Cognitive Stopping Rules in a New Online Reality

Natalie Gerhart
Creighton University
NatalieGerhart@Creighton.edu

John Windsor
University of North Texas
John.Windsor@unt.edu

Abstract:

This research is a conceptual replication of a study by Browne, Pitts, and Wetherbe (2007) that explores information stopping rules in an online search context. Information stopping rules consider the cognitive reasons decision makers determine when enough information is collected to make a decision. Previous research outlines five stopping rules decision makers use and applies them in different decision context. The original research considers three information search tasks (search for a television, map, and job) and hypothesizes the relationship between structure of the task and the stopping rule employed. This research replicates that study in a new information environment with new search methodologies and technology. We find that structured tasks use similar stopping rules to the original study; however further analysis shows distinct differences in the nature of the two tasks presented. Poorly structured tasks potentially involve the use of different stopping rules than previously determined. The updated findings suggest information systems used for poorly structured search tasks might also benefit from highlighting the uniqueness of information in order to encourage a user to continue searching for information.

Keywords: IT use, Surveys, Decision Making, Online Search Behavior

1   Introduction

Thanks, in large part, to the Internet there is unprecedented access to information for individual decision makers. While more information provides more insights, it can also become a hindrance as people approach cognitive limitations to information processing (Schroeder, Driver, & Streufert, 1967). When individuals engage in the decision making process, there are two distinct decision points: deciding to stop collection of information and making the ultimate decision from the accumulated alternatives (Browne et al., 2007). Technology is changing the accessibility of information for decision makers and therefore affecting the search process. As the amount of available information grows, understanding this first decision point in the area of information collection becomes increasingly important.

Studying the decision process is very difficult, as it is not an easily visible process external to the decision maker (Robey & Taggart, 1981). One approach for considering information collection is centered on information search stopping rules. In Information Systems (IS) literature, a large amount of work revolves around four stopping rules originally outlined by Nickles et al. (1995) and expanded by Browne and colleagues (Browne & Pitts, 2004; Browne, Pitts, & Wetherbe, 2005; Browne et al., 2007; Pitts & Browne, 2004).

The five stopping rules employed in prior research each have distinct features. Originally, two rules were identified as judgment rules, during which the decision maker must make a determination to stop collecting information based on the actual information discovered (Nickles et al., 1995). Difference threshold is the first of these rules, which states the decision maker has enough information when there is no longer new value from additional information (Nickles et al., 1995). Essentially, the decision maker makes a determination, while in the act of gathering information, that the new information is not providing new insight. The second judgment rule is magnitude threshold which involves reaching a pre-determined threshold of the amount of information collected (Nickles et al., 1995). The amount of information the decision maker needs is pre-determined, but the judgment is made in the moment of when there is "enough" information.

The second two rules are reasoning-based, or limited by cognitive capacities of individuals, instead of based on the information itself (Nickles et al., 1995). Mental list requires the decision maker to make a list of criteria that he or she is interested in before searching, and when information is collected on these criteria, the search process is complete (Nickles et al., 1995). Fundamentally this is a checklist of the information that needs to be collected. Finally, representational stability is an abstract rule that a decision maker stops searching when the understanding of the decision situation is no longer changing as new information is acquired (Browne et al., 2007). The key to representational stability is the abstract overall view of the decision situation is stagnant as information is gathered.

Based upon further research, a fifth stopping rule developed, termed single criterion (Browne et al., 2005). This rule specifies the need to perform a search on only one criterion before making an ultimate decision between alternatives (Browne et al., 2005). These five stopping rules identified are all cognitive; however other motivational factors might affect the information collection process (Prabha, Connaway, Olszewski, & Jenkins, 2007).

One of the most comprehensive studies in this area involved three different search contexts: search for a television, a map, and a job (Browne et al., 2007). This study opened new doors to understanding the use of different stopping rules in different search contexts; however much has changed in the information search process since 2007. As an example, there were a little over 122 million smartphones sold in 2007, while now that number has grown tenfold (Statista, 2016). As a result of technological advances, younger generations search for information using new methodologies (American Press Institute, 2015).

With the intent of gaining new understanding of information search processes with regards to stopping rules, this research conceptually replicates the well-documented and insightful Browne et al. (2007) research study. Replication research offers confidence in prior theoretical foundations as the paradigm grows (Dennis & Valacich, 2014). Replication is particularly useful if there are changes in the technology over time or large changes in technology awareness of the user (Niederman & March, 2015). With this in mind, we seek to extend the understanding of stopping rules in online search contexts.

The remainder of this paper is laid out as follows. First, the foundations of the original study are presented. Next, the methodology and changes to the original study are outlined. Third, the results of the replication study are presented. Fourth, we outline the areas where our research confirms and expands on the prior research. Finally, we discuss limitations and conclusions of this replication.

2 Overview of Original Research

Based on decision-making theory and cognitive stopping rule backgrounds, Browne et al. (2007) conducted an experiment involving three online search tasks with 115 student participants. The previous research hypothesizes differences between well-structured and poorly structured tasks. Well-structured tasks are identified as searches that are easily understood by the participant, likely due to prior experience conceptualizing what inputs, operations, and outputs are needed (Browne et al., 2007). The authors hypothesize well-structured tasks will use mental list and single criterion stopping rules more than the other three stopping rules (hypothesis 1). Secondly, poorly structured tasks will use magnitude threshold and representational stability more than the other stopping rules (hypothesis 2). Poorly structured tasks involve search situation that are complex and unfamiliar to the decision maker (Browne et al., 2007). The research uses two well-structured tasks: search for a 32-inch television on BestBuy.com and search for a job with Amazon.com. For an unstructured task, participants were asked to find a map of the battlefield of the Battle of Fallen Timbers, a purposefully obscure task. After finding the map, participants were required to draw the battlefield from memory.

The research finds that individuals use different stopping rules based on the nature of the decision situation by finding some support for both hypotheses. Specifically, mental list is used more than the other stopping rules on a structured task. Single criterion is used more than difference threshold on both tasks. The television search task does not find significant differences between single criterion and both representational stability and magnitude threshold, resulting in partial support. The job search task does not significantly use single criterion more than magnitude threshold.

Similarly, hypothesis 2 is mostly supported. Magnitude threshold is used more than the other stopping rules. The only non-significant finding is that representational stability is not significantly used more than mental list on the map search task. Overall, the study shows that different stopping rules are applicable in different search situations.

3 Methodology

In the original study, as well as this replication, participants were students from a large research university. The demographics are similar, although slightly different because it is extremely difficult to get a sample that will be the same in a replication study (Niederman & March, 2015). While original demographic reporting is limited, the participants were 69% male, while this replication study contains 51.6% of those who chose to specify gender were male. In this study, 47.6% were female, while originally 31% were female. In the replication study, one participant specified gender as other. 97.6% of participants are between the ages of 18-34, similar to the original study. Finally, while the original study indicates both undergraduate and MBA students participated, this study included only two graduate students.

Unless otherwise stated, all research methods of the Browne et al. (2007) study were mimicked as closely as possible. A major difference in the prior study and the current study is the laboratory setting. Due to the changing nature of technology, the current study was administered online through Qualtrics survey software instead of using hard copy experiment booklets. Therefore, participants were asked to provide links to their decisions, and in the case of the map, participants were asked to "describe the battlefield in enough detail so that a person could accurately draw the battlefield of the Battle of Fallen Timbers without ever seeing a picture" instead of drawing the map. This change in procedure likely had a direct impact on the resulting usable responses because participants that are not monitored in a lab can easily quit the task without anyone knowing.

 The original two survey questions of "Why did you stop searching for information when you did?" and "How did you decide to stop searching for information?" were both maintained. These questions are purposefully redundant to gain more complete insights of what the decision maker was thinking (Browne et al., 2007).

Slight modifications were made to the tasks to modernize the search. Instead of a 32-inch television, participants were asked to search BestBuy.com for a 4K television. Instead of searching Amazon.com for a job, participants were asked to search Monster.com. With these revisions, the intent was to make the search tasks seem more natural and realistic, not change the fundamental task.

As in the original study, all participants performed all three tasks and the order was randomized for each participant. There were no time limits on the search process to allow participants to determine their own

stopping rule. Participants received course credit for participating and were informed non-serious answers would be removed before credit was assigned with the intent of motivating an effort for the search. A total of 271 students started the survey; however, after removing non-serious responses and incomplete responses (not completing all three tasks), a total of 124 usable responses remained.

To determine which stopping rule was used by the participant, three independent researchers were trained to code the responses. Primarily, the coding came from two methods: first, a through reading of the original Browne et al. (2007) study which includes detailed descriptions and examples; secondly, a verbal description of each stopping rule with real life examples and discussion to allow for questions. Following the training, the coders read through each response and coded the responses into one of the five stopping rule categories. If the response was not clearly one of the stopping rules, an "other" category was also an option. Testing for interrater reliability of the initial coding of all three judges results in an AC1 statistic of 0.396, which indicates fair agreement among raters (Landis & Koch, 1977). While this is acceptable, more rigorous training of the coders may have improved results.

While this is not a bad interrater reliability, agreement was necessary between at least two of the three coders to analyze which rule was in use. If there was not agreement among two of the coders, the coders were asked to look at their responses again to confirm the coding, with the knowledge of the other coders coding. With this knowledge, at least two of the three coders were able to come to an agreement about how discrepancies should be resolved. Having three independent coders individually go through each response insures robust and reliable coding, as in the original study.

4   Results

The frequency of responses for this study (2017) as compared to the prior study (2007) are displayed in Tables 1 and 2. These results on the surface level appear to have similar patterns to the original study. To test the hypotheses, the same method as the original study was used, binomial test of significance which presumes all five stopping rules are equally likely (Browne et al., 2007).

Table 1. Frequency of Stopping Rule Used											
	TV Frequency				Map Frequency				Job Frequency			
	2017		2007		2017		2007		2017		2007	
DT	7	5.6%	0	0.0%	14	11.3%	8	7.0%	3	2.4%	1	0.9%
MT	11	8.9%	15	13.0%	71	57.3%	56	48.7%	21	16.9%	8	7.0%
ML	67	54.0%	63	54.8%	7	5.6%	14	12.2%	62	50.0%	63	54.8%
RS	9	7.3%	10	8.7%	23	18.5%	23	20.0%	7	5.6%	3	2.6%
SC	30	24.2%	18	15.7%	7	5.6%	3	2.6%	21	16.9%	14	12.2%
O	0	0.0%	9	7.8%	2	1.6%	11	9.6%	10	8.1%	26	22.6%
Total	124	100%	115	100%	124	100%	115	100%	124	100%	115	100%
DT= Difference Threshold, MT=Magnitude Threshold, ML=Mental List, RS=Representational Stability, SC= Single Criterion, O=Other 2007= Browne et al. Study, 2017= Current Study											

Table 2. Total Frequencies of Stopping Rule Used for All 3 Tasks				
	2017		2007	
Difference Threshold	24	6.5%	9	2.6%
Magnitude Threshold	103	27.7%	79	22.9%
Mental List	136	36.6%	140	40.6%
Representational Stability	39	10.5%	36	10.4%
Single Criterion	58	15.6%	35	10.1%
Other	12	3.2%	46	13.3%
Total	372	100%	345	100%
2007= Browne et al. Study, 2017= Current Study				

The first hypothesis argues that mental list and single criterion will be used more frequently than the other three stopping rules on well-structured tasks (Browne et al., 2007). Using a binomial test of significance (Daniel, 1990) with a probability below 0.4 (because each rule is equally likely and there are two stopping

rules considered), the results indicate both the job and the television task result in a significant difference (exact binomial (one-tailed) p<0.001). In hypothesis two, magnitude threshold and representational stability are hypothesized as most likely in unstructured tasks. The original research found significance, as did this research (exact binomial (one-tailed) p<.001).

To further analyze individual stopping rules, as in the original study, this research tested differences between individual stopping rules for each task. The results are in tables 3 (television task), 4 (job task), and 5 (map task). Probability of the binomial test for these was set at 0.5 because in each instance there is one of two results that are equally likely (Browne et al., 2007). All significance tests assume that α=0.05. Because of the concern for false conclusions associated with binomial testing, we further included the false discovery rate (FDR), as did the original study, which controls type I errors while reducing type II errors (Verhoeven, Simonsen, & Mcintyre, 2005). Using the FDR test, the results are ordered from lowest to highest p-value, then the FDR threshold is calculated to determine significance (Benjamini & Hochberg, 1995).

Table 3. TV Search Task

Test	Exact Binomial One-Tailed p-Value	FDR Threshold	Significance	Same Result
ML>DT	<.0001	0.0083	Yes	Yes
ML>MT	<.0001	0.0167	Yes	Yes
ML>RS	<.0001	0.0250	Yes	Yes
SC>DT	<.0001	0.0333	Yes	Yes
SC>MT	0.019	0.0417	Yes	No
SC>RS	0.019	0.0500	Yes	No

Table 4. Job Search Task

Test	Exact Binomial One-Tailed p-Value	FDR Threshold	Significance	Same Result
ML>DT	<.0001	0.0083	Yes	Yes
ML>MT	<0.001	0.0167	Yes	Yes
ML>RS	<.0001	0.0250	Yes	Yes
SC>DT	0.092	0.0333	No	No
SC>RS	0.180	0.0417	No	No
SC>MT	1.000	0.0500	No	Yes

Table 5. TV Search Task

Test	Exact Binomial One-Tailed p-Value	FDR Threshold	Significance	Same Result
MT>DT	<.0001	0.0083	Yes	Yes
MT>ML	<.0001	0.0167	Yes	Yes
MT>SC	<.0001	0.0250	Yes	Yes
RS>ML	0.007	0.0333	Yes	No
RS>SC	0.035	0.0417	Yes	Yes
RS>DT	0.248	0.0500	No	No

5   Discussion

Hypothesis 1 considers structured tasks are more likely to use mental list and single criterion stopping rules, which we find the same results as the prior study almost ten years ago (Browne et al., 2007). The mental list and single criterion stopping rules are most useful in well-structured tasks because they involve a priori

understanding of the decision situation (Browne et al., 2005). Despite changes in information available and the physical capabilities of the tools, overall stopping rules employed have not changed. This reflects the structured nature of these tasks. The search for a television, and most consumer products, involves a similar search routine whether it as an online or offline search. Similarly, the search process for employment is a task that follows a certain set of logical steps. The ability to acquire information online is improving, however the basic criteria remains the same, and therefore the schema is the same as a decade ago. Structured tasks follow a similar structure as they have for decades, and thus are not changed by improvements in search technology.

Hypothesis 2 considers which stopping rules are used in poorly structured tasks, and overall the findings of the original study are supported. Tasks that are poorly structured do not necessarily have a mental model pre-defined for the decision maker. As a result, the search process might be more obscure or less linear (Browne et al., 2005). Logically, if a task is less structured, there is more freedom for how a person will perform a search. As a result, the stopping rules that follow these tendencies are used more.

The interesting findings of this replication lie in the changes between individual stopping rules in both hypothesis 1 and 2. By looking closer at the comparative use of individual stopping rules, there are changes in the findings. Specifically, for the television search task, single criterion is significantly used more than magnitude threshold and representational stability, which is different than the original findings. This suggests that more people know that they are looking for one feature, specifically in a product search task, than in prior times. This could be easily understood by considering that people are more adept at online shopping. According to the US Census Bureau, e-commerce sales were $136.4 billion in 2007, making up just 3.4% of total retail sales (Scheleur, King, & Kinyon, 2008). In the third quarter of 2016, e-commerce sales make up 7.7% of total retail sales, more than doubling 2007 numbers (DeNale & Wiedenhamer, 2016). This suggests online shopping is an even more structured task than it was in 2007 and people might know just what to look for. The results of this study indicate that people begin an online shopping task with a lot of information either pre-determined, or deemed irrelevant, suggesting a more hurried search process. This shift could have major implications from a brand loyalty perspective. For example, if companies can create strong brand loyalty, online decision makers might just search within one brand for the best price, or the dimensions they desire. Knowing customers might only search on one criteria, and solidifying a brand loyalty with those customers, might keep loyal customers from performing a more in-depth search.

Interestingly, the other structured task, the job search task, indicated different results. On this task, single criterion is not statistically used more than any of the other stopping rules. This means that the use of mental list is drastically more than single criterion, allowing for the test of hypothesis 1 to still remain significant; however single criterion itself is no longer a useful stopping rule for this task. While the television search task was simplified to one key criterion, the job search task instead requires a list of qualifiers. This could be due to the nature of the outcome of the task. Finding a worthwhile job is important to a person, and they have several criteria to meet. Intuitively, single criterion might make sense if the only concern for a job seeker is location; however, students graduating from college are often more interested in other fulfilling factors and are often not location bound. Finding a television is much simpler, because you will likely have multiple televisions and stick to what you want. Considering both tasks are structured tasks, further research should be done to differentiate different types of tasks, such as professional versus personal tasks, complex versus simple tasks, and tasks with different levels of familiarity to the decision maker.

For hypothesis 2, the overall results confirm the findings, which is likely partially due to the use of magnitude threshold. This stopping rule allows a decision maker to search for information until "enough" information is collected, making the criterion relatively broad and abstract. Thus, this stopping rule is easily able to scale with the amount of information that is available. Because the amount of available information has changed the determination of how much is enough may have changed as well, but to an information seeker now or in the past, the term "enough" is vague so it has also potentially changed in definition. Therefore, it is predictable that this stopping rule is still commonly employed for a poorly structured task.

Contrasting the original study, the results of this research indicate that difference threshold might be a useful stopping rule for unstructured information search tasks because difference threshold is not significantly different from representational stability in a binomial test. As more information becomes available, stopping rule implementation must evolve as well. The difference threshold rule requires the evaluation of the individual pieces of information against each other (Browne et al., 2005). As the decision maker collects the information, he or she will determine if each piece adds value to the list of available alternatives or not, and if not, stop collecting information (Browne et al., 2007). The results of this analysis suggest people may be

using this stopping rule with more frequency. The sheer age of the Internet now allows for more citation amongst different websites, thus creating circular references available online and similar information across sites. Difference threshold could possibly be a more used stopping rule simply because there is more uniform information available on multiple sites and low switching costs for using different sites.

This new finding has implications for the design of information search tools. First, the stopping rules used to collect information for well-structured tasks are similar to prior generations of information search processes; however closer analysis shows that grouping these as simply structured tasks does not provide full understanding. Potentially, future research should consider the user's experience with other tasks. On the other hand, to encourage users to continue searching for information in a poorly structured task, the design of search systems should clearly distinguish the information on how it provides unique value. These recommendations could be implemented several different ways. One example is to focus titles, headings, and substructures of a website based on the purpose of the website. If it is a site that is primarily used when people are performing a structured search, such as buying a television, quickly highlighting these features, as well as returning them in the search results, could prove most useful.

Secondly, search engines could be designed to more intuitively determine the nature of the search. Generally, when search results are returned, the brief descriptions are very similar, and sometimes the same, for different sites. If the information seeker searched on a topic that lends itself to an unstructured topic, it might be helpful for the search engine itself to highlight information that is different between sites. As part of the algorithm, the search process is the same, and returns results that might be very similar; however, the display to the information seeker could be altered to highlight differences. While this would be a challenge to determine the nature of the search, it is not outside the scope of future technologies.

6   Limitations and Future Research

This replication research is subject to certain limitations. Primarily, the methodology of this replication could be related to slight differences in results. Participants were asked to imagine drawing a map, instead of actually drawing a map. While this was a useful update on the methodology, the ability to actually draw the map might have been lost in this replication.

As with the original study, motivation is a concern for this study. While course credit points offer some motivation, true motivation is not clear. Since the tasks are concerned with when a person gives up on a process, motivation is particularly important. In the sake of replication, the methodology and tasks were kept as similar as possible; however, future research should seek tasks that provide some intrinsic motivation.

Determining information used during the decision process is a neglected but increasingly valuable area of research (Savolainen, 2009). As indicated by this research, several more responses were coded as "other" in this study than the original, suggesting other stopping rules might be in use, such as simple fatigue. Current Internet users are likely to believe that search engines are dependable, and therefore, indicated that if they did not find a satisfactory solution, it must not exist. Thus, future studies should explore potential new stopping rules that develop from new expectations and technologies. Further, the methodology of open-ended responses is limited by the ability of the coders, therefore, determining other operationalizations of these identified stopping rules could prove valuable. Finally, a better understanding of the differences between tasks would prove useful in understanding how people determine when to stop collecting information.

7   Conclusion

This study is not an exact replication, as slight modifications to the methodology were made, but it is a very close conceptual replication as outlined by Dennis and Valacich (2014). While small changes were made, the results indicate confirmation of prior understanding of stopping rule use in structured tasks as well as new insights about information stopping rules in unstructured tasks. The constant access to information via the Internet is creating a new information environment that individuals must now cope with to perform efficient searches. Understanding how the user determines when enough information is gathered to make a decision can provide design insights on multiple systems. Our results show that the information stopping rules employed do change over time for poorly structured tasks and that there are other factors that distinguish structured tasks. As a result, the systems employed must also evolve to help information consumers digest the most relevant information as efficiently as possible.

References

American Press Institute. (2015, March 16). How millennials get news. Retrieved from https://www.americanpressinstitute.org/publications/reports/survey-research/millennials-news/.

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1), 289–300.

Browne, G. J., & Pitts, M. G. (2004). Stopping rule use during information search in design problems. *Organizational Behavior and Human Decision Processes*, 95(2), 208–224.

Browne, G. J., Pitts, M. G., & Wetherbe, J. C. (2005). Stopping rule use during web-based search. In Proceedings of the 38th Annual Hawaii International Conference on System Sciences, 1–6.

Browne, G. J., Pitts, M. G., & Wetherbe, J. C. (2007). Cognitive stopping rules for terminating information search in online tasks. *MIS Quarterly*, 31(1), 89–104.

Daniel, W. W. (1990). *Applied nonparametric statistics* (2nd Ed.). Boston: PWS-Kent.

DeNale, R., & Wiedenhamer, D. (2016). Quarterly retail e-commerce sales 3rd quarter 2016. *US Census Bureau News*.

Dennis, A. R., & Valacich, J. S. (2014). A replication manifesto. *AIS Transactions on Replication Research*, 1(1), 1–5.

Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159–174.

Nickles, K. R., Curley, S. P., & Benson, P. G. (1995). Judgment-based and reasoning-based stopping rules in decision making under uncertainty, Working Paper, Wake Forest University, October 1995.

Niederman, F., & March, S. (2015). Reflections on Replications. *AIS Transactions on Replication Research*, 1(1), 1–16.

Pitts, M. G., & Browne, G. J. (2004). Stopping behavior of systems analysts during information requirements elicitation. *Journal of Management Information Systems*, 21(1), 203–226.

Prabha, C., Connaway, L. S., Olszewski, L., & Jenkins, L. R. (2007). What is enough? Satisficing information needs. *Journal of Documentation*, 63(1), 74–89.

Robey, D., & Taggart, W. (1981). Measuring managers' minds: the assessment of style in human information processing. *Academy of Management Review*, 6(3), 375–383.

Savolainen, R. (2009). Information use and information processing: Comparison of conceptualizations. *Journal of Documentation*, 65(2), 187–207.

Scheleur, S., King, C., & Kinyon, D. (2008). Quarterly retail e-commerce sales 4th quarter 2007. *US Census Bureau News*.

Schroeder, H. M., Driver, M. J., & Streufert, S. (1967). Human information processing: Individuals and groups functioning in complex social situations. New York: Holt, Rinehart and Winston, Inc.

Statista. (2016). Global smartphone sales to end users 2007-2014. Retrieved February 2, 2016, from http://www.statista.com/statistics/263437/global-smartphone-sales-to-end-users-since-2007/

Verhoeven, K. J. F., Simonsen, K. L., & Mcintyre, L. M. (2005). Implementing false discovery rate control: Increasing your power. Oikos, 108(3), 643–647.

About the Authors

Natalie Gerhart. is an Assistant Professor of Information Systems at Creighton University. She received her Ph.D. from the Information Technology and Decisions Sciences department at the University of North Texas. She also holds an MBA from the University of Missouri-Columbia in Marketing Analytics and a Bachelor of Science degree from Truman State University in Management Information Systems and Marketing. Previously, she has published in *MISQ Executive*, *Decision Support Systems*, *Journal of Computer Information Systems*, and the *Decision Sciences Journal of Innovate Education*. Her research interests include decision stopping rules, human computer interaction (HCI), business intelligence and analytics (BI&A), and social networking.

John Windsor. is a professor of Information Systems, former Department Chair, and former Director of the Information Systems Research Center at the University of North Texas. He received his Ph.D. from Georgia State University. He has published over six books and 60 articles in such journals as *Data Base*, *IIE Transactions*, *Information & Management* and *Computers & Security*. His research interests include software and data engineering, systems security, collaborative computing, and the organizational impact of information technology.