Naval Research Laboratory
Washington, DC 20375-5320

NRL/7340/MR—2023/1

A Study of Visual Preferences for Map Label Placement

JAELLE SCHEUERMAN

JASON HARMAN

DINA M. ACKLIN

CHRIS J. MICHAEL

Cognitive Geospatial Systems
Ocean Sciences Division

March 16, 2023

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
16-03-2023

2. REPORT TYPE
NRL Memorandum Report

3. DATES COVERED *(From - To)*
13-09-2020 – 12-09-02022

4. TITLE AND SUBTITLE
A Study of Visual Preferences for Map Label Placement

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Jaelle Scheuerman, Jason Harman, Dina M. Acklin, and Chris J. Michael

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER
N2Y9

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Research Laboratory
4555 Overlook Avenue, SW
Washington, DC 20375-5320

8. PERFORMING ORGANIZATION REPORT NUMBER
NRL/7340/MR--2023/1

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Naval Research Laboratory
4555 Overlook Avenue, SW
Washington, DC 20375-5320

10. SPONSOR / MONITOR'S ACRONYM(S)
NRL

11. SPONSOR / MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A: Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
Karles Fellowship

14. ABSTRACT

Digital maps are important for many decision-making tasks that require situational awareness, navigation, or location-specific data. Often, digital mapping tools must generate a map that displays labels near associated features in a visually appealing manner, without occluding important information. Automated label placement systems generally accomplish this nontrivial task through a combination of heuristic algorithms and cartography rules, but the resulting maps often do not reflect the preferences and needs of the map user. To achieve higher quality map views, research is needed to identify cognitive and computational approaches for generating high-quality maps that meet user needs and expectations. In this paper, we present a study designed to better understand the visual preferences of map users and support the development of models for digital map displays. In particular, we found that participants demonstrated preferences for how labels are placed near their point of interest, and that when making trade-offs between alignment and distance of a label, they were more consistent in choosing positions that prioritized alignment over distance. We also investigated if context effects, such as similarity, attraction, or compromise, could be observed when participants chose between equally valued alternatives.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON / 19b. TELEPHONE NUMBER *(include area code)*
U	U	U	U	19	Jaelle Scheuerman / (228) 688-4993

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

This page intentionally left blank.

CONTENTS

FIGURES

TABLES

This page intentionally left blank.

A STUDY OF VISUAL PREFERENCES FOR MAP LABEL PLACEMENT

1. INTRODUCTION

Digital maps are commonly used to plan routes and navigate, such as by car on a street map or by plane on aeronautical charts. Other uses include visualizing and understanding spatial features or making decisions about how to deploy resources in a geographic area. In such scenarios, maps must display large amounts of relevant spatial information to the user in an intuitive display that helps them find the needed information quickly and easily.

For centuries, cartographers have conducted the time-consuming task of designing detailed maps and, in the process, have created many best practices and standards to represent geographic information accurately, while supporting spatial processing and reducing clutter. However, modern digital maps offer functionality that is not available in print maps, such as panning, zooming, and adding layers. This complex environment makes it impossible to manually create map views for every scenario. Mapmakers have turned to algorithmic and machine learning approaches [1, 2] to generate digital map displays.

Label placement algorithms seek to place labels near their respective features in such a way as to minimize collisions between map features and to reduce clutter. Several approaches have tried to incorporate cartography rules into label placement systems, but it is difficult to codify every possible situation. As the complexity of a map increases, the resulting layouts often do not meet the standards set forth by cartographers [3]. Although much of the work has focused on improving the efficiency of label placement algorithms, little has been done to incorporate information about human cognition and preferences into them. Such information could lead to higher quality layouts optimized for a particular user's needs or preferences.

The psychological literature has explored how visual attention, preferences, and other aspects of cognition can affect decision-making and spatial processing. Many examples have shown how psychological theories and models can be used to improve computational algorithms that account for human behavior [4, 5]. However, the incorporation of this information into map generation algorithms has not been well explored.

In this report, we present a study that explores the label placement preferences of map users and shows that people exhibit consistent visual preferences about where labels should be placed in relation to points of interest on a map. We also explore whether context effects, such as similarity, attraction, or compromise, emerge when making decisions between label positions of equal or similar preference value. The resulting dataset provides a basis for exploring visual preferences for label placement in the context of digital maps.

1.1 Related Work

Map label placement is one of the most challenging processes in map creation. To be useful, the labels must be placed near their associated features in a readable way, without colliding with or obscuring other labels [6]. This is a computationally complex task, even in its simplest form [7]. Digital maps make the problem more complex, requiring optimal placements to be recalculated as the map is scaled or layers are

Manuscript approved January 18, 2023.

Fig. 1—An example of the stimuli seen by participants in Part I of the survey. Alignment for both maps is fixed to top left. Distance is manipulated between the two maps, with labels on the left map positioned close to their point of interest, while those on the right are positioned farther away.

added and removed. Many heuristic algorithms have been suggested to approximate label placements more quickly [8]. However, the layouts produced by automatic algorithms often do not meet the quality standards established by cartographers [3].

Existing research has explored how the visual properties of maps influence the effectiveness of communicating spatial information. Reducing clutter is a popular area of study for map displays, with many researchers considering the effects of saliency and color density on clutter perception [9–11]. Further work has shown the effect of visual attention and salience on search strategies in novice and experienced users [12]. Visual properties, such as spatial distribution, number of objects on a screen, and object connectivity, have been used to analytically measure the readability of a map [13].

Decision-making research has explored how utility and preferences affect people's choices. For example, prospect theory has examined the effects of utility and preference on attention and decision-making [14]. Preference researchers have also explored many computational methods to represent and learn preferences [15, 16]. Some of this work has examined the use of cognitive models, such as decision field theory [17], to create preference models that account for the context effects seen in the psychological literature [18–20].

2. STUDY OF MAP PREFERENCES

We build on this research by presenting a study of visual preferences for map displays. It examines what type of label positions people prefer, and also investigates their decision-making when choosing between map displays that are similarly or equally preferred. This study serves as a first step toward understanding preferences in map displays and provides a dataset that supports the development of new models that can account for visual and spatial aspects of a map. The resulting dataset shows that users express consistent preferences considering the alignment and distance of a label from its point of interest. We also explore how people make decisions between equitable options. Finally, we will discuss the future potential of using and extending the dataset to better understand how visual attention and spatial reasoning affect map preferences.

2.1 Methodology

478 college students were asked a series of questions about their map preferences. All participants self-reported that they had vision that was normal or corrected to normal, and 3 participants stated that they were colorblind. Participants were asked how often they use maps in their daily life (9.4% rarely/never, 35.3% a few times per month, 40.0% a few times per week, and 15.3% daily). They were also asked which digital map applications they have experience with (Google Maps – 83.5%, Apple Maps – 76.8%, Waze – 63.6%, Map My Run – 7.8%, MapQuest – 4.8%, Strava – 2.9%, All Trails – 2.5%, Other – 3.1%).

After completing the introductory questions, the participants completed 3 survey blocks totaling 52 questions. Part III of the study included eight questions with one correct answer. If the instructions were understood and the participant was paying attention, they should consistently get these questions correct. In the results and analysis below, we removed inattentive participants who did not get at least 6 of the 8 questions correct, leaving 292 participants.

2.2 Part I: Preference Consistency

Part I of the survey was designed to answer whether users show consistent preferences regarding two attributes of the label: alignment and distance. Participants were presented with a series of binary choice questions in which they were asked to choose the map that they thought was more useful or easy to use in real world scenarios. Each binary choice involved two maps where the labels differed in alignment or distance, but not both (see Figure 1 for an example). The participants responded to choices between every possible combination of two alignment conditions, chosen from {top left, direct left, bottom left} and every combination of two distance conditions chosen from {close, medium, far}. When the alignment conditions were compared, the distance of the labels remained fixed in the medium position. When the distance conditions were compared, the alignment of the labels remained fixed at the top left position. Two backgrounds were used to test whether people's preferences changed when the streets on the map were in a grid pattern compared to a more complicated radial pattern. This resulted in 12 total questions in Part I. In addition to choosing their preferred map, participants were asked to describe the strength of their preference by dragging a marker along a slider with a range of [0,100].

2.3 Part II: Preference Tradeoffs

Part II considers how user preferences across multiple dimensions may be combined when choosing an overall preferred label position. However, instead of choosing between positions that differ in a single attribute, participants were asked to rank three positions that differ in both distance and alignment. We considered two sets of three maps on four different backgrounds (2 with streets in a grid pattern and 2 in a radial pattern). This resulted in a total of 12 questions.

- **Label Set 1**: (top left, medium), (direct left, far), (bottom left, close)
- **Label Set 2**: (top left, far), (direct left, close), (bottom left, medium)

2.4 Part III: Context Effects

Finally, in Part III, we wanted to investigate whether context effects are observed when people choose between equally preferred options. Context effects occur when adding a third option to a choice set causes a violation of expected utility. Three context effects are prominent in the literature, including similarity, attraction, and compromise. The similarity effect is produced by adding a competing alternative that is similar to the choice set [18], while the attraction effect can be observed when adding a dominated alternative (i.e., one that is less preferred to all other alternatives in the choice set) [20]. Finally, the compromise effect is produced by adding an alternative that is less extreme in some dimension, compared to the other choices [19]. These effects have been observed across many paradigms, but to our knowledge, they have not been explored in a rich visual environment such as map displays. To test for these effects, we chose questions that should produce contextual effects if they are present.

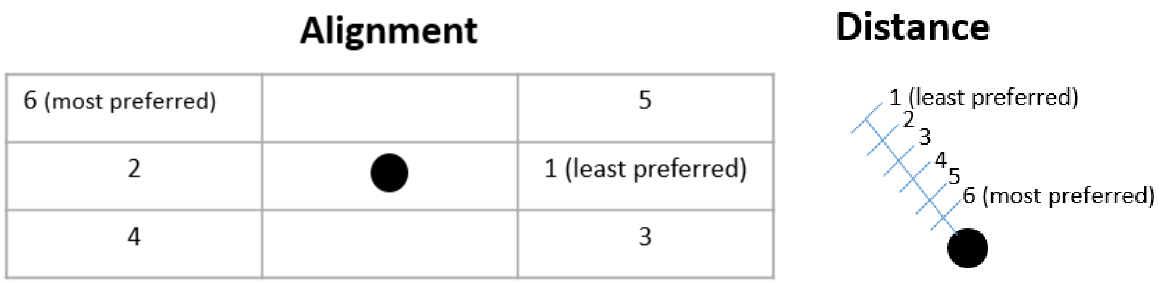

Fig. 2—Boss's preference values for the different alignment and distance options

The participants were presented with a set of preference values for six distances and six alignments of a point label (see Figure 2). They were told that their boss asked them to choose a label's position given these preferences. Then, a set of three maps was shown in which a single label differs in alignment and distance. Participants were asked to choose the map that represented the best trade-off between alignment and distance, using the boss's preferences. In the event of a tie, participants were asked to choose the position of the label they thought was best.

In most of the questions, the three label positions were chosen to be equal to or similar in value for the boss, but the choices were selected in such a way as to identify context effects if they were present. For example, a similarity question would include three label placements where the preferences for each attribute add up to the same value (i.e., $(2,5)$,$(1,6)$, and $(5,2)$). If the similarity effect is present, we expect participants to choose more often the position that is most different (i.e., $(5,2)$, which has a high value in the first dimension and a low value in the second). Table 1 shows the label placement sets that test for similarity effects.

Questions that test for the attraction effect include two label positions where the preference values add up to the same total (i.e. $(1,6)$ and $(4,3)$, while the third position is not quite as good, but similar to one of the other two (i.e. $(1,5)$). In this example, $(1,5)$ is not as good as $(1,6)$, but since it is similar, the attraction effect would lead participants to choose $(1,6)$ more often than the other choices. Table 1 shows the label placement sets that test for attraction effects.

Finally, the compromise questions provide three label positions that have equal total preference values. However, two of the options are extreme in each dimension, while the third is more moderate (i.e. $(1,6)$, (5, 2) have extreme values in both dimensions, while (3, 4) is more moderate). Table 1 shows the label placement set that test for compromise effects.

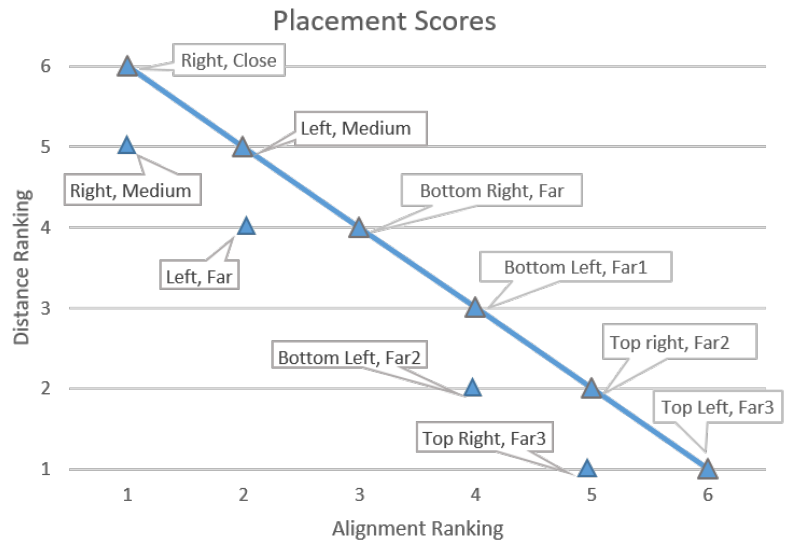

Fig. 3—Preference values for label positions in Part III questions testing for context effect

Set Name	**choice 1**	**choice 2**	**choice 3**
similarity1:	**(direct left, medium): 7**	(top right, far2): 7	(top left, far3): 7
similarity2:	(direct left, medium): 7	**(top right, far2): 7**	(right, close): 7
attraction1:	**(right, close): 7**	(right, medium): 6	(bottom left, far1): 7
attraction2:	(right, close): 7	**(bottom left, far1): 7**	(bottom left, far2): 6
compromise:	(right, close): 7	**(bottom right, far): 7**	(top right, far2): 7
other1:	**(top left, medium): 11**	(direct left, far): 6	(bottom left, close): 9
other2:	**(top left, far): 10**	(direct left, close): 8	(bottom left, medium): 9

Table 1—The label position sets compared in Part III. Each choice lists the alignment, distance, and boss's utility value for that choice. The names correspond to the names along the x-axis in Figures 7 and 8. The expected choice is in bold.

The survey included the label positions described in Table 1 to test for attraction, similarity, and compromise effects. Additional questions included the same alignment and distance trade-offs from Label Sets 1 and 2 in Part II. The label positions in these additional questions did not have equal preference values and therefore had only one correct answer. The seven types of questions were replicated across four different map backgrounds, resulting in 28 total questions.

To ensure that the participants understood the instructions, each person completed a tutorial made up of five questions. To proceed to the next question of the tutorial, the participant had to choose the correct response. If they chose incorrectly, the tutorial would provide feedback on why the choice was wrong. After successfully completing the tutorial, the participant went on to complete the study.

3. RESULTS AND DISCUSSION

In the following sections, we report on the results of each part of the survey, including an analysis of participant responses.

3.1 Part I

When analyzing the results from Part I, we compared the distribution of responses from both of the map backgrounds (grid and radial) to see if there were statistical differences in people's preferences between the two maps. We did not find statistical differences in the responses between the groups. Therefore, we combined these groups into one for the remainder of the analysis. This resulted in 584 responses for each comparison (2 per participant).

When analyzing alignment preferences, we found that the direct left position (66.4%) was preferred to the bottom left (33.6%), with a mean preference strength of 50.1 (sd=29.7). Direct left (54.6%) was also slightly preferred to the top left position (45.4%), with a mean preference strength of 50.9 (sd=28.8). Finally, the top left position (62.3%) was preferred to the bottom left position (37.7%). The mean preference strength was 49.7 (sd=29.2).

When analyzing distance preferences, we found that medium positions (79.3%) were preferred to close positions (20.7%), with a mean preference strength of 51.1 and a standard deviation of 28.8. Far distances (73.5%) was also preferred to close (26.5%), with a mean preference strength of 49.2 and a standard deviation of 29.2. Medium distance positions (67.8%) were preferred to the far positions (32.2%) with a mean preference strength of 44.2 (sd=30.4).

If participants were indifferent between the label positions, we would expect to see approximately 50.0% of responses for each choice. However, the above analysis shows that participants were not indifferent between the position of the labels, and generally preferred labels that were aligned to the left and positioned a moderate amount from their associated point of interest. The next preferred alignment was top left, while far was the next preferred distance position. The bottom left and close positions were the least preferred.

Next, we wanted to understand how people combine different attributes when choosing a position where both alignment and distance are manipulated. Will they maintain the same general preference for distance and alignment when these are changed? If they need to choose labels that are preferred in one dimension but disliked in another, does one factor get weighted higher than the other?

These questions lead to Part II of the study, which considers how participants rank label positions when a label differs in both alignment and distance.

3.2 Part II

The two label sets described in the methodology above were shown on four different map backgrounds. We compared the distribution of responses for each background map group and found no statistical differences between them. Therefore, we combined the four groups into one, resulting in 1168 responses for each trade-off set (4 per participant).

When considering Label Set 1 (see Figure 4), we note that participants ranked (top left, medium) more often in first place (73.4%), while (bottom left, close) is more often ranked in last place (61.7%). In Label

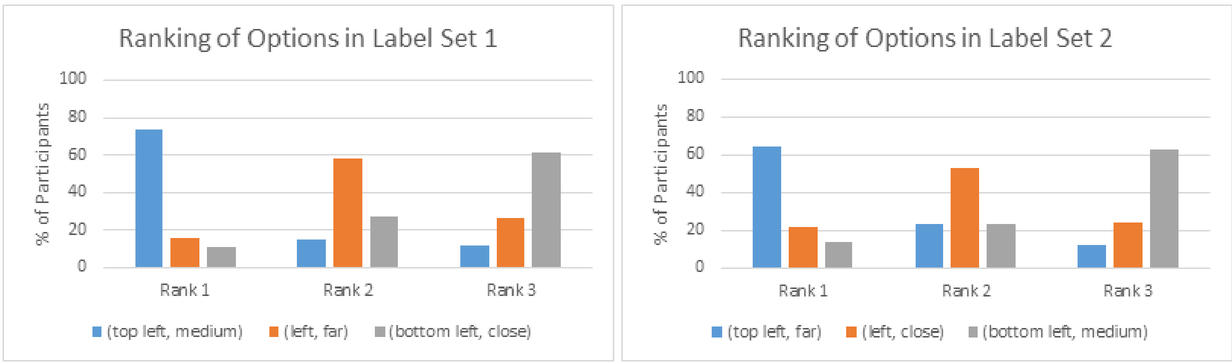

Fig. 4—Rankings of label positions in Label Set 1 and Label Set 2

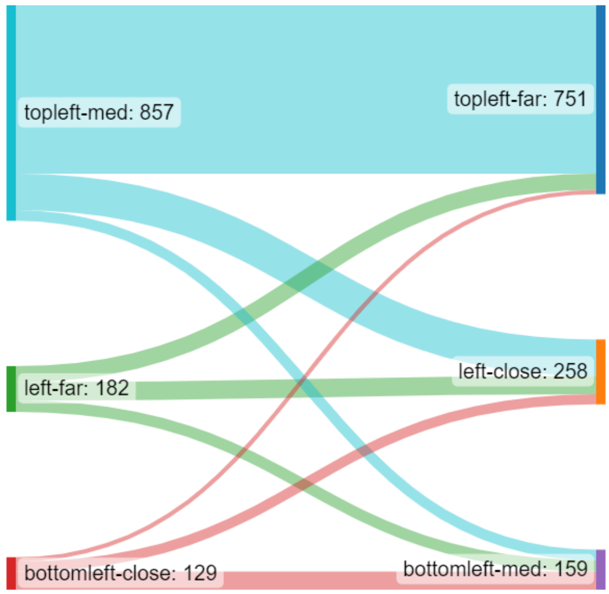

Fig. 5—Proportions of participants' first ranked items in Label Set 1 (left) and Label Set 2 (right).

Set 2 (see Figure 4), we see that there is a clear preference for ranking (top left, far) first (64.2%), (direct left, close) second, and (bottom left, medium) third (63.1%).

Most of the participants appeared to prioritize alignment over distance information in both scenario. A visualization of the options that the participants ranked first in Label Set 1 and Label Set 2 can be seen in Figure 5. As a group, most of the participants (Label Set 1: 73.4%, Label Set 2: 64.2%) ranked the labels placed at the (top left, medium) position. 78.2% of those who chose the top left label in Label Set 1, chose the (top left, far) label position in Label Set 2.

If participants had prioritized distance information, then (top left, far) would not have been so strongly preferred in Label Set 2, compared to the small number of responses (39.6%) that ranked (direct left, far) in Label Set 1. In fact, of the participants who chose (direct left, far) in Label Set 1, the largest proportion (39.6%) continued to prefer left aligned labels (direct left, close) in Label Set 2.

Although most of the participants appeared to focus on the alignment of labels, there was a minority who appeared to take distance information into account. Of those who chose the position (top left, medium) in Label Set 1, 16.9% shifted their decision to (direct left, close) and 4.9% shifted to (bottom left, medium) in Label Set 2. This is in line with the distance preferences we observed in Part I, with people preferring closer labels to those farther away.

When comparing these results with those found in Part I, we found some similarities. In Part I, we found that people generally preferred bottom left positions the least, and this was also seen in Part II. In fact, when considering (bottom left, medium) in Label Set 2, we found that people's preference for medium positions in Part I seems to be overpowered by their preference against the bottom left position. With that exception, we found that people preferred the medium and far placed positions to close positions, as we also found in Part I.

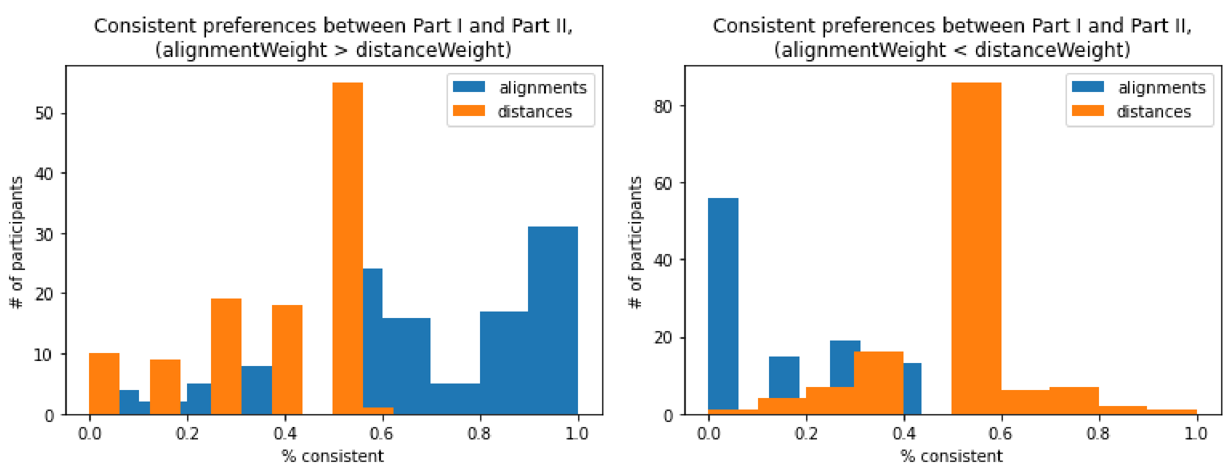

Fig. 6—Each graph represents the number of participant responses in buckets denoting how consistent they were. If % consistent is 1.0, the participant chose the Part II option in every question that aligned with their preferences along the associated dimension in Part I.

Modeling Map Preference Trade-offs To better understand how participants choose which label positions to rank first, we employed a model to estimate the utility for each of the label positions in Part II, fitted to the preference data collected in Part I.

In Part I, participants chose their preferred position between three distance positions and three alignment positions in a series of pairwise comparisons. Using this information, we can generate an order over the positions in A = {top left, left, bottom left} and positions in D = {close, medium, far}, for each map background, such that $a_1 \geq a_2 \geq a_3$ and $d_1 \geq d_2 \geq d_3$.

We assign a utility to each position, assuming that the relationship between them is linear, and that higher ranked items should have a higher utility than lower ranked items. More formally, n is the number of options compared ($n = 3$ in this case), and $r(x)$ returns a value in $\{1, 2, .., n\}$ that corresponds to the rank index of x. Note that it was possible for multiple positions to share the same rank $r(x)$ if the participant was indifferent between them. Therefore, the utility of the alignment position is $\forall a \in A, u(a) = n - r(a)$, and the utility value of the distance position can then be described $\forall d \in D, u(d) = n - r(d)$.

Given these utilities, we can model the rank 1 responses in Part I by aggregating the alignment and distance utilities for the label position, as a weighted sum: $u(P) = w_A u(a_P) + w_D u(d_P)$. Here, w_A and w_D represent the weights applied to the alignment and distance features, respectively.

The weights represent how consistently an individual chose options using a preferred placement found in Part I. For example, if w_D was closer to 1, then the participant's decisions in Part II were more consistent with their distance preferences collected in Part I. If a weight was closer to -1, then their preferred choices in Part II were not heavily influenced by their preference in Part I.

After the utilities for each option have been generated, the model returns the choice with the highest utility as the preferred option. If multiple choices are tied for the highest utility, then the model is indifferent between them and returns one randomly chosen from a uniform distribution.

As seen in Figure 5, participants exhibited a variety of strategies when ranking positions, so we model the preference weights for each individual, rather than for the entire group.

Model Results To model the rankings, we use a grid search to fit the values for w_A and w_D, chosen between -1 and 1 in increments of 0.1 $[-1, -0.9, \ldots, 0, 0.1 \ldots, 0.8, 0.9]$. We fit the values by selecting the pair (w_A, w_D) that most often predicted an individual's choice of ground truth Rank 1 choice over eight total responses (four responses each combined with the two sets of preferences gathered in Part I). Predictions were generated by calculating the utility of the position using the process described above. Since the model broke ties randomly, a variety of outcomes were possible. To ensure that the model works well across a number of different random seeds, we ran the model 20 times. When comparing the runs, there was no statistical difference between them, and the model accurately produced 72.7% (sd = 0.01) of the Rank 1 responses on average. We used one of these runs (accuracy = 72.9%) to perform the weight analysis below.

Weights Analysis We can analyze the weights of the resulting models to understand how consistent participants were in their preference choices between Part I and Part II of the survey. For example, if $w_A > w_D$ (N = 896), then the participants were generally more consistent in choosing the option that reflected their alignment preferences from Part I, compared to distance preferences. Likewise, if $w_A < w_D$ (N = 1040), then participants were more consistent in choosing an option that reflected their distance preferences in Part I. If $w_A = w_D$ (N = 400), then participants were often not very consistent in using their Part I preferences to choose options in Part II. In Figure 6, we can see a summary of these results, showing that people were generally more consistent in choosing options that used their alignment preferences than those that used their distance preferences.

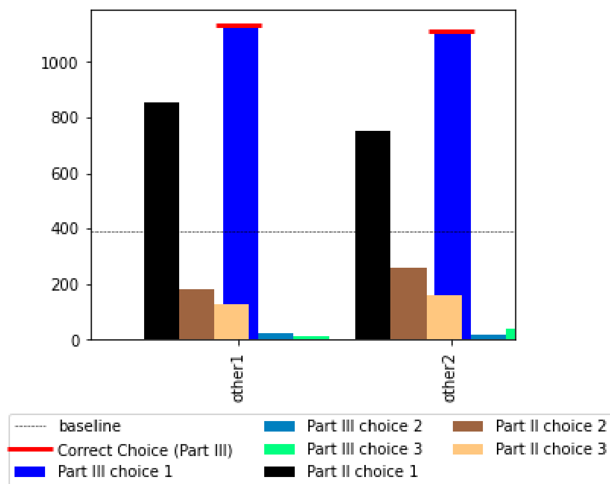

Fig. 7—Comparison of responses from Part II and Part III for Label Sets 1 and 2. Part II choices are those that were ranked first in Part 2. Part III choices were the position selected as the best choice, given the boss's preferences.

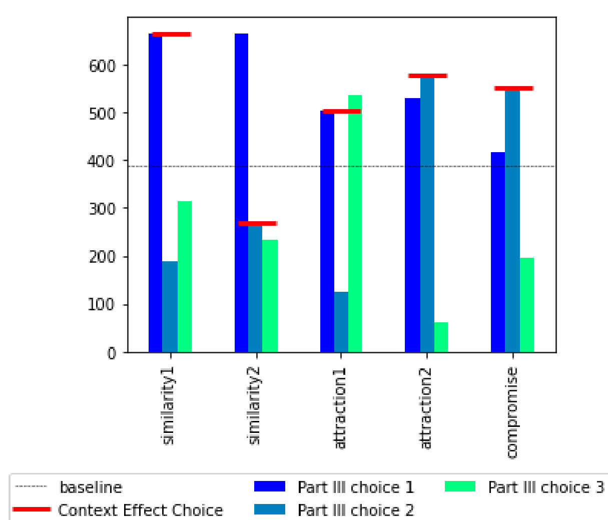

Fig. 8—Results of context effect questions. Red lines mark the choice we would expect to be chosen the most often, if the associated context effect is seen in the data

3.3 Part III

Recall that Part III asked participants to choose the label position that best represented their boss's preferences. Eight of the questions replicated Label Set 1 and Label Set 2 from Part II, on four backgrounds. However, in this part of the survey, participants needed to respond using the boss's preferences, rather than their own. In all of these questions, there was a single correct answer. We compared the distributions of the responses from Part II and Part III and found that they were significantly different ($P < 0.005$). Figure 7 shows a comparison between each group's responses for each question. In Part II, we see that, while there was a clear majority in favor of one particular label position in each question, minority groups still preferred the less popular option in each case. However, in Part III, nearly all participants chose the option that correctly aligned with the boss's preferences. These results seem to indicate that the participants understood the instructions to choose maps using the boss's preferences in Part III.

The remaining questions in Part III of the survey were included to see if similarity, attraction, and compromise effects occur in the rich visual environment of a map. Recall that given the boss's preferences, the three label positions compared in each question either had equal preference values (to test for the similarity and compromise effect), or near equal preference values (to test for the attraction effect). For the questions that were equal (similarity and compromise), we would expect that if the participants made the choice randomly among positions of equal value, each choice would receive 33.3% of the responses. In the case of questions testing the attraction effect, where one option is slightly less preferred by the boss, we would expect the less preferred option to get the least number of responses and the remaining responses to be evenly distributed among the other two.

We can see from Figure 8 that participants did not randomly choose label positions among those with equal preference values. Although the context effect option is the most popular in the first similarity set (57.0%), it is not very popular in the second similarity set (23.1%). Instead, it seems that people chose the (direct left, medium) aligned option by a large majority in both of these questions (similarity1: 57.0%, similarity2: 56.7%). In the attraction questions, the dominated option is chosen much less often than the other two, and the responses are nearly equally distributed between the remaining two options (attraction1 - choice1: 43.2%, choice3: 46.0%, attraction2 - choice 1: 45.4%, choice 2: 49.6%), with the (bottom left, far1) position being slightly preferred in both cases. Compromise questions may show a weak effect, with the contextual choice slightly preferred (47.3%) over the other two options.

Rather than exhibiting context effects in these questions, the participants seem to be choosing a preferred option given some other internal method. The participants were instructed to break ties by choosing the position they thought was best, so their choices may have been moderated by their own preferences, as the boss would have been indifferent between the label positions. Many of the existing studies in the literature use values that are objective (i.e., decisions being based on the strength of eyewitness testimony, or by choosing the area of a rectangle [21]), such that they may not evoke strong personal preferences or be open to interpretation.

4. CONCLUSION

4.1 Limitations and Future Work

Certain limitations exist in the current dataset that should be addressed in the future. Due to limits on the participants' time, we did not collect responses for all possible distance and alignment comparisons in Part I or ask participants to rank all the possible Part III positions in Part II. Since Part III questions made use of alignment and distance options that were not present in Parts I and II, it is not possible to analytically test if people use their personal preferences to break ties between the choices that the boss was indifferent about. To address this, we will collect additional data where people rank all the context effect questions using their own preferences, as they did in Part II. This information would allow us to compare the distributions of their choices before and after being instructed to use the boss's preferences.

The results and analysis presented here represent the first steps in understanding people's visual preferences in digital map displays. Future work will include developing models of preferred label positions using data collected in the survey, as well as incorporating additional cognitive and spatial information such as saliency, clutter perception scores, color density, spatial distribution, or alignment with other map features.

In the future, we plan to test whether using people's label position preferences in map generation impacts their experience or the map's effectiveness for finding information on new maps. It would also be interesting to compare the differences between the preferences of novice and expert map users, to understand the impact of expertise on visual preferences in map displays.

4.2 Conclusions

We have introduced a study of visual preferences for the placement of labels on map displays. This study considered whether people have consistent preferences about the alignment and distance of labels from their associated point of interest. We considered preferences when alignment and distance are manipulated independently and together. In general, we found that people preferred specific label positions and that they were more consistent in choosing positions that prioritized alignment over distance.

We also collected responses where participants chose positions that best represented some provided preferences, and found that they successfully applied these when one position clearly dominated the others. However, when possible positions had the same preference value, participants did not evenly distribute their responses among options, indicating that they used some internal mechanism to break the ties. We tested to see whether psychological context effects, such as compromise, attraction, and similarity, could be observed when participants chose between positions of equal value, but did not observe this outcome. Participants may have broken ties using their own personal preferences. Additional data must be collected to further investigate this.

The dataset resulting from this survey will pave the way for new models of visual preferences in map displays and creates opportunities to consider how visual attention, clutter, and spatial perception influence preferences in digital displays.

REFERENCES

1. H. Freeman, "Automated Cartographic Text Placement," *Pattern Recognition Letters* **26**(3), 287–297 (2005).

2. J. Stoter, D. Burghardt, C. Duchêne, B. Baella, N. Bakker, C. Blok, M. Pla, N. Regnauld, G. Touya, and S. Schmid, "Methodology for Evaluating Automated Map Generalization in Commercial Software," *Computers, Environment and Urban Systems* **33**(5), 311–324 (2009).

3. W. A. Mackaness, A. Ruas, and L. T. Sarjakoski, *Generalisation of Geographic Information: Cartographic Modelling and Applications* (Elsevier, 2011).

4. O. Plonsky, I. Erev, T. Hazan, and M. Tennenholtz, “Psychological Forest: Predicting Human Behavior,” May 2016.

5. D. D. Bourgin, J. C. Peterson, D. Reichman, S. J. Russell, and T. L. Griffiths, “Cognitive Model Priors for Predicting Human Decisions,” Proceedings of the International Conference on Machine Learning (PMLR), 2019, pp. 5133–5141.

6. P. Yoeli, “The Logic of Automated Map Lettering,” *The Cartographic Journal* **9**(2), 99–108 (1972).

7. J. Marks and S. M. Shieber, “The Computational Complexity of Cartographic Label Placement,” Harvard Computer Science Group, 1991.

8. F. Wagner, A. Wolff, V. Kapoor, and T. Strijk, “Three Rules Suffice for Good Label Placement,” *Algorithmica* **30**(2), 334–349 (2001).

9. M. C. Lohrenz, J. G. Trafton, M. R. Beck, and M. L. Gendron, “A Model of Clutter for Complex, Multivariate Geospatial Displays,” *Human Factors* **51**(1), 90–101 (2009).

10. R. Rosenholtz, Y. Li, and L. Nakano, “Measuring Visual Clutter,” *Journal of Vision* **7**(2), 17 (Aug. 2007), ISSN 1534-7362, doi:10.1167/7.2.17.

11. G. Touya, C. Hoarau, and S. Christophe, “Clutter and Map Legibility in Automated Cartography: A Research Agenda,” *Cartographica: The International Journal for Geographic Information and Geovisualization* **51**(4), 198–207 (2016), ISSN 1911-9925.

12. C. Davies, W. Tompkinson, N. Donnelly, L. Gordon, and K. Cave, “Visual Saliency as an Aid to Updating Digital Maps,” *Computers in Human Behavior* **22**(4), 672–684 (July 2006), ISSN 0747-5632, doi:10.1016/j.chb.2005.12.014.

13. L. Harrie, H. Stigmar, and M. Djordjevic, “Analytical Estimation of Map Readability,” *ISPRS International Journal of Geo-Information* **4**(2), 418–446 (June 2015), ISSN 2220-9964, doi:10.3390/ijgi4020418.

14. S. Kim, J. L. Harman, and M. R. Beck, “Diminishing sensitivity and absolute difference in value-driven attention,” *Journal of Vision* **22**(1), 12–12 (2022).

15. F. Rossi, K. B. Venable, and T. Walsh, “A short introduction to preferences: between artificial intelligence and social choice,” *Synthesis Lectures on Artificial Intelligence and Machine Learning* **5**(4), 1–102 (2011).

16. Y. Chevaleyre, F. Koriche, J. Lang, J. Mengin, and B. Zanuttini, “Learning ordinal preferences on multiattribute domains: The case of CP-nets,” in *Preference learning*, pp. 273–296 (Springer, 2010).

17. T. Rahgooy and K. B. Venable, “Learning Preferences in a Cognitive Decision Model,” Proceedings of the International Workshop on Human Brain and Artificial Intelligence (Springer), 2019, pp. 181–194.

18. A. Tversky, “Elimination by aspects: A theory of choice.,” *Psychological review* **79**(4), 281 (1972).

19. I. Simonson, “Choice based on reasons: The case of attraction and compromise effects,” *Journal of consumer research* **16**(2), 158–174 (1989).

20. J. Huber, J. W. Payne, and C. Puto, "Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis," *Journal of consumer research* **9**(1), 90–98 (1982).

21. J. S. Trueblood, S. D. Brown, and A. Heathcote, "The multiattribute linear ballistic accumulator model of context effects in multialternative choice.," *Psychological review* **121**(2), 179 (2014).