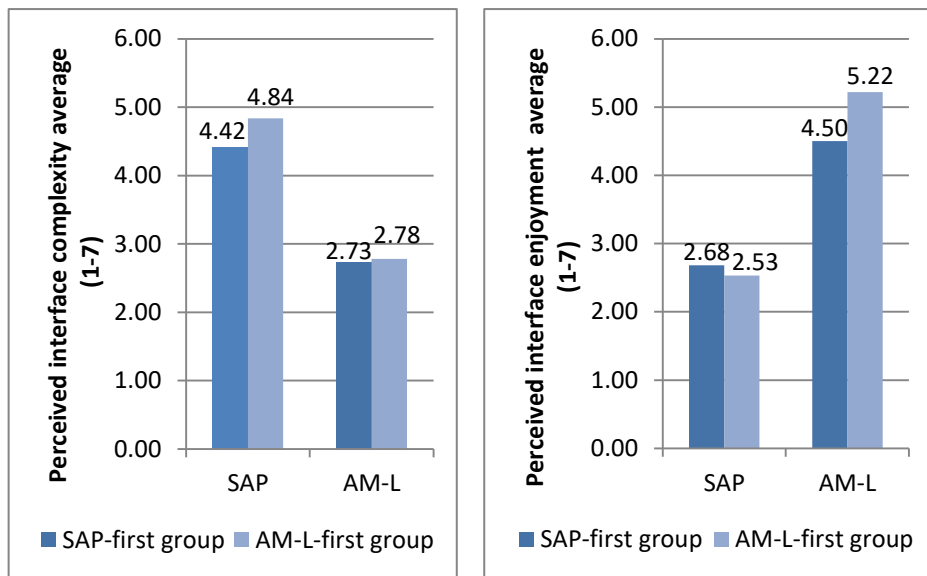
Fig. 17. Perceived interface complexity and enjoyment averages for each interface and in each group.

As shown in Figure 17, the average perception of complexity, measured on a 7-point scale, for SAP's tabular interface is more than one point higher than that of AM-L, regardless of which interface is used first. The average enjoyment score for AM-L is more than one point higher in both groups, and more than two points higher for users

who experienced SAP first. Tables 4 and 5 show the ranges of reported measures of complexity and enjoyment within each group and for each interface.

A paired sample (repeated measures) t-test indicated a significant difference between the composite interface complexity (the average of all four questions) for SAP and AM-L in both groups (at 0.001 significance level). A similar test for enjoyment also indicated a significant difference between the composite enjoyment (the average of all five questions) for SAP and AM-L (at 0.01 significance level in the SAP-first group and 0.001 significance level in the AM-L-first group). Therefore, users clearly perceive AM-L as being less complex and more enjoyable than SAP.

Table 4. Ranges of average reported interface complexity values on a 7-point scale (lower means less complex) per each group for each interface.

Perceived interface complexity (1-lowest - 7 highest)

	SAP	AM-L
SAP-first	4.27 - 4.47 (Avg 4.42)	1.80 - 3.53 (Avg 2.73)
AM-L first	4.27 - 5.12 (Avg 4.84)	2.00 - 3.71 (Avg 2.78)

Table 5. Ranges of average reported values of enjoyment of interface on a 7-point scale (lower means less enjoyable) per each group for each interface.

Interface enjoyment (1-lowest - 7 highest)

	SAP	AM-L
SAP-first	2.40 - 2.87 (Avg 2.68)	4.13 - 4.80 (Avg 4.50)
AM-L first	2.29 - 2.65 (Avg 2.53)	4.82 - 5.59 (Avg 5.22)

5.2 Qualitative Findings

The final portion of the evaluation asked participants to answer the five questions shown in Figure 18. Question 2 was the only one that was not open-ended, with possible answers of AM-L, SAP, or Undecided. In those cases where it was relevant, we distinguish between the responses provided by the 17 participants who used the AM-L interface first and the 15 who used the SAP interface first in performing the required tasks.

1. How would you compare the Association Map to the way of finding the same information in SAP?
2. Assuming you have access to both the Association Map and the SAP interface, which one do you intend to use next time you have to perform similar tasks?
3. Please explain your answer to the above question regarding which interface you intend to use.
4. Do you have any suggestions for improving the Association Map?
5. Do you have any suggestions for improving the SAP interface?

Fig. 18. Open-ended questions on AM-L and SAP interfaces

Question 1. All of the participants who used the AM-L interface first and 12 of the 15 participants who used SAP first found AM-L to be superior to SAP for finding the same information. The most commonly used word was "easier," as is evident from the word cloud of responses shown in Figure 19. This term was always used in reference to the AM-L interface. Participants also found AM-L to be more intuitive and user friendly. Of the three who did not find AM-L to outperform SAP, one commented it was "Difficult, unexiting [*sic*];" another preferred the visual aspect of AM-L but found it to be limited in functionality for defining the search criteria; and the third found AM-L to be more intuitive but still lacking in an easy way to extract results.

Fig. 19. Word cloud derived from responses to Question 1

Question 2. When asked which interface the participant would use for performing tasks similar to those in the study, assuming access to both AM-L and SAP, most participants chose AM-L. The responses to this question are shown in Table 6.

Table 6. Intention to use – number of users preferring each designated interface over another.

Intention to use

	SAP	AM-L	Undecided
SAP-first	1	12	2
AM-L first	0	16	1
Total	1	28	3

Question 3. The reasons for the interface choices in question 2 were very similar to the answers given to question 1 in comparing the interfaces. The primary reason behind the intent to use AM-L over SAP was related to it being easier to use: it was easier to search, sort, learn, and view relationships. Several participants also noted that the visual interface required less typing, was quicker to use, and was aesthetically appealing. The one participant who preferred SAP stated that was due to being "a little familiar with it." The participant who used AM-L first and was undecided thought SAP might be faster for navigating certain functions but would give a slight edge to AM-L. The other two undecideds, both of whom used SAP first, felt that AM-L would be better for simpler tasks involving less data, but SAP could provide more information/better output for complex tasks.

Question 4. All but three of the participants (one of whom used AM-L first and two who used SAP first) had suggestions for improving AM-L. The most commonly requested improvement was to support searches on multiple values for each parameter. Several participants also offered suggestions for improving the clarity of the display, particularly for large datasets. These included more color coding, with the meaning for each color clearly identified; enlarging and displaying in bold the data related to a selection, not just the selected value; providing zoom functionality; using straight rather than curved lines for connectors; improving the alignment of data identifiers with the circles representing them; and having higher order sorting capabilities so that users could limit the scope of what is shown (such as by product family or category). Some participants would also like better support for copy and paste, with one noting that it is hard to copy from the middle column since the display is refreshed when a material is clicked on. Other suggestions that would lessen the need for a pad and pencil would be having the ability to drag and drop to a scratch pad or to export results from a search.

Question 5. Twenty-two participants offered suggestions or areas for improving SAP, seven expressed overall dissatisfaction with the interface, and three simply responded "no." The two most frequently cited areas for improvement were related to searching and sorting. The ability to search on multiple fields and for multiple values per field was a common suggestion. Several participants also suggested improving the visibility and ease of use of the sort functionality; some participants stated that their interface looked different from that in the tutorial, and they were unable to figure out how to sort the data. Another common complaint was that SAP is not intuitive. Suggestions for improvements in that area included getting rid of transaction codes; having icons on buttons that convey meaning; displaying relationships in visual rather than tabular formats; and using colors to differentiate the data.

6   Discussion

The user study revealed a number of interesting findings concerning the use of AM-L versus SAP. One such finding is that the correctness of user responses increased with a bit of practice. For the six simple questions, users did worse with AM-L than with SAP on the first two questions but then did better with AM-L on the remaining four. For the four complex questions, it was only in answering the first one that users fared worse with AM-L. Evidence of a learning curve is not surprising, in that AM-L is likely very different in appearance from other interfaces with which the participants might have prior experience. This issue could be mitigated with time and/or training.

In terms of task completion time, users answered questions more quickly with AM-L. The timing data also hints at a learning curve with SAP, as the time to complete the simple questions fell sharply after answering the first one. The difference in completion time between AM-L and SAP was larger for those who used SAP first. It is likely that lessons learned from using SAP carried over to the use of AM-L.

Overall, mental effort was rated lower when using AM-L, with the highest ratings on mental effort typically from users of SAP who had used AM-L first. AM-L highlights relationships between entities in ways that SAP does not. After becoming accustomed to seeing those relationships represented visually, it stands to reason that having to work through similar relationships presented in a tabular format would present more of a mental challenge. This is particularly true in the absence of a sorting feature, which can be very useful in working out relationships; in this study, none of the participants used SAP's sorting feature because they were either unaware of it or could not find it.

Participants also perceived AM-L to be less complex and more enjoyable to use than SAP. Their responses to open-ended questions about the interfaces provided additional evidence of the clear preference by the majority of participants for AM-L. Equally encouraging was that 24 of 28 of them would choose to use AM-L over SAP if given access to it for performing tasks similar to those in the study.

Most of the improvements suggested for AM-L focused on extending the search functionality, followed by improving the clarity of the display. By contrast, participants often viewed SAP as outdated technology with inadequate support features, particularly in the areas of searching and sorting. Their recommended improvements typically revealed a clear preference for interactive visualizations over tabular data formats for use in performing data-intensive tasks.

Limitations. Given the small sample size in our study, we can only conjecture at this time that the results are generally representative. In addition, the tasks and data used in the study, which were generated by an expert, may not be indicative of actual data and tasks performed in the workplace. Our study participants were students rather than real-life users. Lastly, the longer-term usefulness and effects of our approach has not yet been evaluated.

7   Conclusions

In this study, we evaluated user performance and experience in completing data exploration tasks using an interactive, visual interface (AM-L) versus a tabular data representation. While SAP was used for the latter, we believe our findings are generalizable for any table-based interface used for examining associative relationships between data entities. The participants overwhelmingly preferred AM-L, finding it less complex and more enjoyable to use. The lower mental effort ratings for AM-L may translate into a greater willingness to adopt and use the interface, which was also indicated in answers to a question concerning future use.

User performance was promising, with users faring as well or better with AM-L versus tabular data in terms of time. For correctness, user performance was worse with AM-L on the first two simple questions and the first complex question but then improved to being at or above the level of correctness achieved with SAP. It is expected that minimal training with the interface would help alleviate what appears to be a learning curve issue; this will be investigated in future studies. Our results also corroborate earlier reports that users find visualizations to be helpful in reducing the complexity of problem solving within the enterprise system context.

Outcomes from this study have confirmed the value in continuing with this work. Suggestions and observations from study participants will lead to further improvements in the AM-L interface, to be followed by experimentation for determining the optimal composition of features and interactions.

References

1. C. Lambeck, C. Fohrholz, C. Leyh, and R. Müller, (Re-) Evaluating User Interface Aspects in ERP Systems - An Empirical User Study, in Proceedings of the 47th Hawaiian International Conference on System Sciences, 2014.
2. T. Babaian, W. Lucas, and M. Li, Modernizing Exploration and Navigation in Enterprise Systems with Interactive Visualizations, in Human Interface and the Management of Information. Information and Knowledge Design: 17th International Conference, HCI International 2015, Los Angeles, CA, USA, August 2-7, 2015, Proceedings, Part I, S. Yamamoto, Ed. Cham: Springer International Publishing, 2015, pp. 23–33.
3. T. Babaian, W. Lucas, A. Chircu, and N. Power, Interactive Visualizations for Workspace Tasks, in LNBIP, Springer, forthcoming 2017.
4. H. Topi, T. Babaian, and W. Lucas, Informal Notes on Technology Use as a Mechanism for Knowledge Representation and Transfer, in Proceedings of The Sixth European Conference on Organizational Knowledge, Learning, and Capabilities (OKLC-2005), 2005.
5. M. Iansiti, ERP End-User Productivity: A Field Study of SAP and Microsoft, Key-stone Strategy, LLC, 2007.
6. R. Chang et al., Scalable and interactive visual analysis of financial wire transactions for fraud detection, Inf. Vis., vol. 7, no. 1, pp. 63–76, 2008.

7. Z. Liu, J. Stasko, and T. Sullivan, SellTrend: Inter-Attribute Visual Analysis of Temporal Transaction Data, IEEE Trans. Vis. Comput. Graph., vol. 15, no. 6, pp. 1025–1032, Nov. 2009.

8. M. Sedlmair, P. Isenberg, D. Baur, and A. Butz, Information visualization evaluation in large companies: Challenges, experiences and recommendations, Inf. Vis., vol. 10, no. 3, pp. 248–266, Jul. 2011.

9. W. N. Dilla and R. L. Raschke, Data visualization for fraud detection: Practice implications and a call for future research, Int. J. Account. Inf. Syst., vol. 16, pp. 1–22, Mar. 2015.

10. B. Saket, A. Endert, and J. Stasko, Beyond Usability and Performance: A Review of User Experience-focused Evaluations in Visualization, 2016, pp. 133–142.

11. J. D. Xu, I. Benbasat, and R. T. Cenfetelli, The Nature and Consequences of Trade-Off Transparency in the Context Of Recommendation Agents., MIS Q., vol. 38, no. 2, 2014.

12. Q. Wang, S. Yang, M. Liu, Z. Cao, and Q. Ma, An eye-tracking study of website complexity from cognitive load perspective, Decis. Support Syst., vol. 62, pp. 1–10, Jun. 2014.