Behavior-Based Cognitive Trait Modeling Via Node-Link Diagram Interactions

A thesis

submitted by

Melissa Iori, B.S.

In partial fulfillment of the requirements
for the degree of

Master of Science

in

Computer Science

TUFTS UNIVERSITY

May 2018

Advisor: Remco Chang

For Pepper & Percy.

Behavior-Based Cognitive Trait Modeling Via Node-Link Diagram Interactions

Melissa Iori

Advisor: Remco Chang

Individual traits are not always considered when designing data analytic software. However, previous research has shown that individual differences do matter when presenting visual content. Our experimental results demonstrate the ability to determine a user's extraversion in real time by analyzing their interactions with a data visualization found on many platforms, the node-link diagram. The results show that users with higher extraversion moused-over more nodes than those with lower extraversion, showing a distinct difference in behavior between user trait groups. I then discuss the results of the experiment, and the impact this finding has for user trait modeling.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Most data analysis dashboards assume that all users will have similar needs for information presentation and availability. However, previous research has demonstrated an increase in the level of user satisfaction with the introduction of adaptive software features [GC17]. Evidence also exists indicating performance differences based on personality traits when performing visual searches of data [BOZ+14]. In order to maintain or even improve user satisfaction and performance for data visualization tasks such as visual search, I propose a user personality model generated only from the user's interactions with the visualization itself. In this paper, I focus on node-link diagram visualizations. With this proposed approach, once a system identifies a user's trait via the model, it can adapt in multiple ways throughout the software, including adapting help systems, providing more useful recommendations, or showing different visualizations of the same underlying data based on trait [SCC13].

There are several advantages to these user models being developed. Many data-analytic user interfaces already include node-link diagrams, especially those made for the domains of social network analysis, road analysis, ontology analysis, fraud analysis, and other domains [KEC06]. Developers will save time by instrumenting the already-existing visualizations to obtain user insights, rather than having to implement a more complex tool or survey instrument. Other methods for generating personality trait models exist, but some involve use of cameras, or biometric / physiological input, which can be invasive for the user [MML07]. By basing the model on only interactions, trait information can be gained in a non-invasive, organic way. Traditionally, trait surveys have required users to report their own traits by answering a series of questions. This can lead to data biased by the user's own choice of answers, which makes it difficult to compare users and determine traits. Additionally, as more trait adaptations are generated, more traits will need to be determined, and users will have to fill out more surveys, which will take significantly more time without a user model. Traits can also change over time, and collecting survey data at a fixed point in time is not a sufficient method of capturing these changes. A behavior-based user model has the capability of tracking changes in trait over time if necessary.

Previous research has shown that individual differences alter strategies when interacting with data visualizations [OYC15]. Visual search task patterns can vary in choice of action based on personality traits of the user (e.g. locus of control), which suggests that personality trait insights can be gained by analyzing interactions from visual search tasks [BOZ+14]. Our shows that a model based on interaction data alone can be more accurate than the baseline of 50% accuracy when predicting binary user traits - traits that can divide users into two groups, such as a "high" group and a "low" group. The system could randomly guess the user's trait, a method that would provide only 50% accuracy. Using a model based on our insights about behavior and user interactions would yield a more informed and therefore more accurate model.

A personality trait model could be used to improve many current systems, and to create trait-aware systems that don't currently exist. At least two types of system would benefit from trait user modeling - recommender systems that tailor content to things that a particular trait group is interested in, since traits and their trait preferences for content such as advertising have been established [CHK+15], and help systems that identify when the user's personality type happens to correlate with lower performance for the task that they are attempting to complete (and provides additional assistance in this case) since trait differences due to performance have been established [ZOC+12] [OYC15].

An appropriate real-world setting in which to utilize an instrumented node-link diagram with a trait model is in a hypothetical ontological visualization interface. Ontology data is often a network of related terms forming the "jargon" of a particular domain and how that jargon is connected [LSR09]. This data is similar in nature to the WordNet diagrams that I utilized for my own experiments, making this particular use case a good fit. The ontologies could be important terms for specific domains such as medicine, programming languages, or psychology, among others. A hypothetical adaptive ontological visualization interface could utilize trait models to detect the user's trait while the user is freely exploring, as they did in my experimental tasks. Performance differences for visual search are known in extraverts and introverts - extraverts have been observed to act more quickly, whereas introverts take time to complete tasks [ZOC+12]. However, quick searches do not always provide accurate results, so the system could be designed to let extraverts know to slow down, while introverts can be provided with features that remind them of what nodes they've already visited in the ontology graph, such that they can easily seek out the new information that they wish to see, complementing their trait. Such a system hasn't been implemented or evaluated at the time of this work, but would be a natural use case for the proposed models based on visualization interactions.

In this paper, I describe preliminary findings that demonstrate extraversion trait modeling ability through a simple visualization experiment. Our experimental results show a significant difference between user's high and low in extraversion when interacting with a node-link diagram.

Need for cognition analysis produces a notable result in which the individual user's need for cognition had a negligible effect on how many unique nodes they visited. Previous research has shown that, while higher need for cognition users prefer cognitive tasks and demonstrate higher curiosity, there are circumstances in which behavior may not be as expected, such as Wu's study of Need For Cognition (NFC) and search results, where higher NFC users paginated less frequently and paid less attention to lower-ranked information [WKS14].

Chapter 2

Related Work

2.1 Related Work

As previous research has shown, individual differences have an impact on user interaction patterns with data visualizations [ZOC+12]. Previous work has established that data visualization search strategies vary based on the traits of the user [ZOC+13] [Che00] [OYC15]. Given these differences, a model can be constructed based on decision trees to determine the user's speed in completing a task [BOZ+14].

Measuring user behavior quantitatively would be more difficult if the interaction behavior data lacked a structure or common patterns. However, the theory of information foraging has been researched as it regards search strategies, and Pirolli et. al. have theorized that information foraging plays a critical role in the decisions that users make when searching for information, such as in a visualization [Pir07] [PCW01]. Information foraging is a cognitive theory that suggests a structure to the methods people use when searching for informative content. Applying this idea to a visual search provides an interesting possibility for user modeling - a model based on structured search behavior for different users, which could be quantified through clickstream data.

Chen and Czerwinski investigated individual differences and search strategies in a "semantic space" (a node-link visualization of reearch papers), and discovered that multiple strategies exist, which demonstrates that different strategies can be quantitatively measured for different users [CC97]. The related research described in this section has demonstrated the capability for individual search strategies in visualizations to be discerned through interaction analysis. Traits such as locus of control have also been shown to affect search strategy [BOZ+14]. However, little evidence currently exists for the use of interaction data from exploratory node-link diagram searches to predict traits such as extraversion. More data needs to be gathered from participants who are placed in an ecologically valid, open-ended (with multiple "correct" answers) experiment, in order to determine how trait behaviors actually manifest themselves in software interactions. In this work, I provide an ecologically valid study that takes place in the participant's Web browser and tests the hypothesis that traits can be discerned from interaction data alone.

2.2 Personality Traits

In this work, I attempt to collect data on user behavior with a node-link diagrams and compare the gathered clickstream data with validated trait survey instruments. Extraversion is one personality trait of the Big Five personality traits, and it measures one's external sociality [big]. Research exists that suggests that extroverts and introverts approach problems differently [GC17]. Need for Cognition (NFC) is a trait linked to curiosity and preference towards "thinking" tasks [CP82].

Chapter 3

Approach

3.1 Measurement

The most accepted scale to measure user extraversion is The Big Five Factor Markers survey [big]. A subject provides answers to questions on the extraversion scale. We can then produce a numerical score for that subject on the extraversion scale. I calculated the median score out of all subjects, then assigned a binary label to subjects based on their score - "high" scores were above median, and "low" scores were below median.

3.2 Search

To create an ecologically valid search setting, I used JavaScript to create node-link diagram visualizations in the Web browser. Thus, the experience of subjects resembled a real-life interaction scenario. This ecological validity is crucial to understanding what real behaviors would be observed outside of an experimental scenario.

The visual search task involved a node-link diagram visualization that contains nodes, each labeled with a word or short phrase found in the English language. This data originates from the WordNet dataset.

Subjects are provided with a dictionary definition of a particular word, but not the word it pertains to. They are tasked with locating the word in the diagram that, in their opinion, best fits the definition. Multiple answers are possible.

3.3 Hypotheses

Previous research suggests user extraversion can affect the patterns with which users search for data [GC17] [ZOC$^+$13]. As well, users higher in need for cognition are "more curious", and due to this, I hypothesized that they would utilize node mouseovers significantly more than users in the "low" group.

H1: Higher extraversion in the user can be determined from mouseover count.

H2: Higher need for cognition in the user will yield more mouseovers on nodes.

Chapter 4

Experiment Design

An experiment was performed to investigate traits and behaviors. I used an instrumented website to gather the data, recruiting subjects via Amazon's Mechanical Turk. Each subject was paid $15.00 for participating. The website uses D3.js [BOH11] to generate node-link diagrams of related words, based on WordNet [Mil95]. It also logs all interactions with the visualization, including mouse clicks, mouse overs, and other DOM events by using UserAle.js. [apa]

Using Amazon's Mechanical Turk service, 50 people participated in the study and were compensated with $15.00 each. Participants filled out a demographic survey at the beginning of the study. Demographics out of 50 participants were 27 male, 23 female, with the average age range being 30-39. 80% of participants rated their computer experience as "somewhat experienced" to "very experienced", and 88% of participants rated their data visualization experience as "a little experienced" to "somewhat experienced".

After completing the demographic survey, each participant filled out cognitive surveys for various traits, allowing for data collection on the traits of the participants according to the current surveys [big] [CP82].

The experiment consisted of two node-link diagram tasks. In our node-link diagrams, each node represents a word or phrase from the WordNet dataset. The word or phrase associated with each node can only be viewed when the mouse cursor is hovering over that node. Otherwise, it is not visible to the user. Each

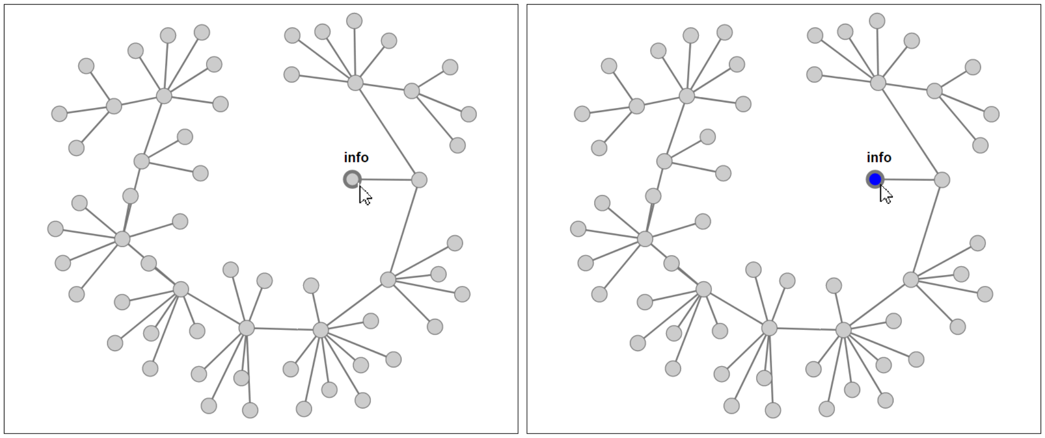

Figure 4.1: One of the task diagrams. On the left, the user is mousing over a node and can view the associated word, "info". On the right, the user has clicked on the "info" node, which changes its color to dark blue.

task was preceded by a tutorial explaining the goal of the task and how to interact with the node-link diagram, including a sample task. The tasks allow the user to freely explore the node-link diagrams while not imposing a lot of restrictions on their mouse movement. By having words only appear when the mouse cursor is hovering over the matching node, I imposed a restriction on exploration that allowed me to quantitatively measure the user's actions when seeking information in the diagram. The size of the graphs depicted (in terms of number of nodes) was determined through pilot studies and conformance to a minimum screen size for the experiment application. 60% of participants self-reported their mental load on this task as 65 or greater on the NASA TLX scale [HS88] (out of 100 maximum mental load score), demonstrating that the task provided enough mental load such that natural problem-solving behaviors could be observed.

The first task, "A", displayed a node-link diagram to the participant. Each node-link diagram had approximately 60 nodes. The nodes are by default all colored in light gray. Clicking on a node colors it in dark blue. The text of a definition of a word in the English language is shown on screen next to the diagram. The participants were tasked with locating the node of the word that they thought best fits the given definition. If the participant chooses to de-select a previously clicked node, they click on it again and its color reverts to gray. They can also freely select another node and the previous selection becomes gray. There were four trials of this task, with multiple orderings distributed to counteract possible ordering effects.

The second task, "B", also displayed a node-link diagram to the participant. Each participant completed two conditions of task A, one with a diagram size of 25 nodes and a 3-node path, and one with a diagram size of 60 nodes and an 8-node path (including start and end). The nodes are by default all colored in light gray. One node is colored dark blue, this represents the "start node". An "end node" is specified for the participant to locate. The participant is given the task of joining two nodes together by clicking on nodes in between. The start node and end node are both given in the task text. If the participant chooses to de-select a previously clicked node, they click on it again and its color reverts to gray.

Following each task, the participant filled out the NASA TLX cognitive load scale [HS88] and a brief questionnaire seeking feedback on the task difficulty and other comments. The process took about 30 minutes to 60 minutes for each participant.

Chapter 5

Results

5.1 Results

All results described here are from the Task "A" analysis. The independent variables being tested are user traits. I compared two groups for each trait, "high" vs. "low", where the "high" group scored above the median on the survey for that trait, and the "low" group scored below the median. Subjects whose scores fell directly on the median were excluded from analysis.

Using the UserAle.js library for JavaScript, the experiment application recorded all mouse and keyboard events as the participants interacted with the node-link diagrams in the experiment. This analysis compares counts of "node mouseover" events, defined as events logged when a participant moved their mouse cursor over a node in

Condition	*Metric*	*Trait*	*Extravert Mean*	*Introvert Mean*	*p*
Task A All	Unique Node Hovers	Extraversion	124.474	147.913	0.0497
Task A All	Node Re-Visit Percentage	Extraversion	56.790	51.870	0.1433
Task A All	Total Node Mouseovers	Extraversion	306.526	328.957	0.540
Task A All	Unique Node Hovers	Need For Cognition	134.632	139.522	0.6891

Table 5.1: A significant (before post-hoc testing) deviation in node hover amount was identified between high and low extraversion users, and a near-significant deviation in node re-visit percentage between high and low extraversion users. The overall difference in mouseover action amounts does not show a significant difference. No significant effect was found for unique node hovers and need for cognition. (n = 42 for all results.)

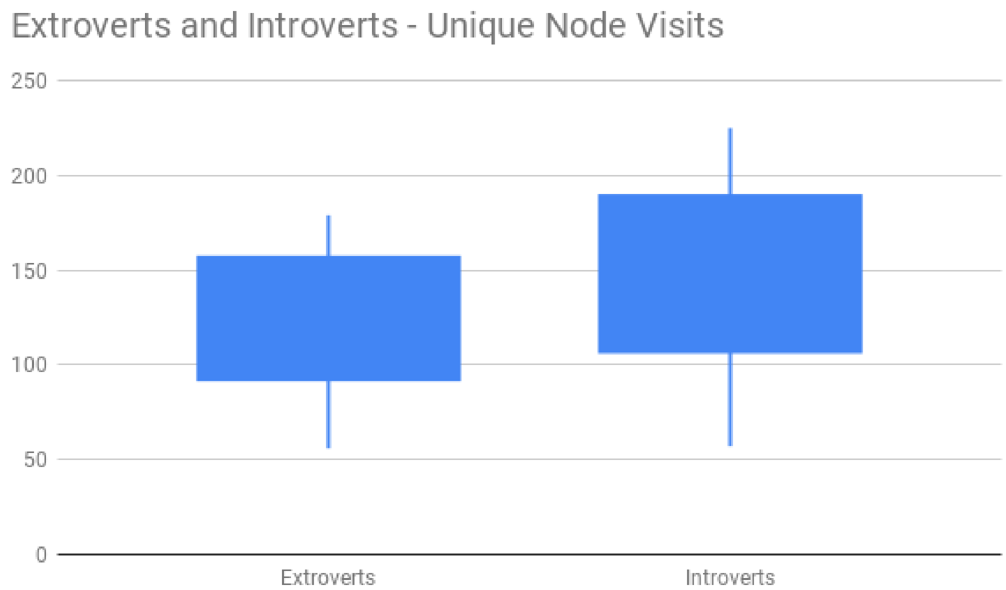

Figure 5.1: This boxplot shows the average count of node mouseover actions when the subject hovered over (viewed the word of) a new node that they had not previously visited. Extroverts tended to view new nodes significantly less often than introverts.

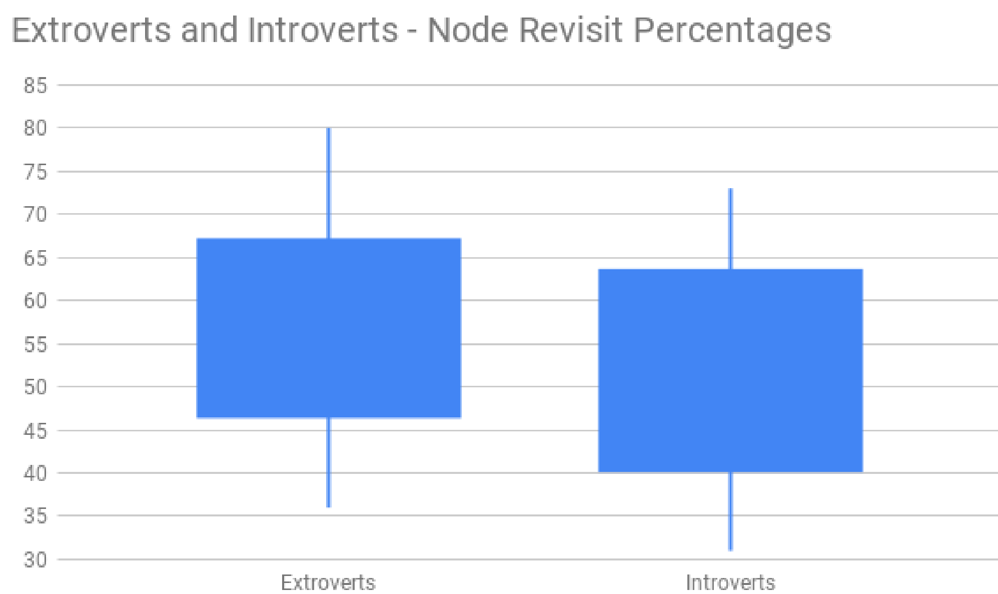

Figure 5.2: This boxplot shows average percentages out of the total mouseover count for re-viewing the same nodes. That is, out of the total number of mouseovers, what percentage were re-visits to the same nodes. Extroverts re-visited nodes more than introverts did, and a greater percentage of their total mouseovers were re-visits.

the diagram. "Unique" node mouseovers are node mouseover events where that particular participant has not yet viewed the particular node (word) being interacted with. "Re-visits" are defined as node mouseover events that "re-visit" a previously viewed node (word). In this analysis, I view the "re-visit percentage" as the percentage of the total count node mouseovers for a participant that are "re-visits" of previously visited nodes. When comparing the number of unique node mouseover events on nodes in the diagram, a significant (before post-hoc testing) difference was found between high (mean=124.474, SD=32.220) and low (mean=147.913, SD=41.089) extraversion users. (p=0.0497, n=42) (See Fig. 2). This finding indicates that the two groups can be discerned via clickstream data alone. The Bonferoni Correction sets the confidence interval from 0.05 to 0.025, which would show this result to be out of the range of statistical significance. The effect size of this finding is 0.635 (Cohen's d), which is considered slightly higher than a "medium" effect size [Coh88] [AS13], and to achieve a statistical power of 0.8, considering the corrected p-value of 0.025, would require a total of 98 participants, with at least 49 extraverts and 49 introverts. [sop]

Out of the total count of node mouseover actions performed by each user, I measured the percentage of those actions that were re-visits of nodes that the user had previously viewed. When comparing these percentages, a near-significant difference was found between high (mean=56.790, SD=10.223) and low (mean=51.870, SD=11.581) extraversion users. (p=0.1433, n=42) (See Fig. 3). The overall finding for mouseover action amounts shows no significant difference between groups, however, overall mouseover count as a metric does not contain much information about user search choices and strategy. As evidenced by this analysis, features that are indicative of the particular choices that users make (to look at previously viewed information versus seeking new information) show a greater disparity between trait groups, which allows traits like extraversion to be more easily distinguished from this data.

When comparing the number of unique node mouseover events, no significant difference was found between high (mean=134.632, SD=44.370) and low (mean=139.522, SD=34.285) need for cognition users. (p=0.6891, n=42). Relatively small sample size could explain this result, which is contrary to previous research and our hypothesis.

Data for 42 out of the 50 participants is considered in this analysis, with eight outliers excluded due to demonstrable task performance and data quality issues, as these outliers had significantly fewer logged interactions with the experiment and had notably lower quality task responses than others.

5.2 User Models

Linear regression models were constructed for two models: one based on unique node hovers, and one based on node re-visit percentage. Despite the significant difference observed between "high" and "low" groups, the correlation between extraversion and the metrics is not as strong as predicted. A Naive Bayes model can use both metrics to predict extraversion correctly at a rate of 59.5%, though it only divides users into "high" or "low" bins.

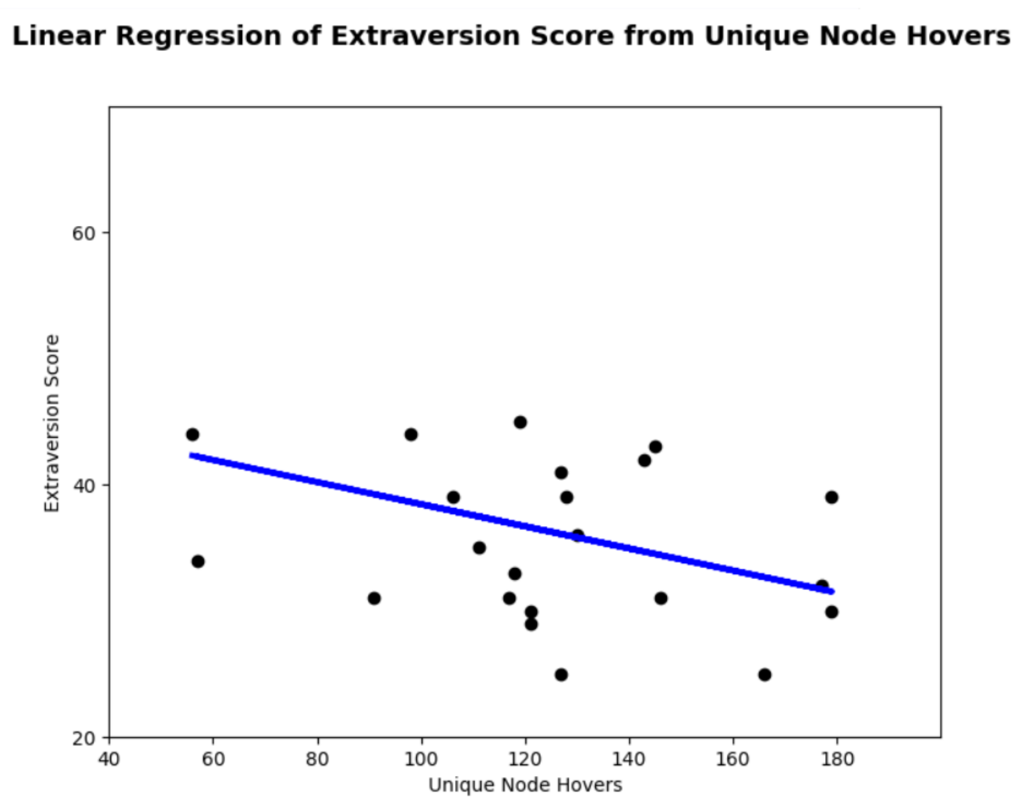

Figure 5.3: This linear regression takes in the unique node hovers as input and attempts to generate a best fit line for extraversion score.

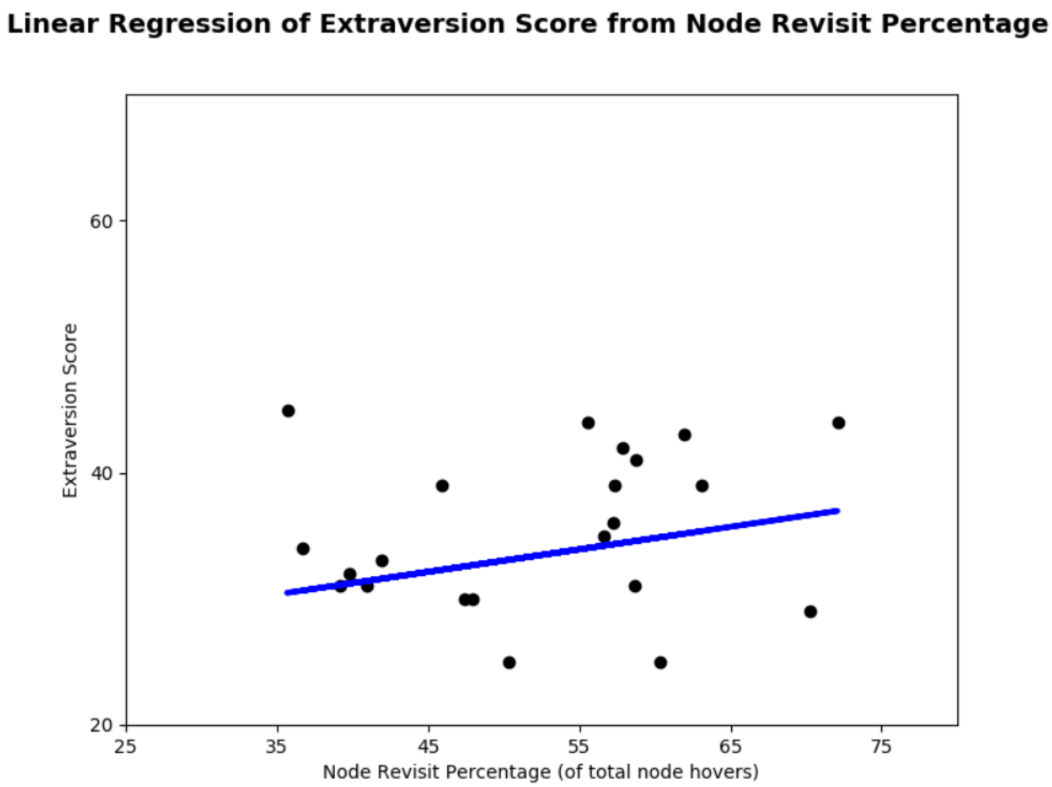

Figure 5.4: This linear regression takes in the node re-visit percentage as input and attempts to generate a best fit line for extraversion score.

Chapter 6

Discussion

The task shows how information foraging [Pir07] may vary in a subjective task. Users are free to choose whichever node they believe is most relevant to the given definition. Higher and lower extraversion users addressed this differently based on the number of mouseover events. The introverted users tended to seek out more unique nodes, even when the total mouseover count was relatively constant between the two groups. Previous research leads us to speculate that this difference is observable due to differences in how introverts and extroverts approach visual tasks [ZOC$^+$13]. Introverts may spend more time seeking out and evaluating different options.

In an applied setting, the amount of unique node mouseovers could be measured per unit time, such as every minute. Collecting this data would allow a user model, preferably trained on data from typical users, to predict a new user's extraversion.

There are multiple use cases for an extraversion model. One is the development of adaptive systems, featuring adaptations tailored to the user's extraversion. Split adaptive interfaces have been shown to have different utilization by users with different levels of extraversion, so an extraversion user model would be able to provide the user trait information needed for an adaptive interface to know how to adapt to which user [GC17].

In a more speculative sense, I believe that these findings may indicate some evidence that extraversion is linked to visual search interaction patterns, particularly relating to user choices when seeking new information. Though this evidence is not conclusive, the findings relating to unique node hovers and node re-visit percentage suggest that a measurable relationship between extraversion and these metrics does exist. With this work, my aim was to design an open-ended, exploratory, and highly ecologically valid task using node-link diagrams. In so doing, I found that datasets gathered from monitoring behaviors tend to be highly variable in nature, since the exact behaviors chosen are left up to the user. However, the amount of possible interactions was limited by my experimental interface. I believe that future work in this area should be focused on a variety of diverse visualization scenarios, including hierarchical data visualizations and scatterplots. Further experiments would help us to fully understand how search behavior can predict extraversion and possibly other traits.

In the case of need for cognition modeling, I would consider further analysis of interactions that indicate the user's interest level or level of engagement with the task. Since users with higher need for cognition tend to demonstrate a higher level of interest in tasks that require deep thought, my next hypothesis would be that higher need for cognition is able to be discerned through time spent and number of node mouseovers per unit time. However, more time spent doesn't always accurately represent more interest, so I would need to collect more data via a followup questionnaire asking the user to report their interest level, in order to analyze the correlation between time and interest. If the users who spent more time are generally also more interested, then I could use time as a standalone metric predicting need for cognition, but pending proper analysis, the relationship is currently not yet known.

Chapter 7

Conclusion

After completing this study, I observed a difference in interactions between high and low extraversion groups.

Being able to differentiate between high and low extraversion users gives an interesting insight on the link between extraversion, information foraging, and visualization. It shows that interactions are different among different people, and that the behavior stemming from this difference can be quantifiable and measurable enough to create a user model. This user model can help adaptive systems to understand users, providing a rich data source that feeds back into the system.

The creation of a user extraversion model requires further study into possible scenarios under which this result can be replicated. As well, more advanced machine learning techniques such as interaction n-gram analysis, as has been explored in previous work [BOZ$^+$14]. should be investigated for their use in model creation.

Continuing experiments will focus on the traits of extraversion and other traits to determine other differences that could be used for modeling. User models resulting from this research would be evaluated by integrating with adaptive software.

7.0.1 Acknowledgements

This work was supported by a grant from Draper.

Bibliography

[apa] Apache senssoft. `http://senssoft.incubator.apache.org/`.

[AS13] Milton Abramowitz and Irene A. Stegun. *Handbook of mathematical functions with formulas, graphs, and mathematical tables.* Dover Publ., 2013.

[big] Big-Five Factor Markers. `http://ipip.ori.org/newBigFive5broadKey.htm/#Extraversion`.

[BOH11] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. D3 data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, December 2011.

[BOZ+14] Eli Brown, Alvitta Ottley, Jieqiong Zhao, Quan Lin, Alex Endert, Richard Souvenir, and Remco Chang. Finding waldo: Learning about users from their interactions. *IEEE Conference on Visual Analytics (VAST)*, 2014.

[CC97] Chaomei Chen and Mary Czerwinski. Spatial ability and visual navigation: an empirical study. *New Review of Hypermedia and Multimedia*, 3(1):6789, 1997.

[Che00] Chaomei Chen. Individual differences in a spatial-semantic virtual environment. *Journal of the American Society for Information Science*, 51(6):529542, 2000.

[CHK+15] Jilin Chen, Eben Haber, Ruogu Kang, Gary Hsieh, and Jalal Mahmud. *Making use of derived personality: The case of social media ad targeting*, pages 51–60. AAAI Press, 2015.

[Coh88] Jacob Cohen. *Statistical power analysis for the behavioral sciences.* Lawrence Erlbaum Associates, 1988.

[CP82] John T. Cacioppo and Richard E. Petty. The need for cognition. *Journal of Personality and Social Psychology*, 42(1):116131, 1982.

[GC17] Krzysztof Z. Gajos and Krysta Chauncey. The influence of personality traits and cognitive load on the use of adaptive user interfaces. *Proceedings of the 22nd International Conference on Intelligent User Interfaces - IUI 17*, 2017.

[HS88] Sandra G. Hart and Lowell E. Staveland. Development of nasa-tlx (task load index): Results of empirical and theoretical research. *Advances in Psychology Human Mental Workload*, page 139183, 1988.

[KEC06] Ren Keller, Claudia M. Eckert, and P. John Clarkson. Matrices or node-link diagrams: Which visual representation is better for visualising connectivity models? *Information Visualization*, 5(1):6276, 2006.

[LSR09] M. Lanzenberger, J. Sampson, and M. Rester. Visualization in ontology tools. In *2009 International Conference on Complex, Intelligent and Software Intensive Systems*, pages 705–711, March 2009.

[Mil95] George A. Miller. Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41, November 1995.

[MML07] Scott W. Mcquiggan, Bradford W. Mott, and James C. Lester. Modeling self-efficacy in intelligent tutoring systems: An inductive approach. *User Modeling and User-Adapted Interaction*, 18(1-2):81123, 2007.

[OYC15] Alvitta Ottley, Huahai Yang, and Remco Chang. Personality as a predictor of user strategy. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI 15*, 2015.

[PCW01] Peter Pirolli, Stuart K. Card, and Mija M. Van Der Wege. Visual information foraging in a focus context visualization. *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI 01*, 2001.

[Pir07] Peter Pirolli. Stochastic models of information foraging by information scent. *Information Foraging Theory*, page 137146, Mar 2007.

[SCC13] Ben Steichen, Giuseppe Carenini, and Cristina Conati. User-adaptive information visualization. *Proceedings of the 2013 international conference on Intelligent user interfaces - IUI 13*, 2013.

[sop] A-priori sample size calculator for student t-tests [software]. `http://www.danielsoper.com/statcalc`. Accessed: 2018-04-22.

[WKS14] Wan-Ching Wu, Diane Kelly, and Avneesh Sud. Using information scent and need for cognition to understand online search behavior. In *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval*, SIGIR '14, pages 557–566, New York, NY, USA, 2014. ACM.

[ZOC+12] C. Ziemkiewicz, A. Ottley, R. J. Crouser, K. Chauncey, S. L. Su, and R. Chang. Understanding visualization by understanding individual users. *IEEE Computer Graphics and Applications*, 32(6):88–94, 2012.

[ZOC+13] C. Ziemkiewicz, A. Ottley, R. J. Crouser, A. R. Yauilla, S. L. Su, W. Ribarsky, and R. Chang. How visualization layout relates to locus of control and other personality factors. *IEEE Transactions on Visualization and Computer Graphics*, 19(7):11091121, 2013.