

# Toward Foraging for Understanding of StarCraft Agents: An Empirical Study

Sean Penney, Jonathan Dodge, Claudia Hilderbrand,  
 Andrew Anderson, Logan Simpson, Margaret Burnett  
 Oregon State University, Corvallis, OR, 97331 USA  
 { penneys, dodgej, minic, anderan2, burnett }@eecs.oregonstate.edu

## ABSTRACT

Assessing and understanding intelligent agents is a difficult task for users that lack an AI background. A relatively new area, called “Explainable AI,” is emerging to help address this problem, but little is known about how users would forage through information an explanation system might offer. To inform the development of Explainable AI systems, we conducted a formative study — using the lens of Information Foraging Theory — into how experienced users foraged in the domain of StarCraft to assess an agent. Our results showed that participants faced difficult foraging problems. These foraging problems caused participants to entirely miss events that were important to them, reluctantly choose to ignore actions they did not want to ignore, and bear high cognitive, navigation, and information costs to access the information they needed.

## CCS Concepts

•Human-centered computing → User studies;  
 •Computing methodologies → Intelligent agents;

## Author Keywords

Intelligent Agents; Explainable AI; Intelligibility; Content Analysis; Video Games; StarCraft; Information Foraging

## INTRODUCTION

Real-time strategy (RTS) games are a popular test bed for artificial intelligence (AI) research, and platforms supporting such research continue to improve (e.g., [41]). The RTS domain is challenging for AI due to real-time adversarial planning requirements within sequential, dynamic, and partially observable environments [30]. Since these constraints transfer to the real world, improvements in RTS agents can be applied to other domains, for example, mission planning and execution for AI systems trained to control a fleet of unmanned aerial vehicles (UAVs) in simulated environments [38]. However, the intersection of two complex domains, such as AI and flight, poses challenges: who is qualified to assess behaviors of such

a system? For example, how can a domain expert, such as a flight specialist, assess whether the system is making its decisions *for the right reasons*?

If a domain expert making such assessments is not an expert in the complex AI models the system is using, there is a gap between the knowledge they need to make such assessments vs. the knowledge they have in the domain. To close this gap, a growing area known as “Explainable AI” aims to enable domain experts to understand complex AI systems by requesting explanations. Prior work has shown that such explanations can improve mental models [17, 19], user satisfaction [15], and users’ ability to effectively control the system [1, 3, 18].

However, little is known about what an RTS domain expert’s information needs are – what they need to have explained, in what sequence, and at what cognitive and time costs. Therefore, to inform explanation systems in this area, we conducted a formative study of how experienced RTS players would go about trying to understand and assess an intelligent agent playing the RTS game of StarCraft.

Our setting was StarCraft replay files. A StarCraft replay file contains an action history of a game, but no information about the players (i.e., no pictures of players and no voice audio). This anonymized set-up enabled us to tell our participants that one of the players was an AI agent. (We detail this design further in the Methodology section.) In addition, the



Figure 1. A screenshot from our study, with participants anonymized (bottom right corner). Superimposed red boxes point out: (1, bottom left) the *Minimap*, a birds-eye view enabling participants navigate around the game map; (2: top left) a drop-down menu to display the *Production tab* for a summary of build actions in progress; (3, middle right) *Time Controls* to rewind/forward or change speed.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

IUI 2018, March 7–11, 2018, Tokyo, Japan.

Copyright © 2018 ACM ISBN 978-1-4503-4945-1/18/03 ...\$15.00.

<http://dx.doi.org/10.1145/3172944.3172946>

participants had functionality to seek additional information about the replay, such as navigating around the game map, drilling down into production information, pausing, rewinding, fast-forwarding, and so on (Figure 1).

However, we wanted a higher level of abstraction than features specific to StarCraft. Specifically, we aimed for (1) applicability to other RTS environments, and (2) connection with other research about humans seeking information. To that end, we turned to Information Foraging Theory (IFT).

IFT has a long history of revealing useful and usable information functionalities in other information-rich domains, especially web environments (e.g., [34]) and software development environments (e.g., [9, 32]). Originally based on classic predator-prey models in the wild, its basic constructs are the *predator* (information seekers like our participants) seeking *prey* (information goals) along pathways marked by *cues* (signposts) in an *information environment* (such as the StarCraft replay environment). The predator decides which paths to navigate by weighing the expected cost of navigating the path against the expected value of the location to which it leads.

Drawing from this theory, we investigated how people would assess a (purported) intelligent agent, via the following research questions (RQs):

**RQ1** *The Prey*: What kind of information do domain experts seek, how do they ask about it, and for what reasons?

**RQ2** *The Foraging Paths*: What paths do domain experts follow in seeking their prey, why, and at what cost?

**RQ3** *The Decisions and the Cues*: What decision points do domain experts consider to be most critical, and what cues lead them astray from these decision points?

### BACKGROUND AND RELATED WORK

When assessing whether an AI agent is making its decisions for the right reasons, humans automatically develop mental models of the system [29]. Mental models, defined as “internal representations that people build based on their experiences in the real world,” enable users to predict system behavior [29].

Ideally, mental models of a system would help people gain the understanding they need to assess an AI agent, but this is not always the case. Tullio et al. [39] examined mental models for a system that predicted the interruptibility of their managers. They found that the overall structure of their participants’ mental models was largely unchanged over the 6 week study, although they did discount some initial misconceptions. However, their study did not deeply engage in explanation; it was mostly visualization. In other work, Bostandjiev et al. [3] studied a music recommendation system and found that explanation led to a remarkable increase in user-satisfaction. In an effort to improve mental models by increasing transparency of a machine learning system, Kulesza et al. [17] identified principles for explaining (in a “white box” fashion) how a machine learning based system makes its predictions more transparent to the user. In their study, participants using a prototype following the principles observed an improvement of their mental model quality by up to 52%.

Several studies have also found that explanations have been able to improve users’ ability to actually *control* the system.

Stumpf et al. [37] investigated how users responded to explanations of machine learning predictions, finding that participants were willing to provide a wide range of feedback in an effort to improve the system. Kulesza et al. [18] found that the participants who were best able to customize recommendations were the ones who had adjusted their mental models the most in response to explanations about the recommender system. Further, those same participants found debugging more worthwhile and engaging. Kapoor et al. [15] found that interacting with explanations enabled users to construct classifiers that were more aligned with target preferences, alongside increased satisfaction. Beltran et al. [1] presented a novel gestural approach to querying text databases, allowing users to refine queries by providing reasons why the result was correct or incorrect. Their results indicated that action explanation allowed for more efficient query refinement.

However, in the domain of intelligent agents in RTS games, although there is research into AI approaches [30], there is only a little research investigating what *humans* need or want explained. Cheung et al. [5] studied how people watch the RTS genre, creating personas for various types of viewers. Metoyer et al. [26] studied how experienced players explained the RTS domain to novice users while demonstrating how to play the game. Finally, McGregor et al.’s [25] work is also pertinent, describing support for testing and debugging in settings with thousands of decisions being made sequentially. The work most similar to our own is Kim et al.’s [16] study of intelligent agent assessment in StarCraft. Their study invited experienced players to assess skill levels and overall performance of AI bots by playing against them. They observed that the humans’ ranking differed from an empirical ranking based on the bots’ win rate at AI competitions. Our study differs from theirs in that our participants did not play, but instead strove to *understand and explain* by interacting with a game replay.

In everyday conversation, people obtain explanations by asking questions. Drawing upon this point, Lim et al. [23] have categorized questions people ask about AI systems in terms of “intelligibility types.” In their work, they investigated participants’ information demands about context-aware intelligent systems powered by decision trees, determining which explanation types provided the most benefit to users. They found the most often demanded questions were *Why* and *Why not* (Why did or didn’t the system do X?). We provide more details of that work and build upon it in when we discuss RQ1 results.

In recognition of the particular importance of these two types of questions, researchers have been working on *Why* and *Why not* explanations in domains such as database queries [2, 14], robotics [13, 24, 35], email classification [20], and pervasive computing [40]. These types of explanations have also attracted attention from the social sciences, which seek to help ground AI researchers’ efforts in cognitive theories [27]. Other research has demonstrated that the intelligibility type(s) the system supported impact which aspects of users’ attitudes are affected. For example, Cotter et al. [7] found that justifying *why* an algorithm works the way it does (but not *how* it works) increased users’ confidence (blind faith) in the system — but not for improving their *trust* (beliefs which inform a full

cost-benefit analysis) in the system. Further, it seems that the relative importance of the intelligibility types may vary from one domain to another. For example, Castelli et al. [4] found that in the smart homes domain, users showed a strong interest in *What* questions, but few of the other intelligibility types.

We drew upon Information Foraging Theory (IFT) to investigate the information that people would seek in the RTS domain. In IFT terms, when deciding where to forage for information, predators (our participants) make cost/benefit estimates, weighing the information value per time cost of staying in the current *patch* (location on the game map or tab with supplemental information) versus navigating to another patch [34]. However, predators are not omniscient: they decide based on their *perceptions* of the cost and value of the available options. Predators form these perceptions using their prior experience with similar patches [32] and the *cues* (signposts in their information environment like links and indicators) that point toward various patches. Of course, predators' perceived values and costs are often inaccurate [33].

IFT constructs have been used to understand humans' information-seeking behavior in other domains, particularly web navigation [6, 10], debugging [9, 21, 32], and other software development tasks [28, 31, 33, 36]. However, to our knowledge, it has not been used in RTS environments like StarCraft. Our paper aims to help fill this gap.

## METHODOLOGY

We conducted a pair think-aloud study, where participants worked to understand and explain the behavior of an intelligent agent playing StarCraft II, a real-time strategy (RTS) game. We brought participants in as pairs to keep them talking, by leveraging the social convention of discussing a shared task.

StarCraft II is a popular RTS game [30] that has been used for AI research [41]. We chose a match from a top level tournament, featuring professional players<sup>1</sup>. The replay we chose to analyze was a representative sample in terms of game flow, e.g., initially building up economy, some scouting, transitioning to increasing combat [30].

Because we were interested in how participants would go about understanding an intelligent agent's behaviors, we hid the players' names and told participants that one of the players was under AI control — even though that was untrue.

Instead of using deception to simulate an intelligent agent with a human player, an alternative design might use a replay of a game with an intelligent agent playing. However, we needed replay files with both interactive replay instrumentation and high-quality game-play. We were unable to locate an intelligent agent in the RTS domain with high enough quality for our investigation, i.e., without limitations like exploiting “a strategy only successful against AI bots but not humans” [16].

<sup>1</sup>We used game 3 of the match between professional players *ByuL* (displayed as “CPU1”) and *Stats* (displayed as “Human”) during the *IEM Season XI - Gyeonggi* tournament. The IEM tournament series is denoted as a “Premier Tournament,” by TeamLiquid, a multi-regional eSports organization that takes a keen interest in professional StarCraft II. The replay file is public at: <http://lotv.spawningtool.com/23979/>

Participant: Age, Gender, Major	SC2 Hours		RTS Hours			
	Casual	Comp.	Casual	Comp.		
Pair1-P1	41 M	EE <sup>1</sup>	200	100	500	300
Pair1-P2	20 M	ECE <sup>2</sup>	50	20	30	30
Pair2-P3	23 M	CE <sup>3</sup>	10	5	25	55
Pair2-P4	23 M	ME <sup>4</sup>	100	200	50	0
Pair3-P5	21 M	EE	50	0	150	12
Pair3-P6	27 M	CE	15	2	150	10
Pair4-P7	23 M	CE	40	20	20	30
Pair4-P8	28 F	EnvE <sup>5</sup>	200	100	300	30
Pair5-P9	21 M	BE <sup>6</sup>	40	40	100	0
Pair5-P10	19 M	ECE	700	300	50	0
Pair6-P11	22 M	BE	100	2	160	100
Pair6-P12	22 F	EnvE	0	70	0	0
Pair7-P13	22 M	CE	15	60	100	50
Pair7-P14	20 M	BE	35	3	40	0
Pair8-P15	23 M	CE	10	0	100	5
Pair8-P16	22 M	BE <sup>7</sup>	16	1	15	0
Pair9-P17	21 M	PS <sup>8</sup>	90	5	500	80
Pair9-P18	19 M	ME	100	0	20	0
Pair10-P19	24 F	FA <sup>9</sup>	5	5	0	0
Pair10-P20	23 M	EdEn <sup>10</sup>	80	15	50	0

**Table 1. Participant demographics and their casual vs. competitive (Comp.) experience. SC2 is StarCraft II, and RTS is any other Real-Time Strategy game.**

<sup>1</sup>Electrical Engr. <sup>2</sup>Electrical & Computer Engr. <sup>3</sup>Chemical Engr. <sup>4</sup>Mechanical Engr. <sup>5</sup>Environmental Engr. <sup>6</sup>Biological Engr. <sup>7</sup>Business Entrepreneurship <sup>8</sup>Political Science <sup>9</sup>Fine Arts <sup>10</sup>Education & English

However, for us to obtain data on how people assess AI agents, our participants needed only to *believe* that a player was an AI. Fortunately, participants were generally convinced of the deception. For example, Pair5-P10 speculated about implementation, “*he must have been programmed to spam*”, and Pair10-P20 decided “*...the AI <is> trying to act like a human.*”

## Participants

We wanted participants familiar with the StarCraft user interface and basic game elements, but without knowledge of machine learning or AI concepts, so we recruited StarCraft players at Oregon State University with at least 10 hours of experience — but excluding computer science students. Also, to avoid language difficulties interfering with the think-aloud data, we accepted only participants with English as their primary language. As per these criteria, 20 undergraduate students participated (3 females and 17 males), with ages ranging from 19–41, whom we paired based on availability. Participants had an average of 93 hours of casual StarCraft experience and 47 hours of competitive StarCraft experience (Table 1).

## Procedures

### Main Task's Procedures

For the main task, each pair of participants interacted with a 16-minute StarCraft II replay while we video-recorded them. The interactive replay instrumentation, shown in Figure 1, allowed participants to actively forage for information within the replay, and we gave them a short tutorial of its capabilities. Examples of ways they could forage in this environment were to move around the game map, move forward or backward in time, find out how many units each player possessed, and drill down into specific buildings or units.

**When a participant paused the replay...**

- What about that point in time made you stop there?
- Did you consider stopping on that object at any other point in time?

**When a participant navigated to a patch...**

- What about that part of the game interface/map made you click there?
- Did you consider clicking anywhere else on the game interface/map?

**When a participant navigated away from a patch (or unpaused)...**

- Did you find what you expected to find?
- What did you learn from that click/pause?
- Did you have a different goal for what to learn next?

**Table 2. Interview questions (drawn from prior IFT research [33]), and the triggers that caused us to ask them.**

Participants watched and foraged together as a pair to try to make sense of the agent’s decisions. One participant controlled the keyboard and mouse for the first half of the task, and the other participant took control for the second half. To help them focus on the decisions, we asked them to write down *key decision points*, defined for them as, “an event which is critically important to the outcome of the game.” Whenever they encountered what they thought was a key decision point, they were instructed to fill out a form with the decision’s time stamp, a note about it, and which player(s) effected it.

*Retrospective Interview’s Procedures*

After the main task, we conducted a multi-stage interview based on the actions the participants took during the main task. To add context to what participants wrote down during the main task, we played parts of our recording of their main task session, pausing along the way to ask why they chose the decision points they did. The wording we used was: “*In what way(s) is this an important decision point in the game?*”

We went through the main task recording again, pausing at their navigations to ask the questions in Table 2. Since there were too many to ask about them all, we sampled pre-determined time intervals to enable covering several instances of each type of navigation (e.g. temporal, spatial, UI) for all participant pairs.

**Analysis Methods**

To answer RQ1, we qualitatively coded instances in the main task where participants asked a question out loud, regardless of whether they were discussing the human or the intelligent agent, using the code set outlined later in Table 3. Our researchers had used this code set on a different corpus [8] in which they independently coded 34% of the corpus and achieved 80% inter-rater reliability (IRR). In this study, the same researchers who achieved this IRR split up the coding of the current data.

To answer RQ2, we qualitatively analyzed the participants’ responses to the retrospective interview questions. To develop a code set to answer the question “*Why was the participant seeking information?*” we started with affinity diagramming to generate groups of answers, which was performed by a group of four researchers. The affinity diagram led to the following codes: *Monitoring State*, *Updating Game State*, *Obsolete Domain*, and *New Event*, as shown later in Table 5. Two researchers then individually qualitatively coded the participants’ responses using this code set on 20% of the data. Given that our IRR on this portion was 80%, one researcher then completed the rest of the coding alone.

Intelligibility Type	Freq
What: What the player did or anything about game state. -Pair3-P5: “So he just killed a scout right?”	148
What-could-happen: What the player could have done or what will happen. -Pair5-P10: “What’s he gonna do in response?”	16
Why-did: Why the player performed an action. -Pair10-P20: “What was the point of that?”	14
How-to: Explaining rules, directives, audience tips, high level strategies. -Pair3-P5: “You have to build a cybernetics core [structure that unlocks new units], right?”	9
*How-good/bad-was-that-action: Evaluation of player actions. -Pair10-P19: “Like, clearly it didn’t work the first time, is it worth it to waste four units the second time?”	8
Why-didn’t: Why the player did not perform an action. -Pair10-P20: “Why aren’t they attacking the base?”	7

**Table 3. Intelligibility type code set, frequency data, and examples. The code set is slightly modified (denoted by the asterisk) from the schema proposed by Lim & Dey: We added How-good/bad-was-that-action because the users wanted an evaluation of agent actions.**

To answer RQ3, we qualitatively coded the decision point forms the participants used during the main task. Here again, we developed a code set using affinity diagramming. The four higher level codes we used were building/producing, scouting, moving, and fighting. We coded 24% of the 228 identified decision points according to this code set, and reached IRR of 80%, at which point one researcher coded the rest of the data.

**RESULTS**

**RQ1: The Prey**

To understand how predators seek prey in the RTS domain, we analyzed questions participants asked during the main task. To situate our investigation in the literature of humans trying to understand AI, we coded the utterances using the Lim & Dey intelligibility types [22] (Table 3).

The results were surprising. Although prior research has reported Why questions to be much in demand [22, 23], only 10% of our participants’ questions fell into the Why did and Why didn’t categories combined (Table 4). Over 70% of our participants’ questions pertained to What.

To get a sense of how representative our participants’ questions were, we turned to the experts — namely, professional explainers in this domain, known as “shoutcasters.” Because we were interested in the very best explainers in this domain, we restricted our search for shoutcaster videos to those that fit the description in the Methodology (top level tournament, professional players). From this pool, we used the same code set as Table 3 to analyze two professionally explained games: Byun vs. Iasonu<sup>2</sup> and Nerchio vs. Elazer<sup>3</sup>.

Consistent with our participants’ questions, the shoutcasters’ explanations were dominated by answers to What questions. In the Nerchio vs. Elazer game, shoutcasters answered What questions 54% of the time, and in Byun vs. Iasonu they

<sup>2</sup>Game 2 of *Byun vs. Iasonu* in the 2016 IEM Gyeonggi tournament, available at: <https://sc2casts.com/cast20681-Byun-vs-Iasonu-B03-in-1-video-2016-IEM-Gyeonggi-Group-Stage>

<sup>3</sup>Game 2 of *Nerchio vs. Elazer* in the 2016 WCS Global Finals tournament, available at: <https://sc2casts.com/cast20439-Nerchio-vs-Elazer-B03-in-1-video-2016-WCS-Global-Finals-Group-Stage>

Question	Total	Pair 1	Pair 2	Pair 3	Pair 4	Pair 5	Pair 6	Pair 7	Pair 8	Pair 9	Pair 10
What	148	2	3	41	1	6	14	10	1	8	62
What-could-happen	16		1	1		3	1	1		2	7
Why-did	14			2			3	1			8
How-to	9		1	3							5
How-good/bad-was-that-action	8			3		3					2
Why-didn't	7		1	1							5
Total	202	4	4	51	1	12	18	12	1	10	89

**Table 4.** Frequency of Lim & Dey questions participants asked each other, by session. Note how often What questions were asked, both by the population of participants as a whole, and by few individual pairs, where it was particularly prevalent.

provided What answers 48% of the time. Since shoutcasters are hired to provide what game audiences want to know, their consistency with our participants' questions suggest that this distribution of questions was typical for the domain.

#### The many flavors of "What" prey

Why such a difference from prior research results? One hypothesis is that, in this kind of situation, participants' prey was simply "play-by-play" information. However, this hypothesis is not well supported by the data. Although participants did seek some play-by-play information (Pair3-P5: "...so he just killed a scout, right?"), several common prey patterns in their What questions went beyond play-by-play. Three of these patterns accounted for about one-third of the What questions.

*The "drill-down What" of current state:* One common question type participants asked when pursuing prey were questions that involved drilling down to find the desired information. Half of the pairs asked drill-down What questions about the game players' unit production or composition. There were 21 instances of this type of What question alone, accounting for almost 15% of the total What questions. For example, the following question required the participants to drill down into several structures on the map to answer it:

Pair3-P6: "Is the human building any new stuff now?"

Navigating in pursuit of this kind of prey was often costly. The least expensive way was navigating via a drop-down menu (2 clicks) in region 2 of Figure 1, but participants instead often foraged in other ways. For example, to find the answer to a question (like Pair3-P6's earlier) participants sometimes navigated to several unit producing structures on the map, into a structure, and then on to the next. For example, Pair 3 made seven navigations to answer their question about "building new stuff."

Shoutcasters' comments closely matched the participants' interest in drilling down: 18% of shoutcasters' What comments answered drill-down questions, compared to the 15% of our participants' drill-down Whats. As an example of the match to shoutcasters' comments, Pair6-P12's drill-down question about unit composition: "I think, well, we have a varied composition, besides roaches [ground military unit], and what are these?" would be well-matched to shoutcaster explanations such as these:

Shoutcaster for Byun v Iasonu: "[the player has] 41 zerglings [cheap military unit] at the moment."

Shoutcaster for Nerchio v Elazer: "And 12 [zerglings] as well, [and] that's just a few more [zerglings] than you normally see."

This suggests using shoutcasters as a possible content model for future explanation systems, given that shoutcasters' "supply" of explanations seem to match well with participants' "demand" for explanations of this type.

*The "temporal What" of past states:* A second common prey pattern, used by almost half (4 of the 10) participant pairs, was asking What questions to fill in gaps regarding past states. These accounted for 15 instances (about 10%) of their Whats.

Pair3-P6: "When did he start building [a] robotics facility [unit producing structure]?"

As in the drill-down pattern, the participants' demand for temporal Whats matched well with the shoutcasters' supply: about 10% of the shoutcasters' What explanations reminded listeners of some past event pertinent to current game state. For example:

Shoutcaster for Byun v Iasonu: "...the +1 carapace [defensive unit upgrade] early upgrade ...[is] paying off."

*The "higher-level What":* The third common prey pattern was at a higher level of abstraction than the specific units or events, aiming instead toward more general understanding of what was going on in the game. These What questions arose 12 times (about 8%) across all instances. For example, Pair10-P20 asked, "What's going on over there?" in which "there" referred to a location on the map with military units that could have been gearing up for combat. The shoutcasters seemed enthusiastic about providing this kind of information<sup>4</sup>, perhaps because it provided opportunities to add nuance and insight to their commentary. For example:

Shoutcaster for Byun v Iasonu: "This is about to get crazy because [of] this drop coming into the main base [and] the banelings [suicide bomber unit] trying to get some connections in the middle."

Shoutcaster for Nerchio v Elazer: "I like Elazer's position; he's bringing in other units in from the back..."

#### Questioning the unexpected

Lim & Dey's reported that when a system behaved in unexpected ways, users' demand to know Why increased [22]. Consistent with this, when our participants saw what they expected to see, they did not ask Why or Why-didn't questions. For example, Pair 4 and Pair 5 did not ask any Why or Why-didn't questions at all. Instead, they made remarks like:

Pair4-P7: "the Zerg is doing what they normally do."

Pair4-P8: "[The agent is] kind of doing the standard things."

Pair5-P10: "This is a standard build."

<sup>4</sup>We did not count the number of shoutcaster comments that answered this question because we could not narrow them down in this way. That is, although many of their comments could be said to be applicable to this type of question, the same comments were also applicable to more specific questions.

However, in cases of the unexpected, a fourth What prey pattern arose, in which participants questioned the phenomena before them. We counted 9 What questions of this type:

- Pair9-P17: "...interesting that it's not even using those."
- Pair10-P19: "I don't get it, is he expanding?"
- Pair10-P19: "Wow, what is happening? This is a weird little dance we're doing."
- Pair10-P20: "<when tracking military units> What the hell was that?"

The unexpected also produced Why questions. About half of the participants' Why and Why-didn't questions came from seeing something they had not expected or not seeing something they had expected. For example:

- Pair1-P1: "<noticing a large group of units sitting in a corner> Why didn't they send the big army they had?"

*Implications for a Future Interactive Explanation System*

Using the Lim & Dey intelligibility types (What, Why, etc.) to categorize the kinds of prey our participants sought produced implications for shoutcasters as possible "gold standards" for informing the design of a future automated explanation system in this domain. For example, the high rate of What questions from participants matched reasonably well with a high rate of What answers from shoutcasters. Drawing explanation system design ideas from these expert explainers may help inform the needed triggers and content of the system's What explanations.

Also, the dominance of What questions point to participants' prioritizing of state information in this domain. Drill-down Whats were about state information they hadn't yet seen, temporal Whats were about past states they either hadn't seen or had forgotten, and higher-level Whats were about understanding the purpose of a current or emerging state. Further, the shoutcasters matched and sometimes exceeded the participants' rate of Whats in each of these categories with their explanations. This suggests that in the RTS domain, an explanation system's most sought-after explanations may be its explanations relating to state.

Also, as noted in prior research, unexpected behaviors (or omissions of expected behaviors) led to increases in questions of both the What and the Why intelligibility types [22]. If an explanation system can recognize unexpected behavior, it could then better predict when users will want Why and What explanations to understand the deviation from typical behavior.


Finally, the cost of navigating to some of the prey at times became expensive, which points to the need for explanation systems to keep an eye on the cost to users of obtaining that information. In this section, this came out in the form of navigation actions. The next section will point to costs to human cognition as well.

**RQ2: The Foraging Paths**

Various cognitive costs were incurred by participants by following paths to find prey. As an information environment, RTS games have foraging characteristics that set them apart from other information environments that have previously been studied from an IFT perspective, such as web sites [34]

Reasons for participants' path choices code set	Freq
Monitoring State: Continuous game state monitoring, such as watching a fight. -Pair4-P7: "I wanted to see how the fight was going."	65
New Event: Attending to a new event for which participant wished to satisfy curiosity about. -Pair2-P4: "I saw there was a new building."	36
Update Game State: Updating potentially stale game information that the participant explicitly stated prior knowledge about. -Pair1-P2: "I was mainly looking at the army composition, seeing how it had changed from the last fight."	29
Obsolete Domain: Explicitly using domain info that may not be current, such as game rules (e.g., what buildings can produce). -Pair3-P6: "I mainly clicked on the adept because I'm more familiar with [a previous version of the game]."	11

**Table 5. Reasons for participants' path choices code set, with examples and frequency data, to answer the question "Why was the participant seeking that information?"**



**Figure 2. Conceptual drawing of foraging in the RTS domain vs. previously studied foraging. (Left): Information environments in prior IFT literature: the predator considers few paths, but the paths are sometimes very deep. This figure inspired by an IDE foraging situation in [[33] Fig. 5]. (Right): Foraging in RTS, where most navigation paths are shallow, but with numerous paths to choose from at the top level.**

and programming IDEs [33]. These previously studied domains are relatively static, with most changes occurring over longer periods. In contrast, an RTS information environment changes rapidly and continually, driven by actions that do not originate from the foragers themselves. As we will see, this caused participants to spend some time monitoring the overall game state, waiting for a suitable cue to appear for them to investigate further.

The number of paths a forager might follow in an RTS information environment increases with the complexity of the game state, but path lengths tend to be short. This is conceptualized in Figure 2. This means that most questions are answered within a few navigations. However, in foraging environments like IDEs, there might only be a few interesting links from any one information patch, but some can lead to lengthy sequences of navigations (e.g., the "Endless Paths" problem [33]).

*Foraging in the RTS domain*

Interestingly, there was hardly any difference between RTS foraging and other environments at first. During the early stages of a game, there are very few units, buildings, or explored regions for users to navigate to, so foraging is relatively straightforward. As one participant put it:

- Pair7-P14: "There is only so many places to click on at this point."

As long as this remained the case, each relevant path could potentially be carefully pursued, similarly to an IDE. Four participant pairs (2, 4, 7, 9) paused the replay for an average

of 90 seconds within the first 1:30. They studied individual objects and actions with a great deal of scrutiny, which was surprising considering the sparse environment. In contrast, later on in the game, when 50 of the same unit existed, they received much less attention than when there was just one.

Choosing among many available paths created cognitive challenges for participants. Participants needed to keep track of an increasing amount of information as the match progressed. Each time a player performed an action, which added information, the participants could forage for this new information. If the participant did so, we coded their navigation as a *New Event*, which accounted for 26% of our interviewed navigations (Table 5). For example:

Pair10-P19: “...noticed movement in the Minimap, and that the Zerg troops were mobilizing in some fashion. So I guess I just preemptively clicked...”

The rate of path creation exacerbated the “many paths” problem. Professional StarCraft players regularly exceed several hundred actions per minute (APM) [43]. This meant that players performed rapid actions that changed the game state. Each of these actions not only potentially created new paths; they potentially updated the existing ones. This caused the knowledge the participants had about paths that had not been recently checked to become stale, which in turn led to a strong prevalence of two behaviors. *Update Game State* was very common in our data set, indicating that participants often needed to check on paths that may have been updated (21% of interviewed navigations, Table 5).

Pair8-P15: “...there’s a big force again. Just checking it out to see if anything has progressed from earlier.”

Pair1-P2: “I was mainly just looking at the army composition, seeing how it had changed from the last fight, see if they had made any serious changes ...”

Note that this is slightly different from our second behavior, *Monitoring State*, which is like updating game state, but with a nonspecific goal. *Monitoring State* was the most common reason for interviewed navigations (46% of navigations were for the purpose of monitoring, Table 5), for example:

Pair5-P9: “I noticed like the large mass of units on the map and I wanted to know what the player was doing with them.”

Pair8-P16: “I was just kinda checking on things. Sort of due diligence keeping an eye on the different happenings that the AI was doing at the time.”

Since each event and its corresponding cues were only visible for a limited time, paths not chosen right away by our participants quickly disappeared. Further, paths are numerous, and frequently updated. Thus, there is a large risk for paths of inquiry to be forgotten or going unnoticed as the game proceeds, as in these examples:

Pair7-P14: “Oh my gosh, I didn’t even notice he was making an ultralisk den.”

Pair3-P6: “I didn’t notice they canceled the assimilator [mining structure]”

#### *Many Rapidly Updating Paths: Coping Mechanisms*

Our participants responded to this issue in several ways. First, some participants chose a path and stuck to it, ignoring the others. Note that this required paying an information cost, because contextual information that may have been very important for future decisions could be discarded in the process. This strategy was exclusively followed by 3 pairs (2,7,8), who made barely any temporal navigations during the study, as described in Table 6. These participants analyzed the replay using not much more time than shoutcasters spend. However, achieving this speed of analysis required participants to *ignore* many game events.

For example, when asked about desire to click anywhere else, one participant volunteered:

Pair10-P19: “Mmm, if I had multiple, like, different screens yeah. But no, that seemed to be where the action was gonna be.”

In this fashion, participants chose to triage game events based on some priority order. In both of the following examples, the participants navigated away from the conclusion of a fight:

Pair6-P11: “I wanted to check on his production that one time, because he just lost most of his army, and he still had some [enemies] to deal with.”

Pair3-P5: “I was trying to see what units they were building, after the fight, see if they were replenishing, or getting ready for another fight.”

The second method our participants used to manage the complexity of paths was to use the time controls to slow down, stop, or rewind the replay. Although pausing to assess the state was fairly common in all groups, rewind behavior yielded more information. Pairs 3 and 10 rewound the most often (Table 6), and paid higher *navigation* costs to do so, but they viewed these navigations as worthwhile to providing necessary information:

Pair6-P11: “I looped back to the beginning of the final fight ... to see if there was anything significant that we had missed the first time around.”

However, the cost of doing so was more than just time, because the more paths they monitored, the greater the cognitive load:

Pair10-P19: “There’s just so much happening all at once; I can’t keep track of all of it!”

#### *Implications for a Future Interactive Explanation System*

Assessing an agent required considering a great *many* paths, and choosing one (or few), though most paths followed were not particularly *long*. Note that this contrasts with previous literature in software engineering, which is characterized by “miles of methods [33],” such as a long sequence of methods in the stack trace. Thus, rapid evaluation and pruning of paths is critical in the RTS domain, but less so in software engineering, where the options to consider are fewer and time pressure is lighter. One solution could be a recommender system to help the user triage which path to follow next.

During assessment, participants’ often forgot about or otherwise interrupted their paths of inquiry. For example, if a new important path appeared, such as a critical battle, either that

	Task Time	Real-Time Ratio	Rewinds	Timestamp Rewinds	Context Notes
Pair 1	20:48	1.3	3	1	Rewatched 1 fight
Pair 2	20:40	1.3			Extensive pause around 1:00 to evaluate game state
Pair 3	55:08	3.4	12	9	Rewatched fights and fight setup. Slowed down replay during 1 combat.
Pair 4	32:16	2.0	2		Rewatched opening build sequence and evaluated information available to the agent at a key moment. Many pauses to explain game state.
Pair 5	24:23	1.5	5		Rewatched unit positioning, AI reaction to events, and scouting effectiveness.
Pair 6	31:56	2.0	4	2	Rewatched 2 fights.
Pair 7	29:49	1.9			Made no use of time controls other than pausing to write down decision points.
Pair 8	21:27	1.3			Made no use of time controls other than pausing to write down decision points.
Pair 9	39:17	2.4	2	1	Rewatched 1 fight. Slowed down replay for the entire task.
Pair 10	61:14	3.8	Lots	Some	Rewound extensively, in a nested fashion. Changed replay speed many times.

**Table 6.** Participant task time (33:42±14:18 minutes) and time control usage information. Note that the replay file was just over 16:04, so dividing each pair’s time by 16 yields the third column, “Real-Time Ratio” (2.1±0.89). Some of times participants rewound the replay were because we requested timestamps for events, shown in the fourth column, “Timestamp Rewinds.” The last column provides any additional context in which replay and pause controls were used.

path or the current path had to be dropped. In another domain (spreadsheet debugging), participants faced with branching paths with multiple desirable directions became more effective when the environment supported a strategy they call “to-do listing” [12]. Because to-do listing was supported on its own or in composition of other problem-solving approaches, it could also act as a strategy enhancer. Perhaps in the RTS domain, a similar strategy could enable users to carry on with their current path uninterrupted — but also keep track of the critical battle to come back to later.

**RQ3: The Decisions and the Cues**

When participants were not heading down the “right” path, what cues did they instead follow toward some other path? Also, what did they consider the “right” cues to follow?

In the RTS domain, players and intelligent agents make thousands of sequential decisions, and there is a paucity of literature that considers humans trying to understand AI decisions in such a setting. (A notable exception is McGregor et al. [25].) There is, however, literature that starts with the AI’s perspective: instances of its decision-making system components (i.e., neurons) that are interpretable by humans [44, 45]. In contrast, here we wanted to start with the human’s perspective and the foraging paths that result from it: namely, how participants would identify behaviors that were not only potentially human interpretable, but also of interest.


Thus, we asked participants to write down what they thought were the important game events. We defined the term *key decision points* to our participants as “an event which is critically important to the outcome of the game,” to give participants leeway to apply their own meaning. Since all participant pairs were examining the same replay file, we were then able to compare the decision points the different participants selected. That is, the cues in the information environment were the same for all the participants — whether they noticed them or not.

Key decision points fell into four main categories: *building/producing*, *fighting*, *moving*, and *scouting*. The participants were in emphatic agreement about the most important types of decision points to pursue. Of the 228 total decision points participants identified, *Fighting* and *Building* made up 85% (Table 7).

In fact, participants showed remarkable consistency about the importance of the *Expansion* subcategory of *Building*.

Code	Total	Pair 1	Pair 2	Pair 3	Pair 4	Pair 5	Pair 6	Pair 7	Pair 8	Pair 9	Pair 10
Expansion	52	7	7	8	8	-	6	3	-	7	6
Building - Rest	69	6	4	7	15	1	7	11	2	12	4
Building - All	114	13	11	15	21	1	12	12	2	18	9
Fighting - All	98	8	4	11	8	4	10	6	8	11	28
Moving - All	26	3	1	5	1	2	2	2	2	3	5
Scouting - All	23	1	2	1	5	1	1	3	-	3	6
Total	228	34	20	40	45	11	31	39	16	42	65

**Table 7.** Summary of decision points identified by our participants. Sums may exceed totals, since each decision point could have multiple labels. Note how prevalent *Expansion* was within the *Building* category.



**Figure 3.** Dots show the *Building-Expansion* decision points each participant pair (y-axis) identified over time (x-axis). Red lines show when *Expansion* events actually occurred. (Participants noticed most of them.) The red box shows where Pair 4 failed to notice an event they likely wanted to note, based on their previous and subsequent behavior.

Eight of the ten participant pairs identified *Expansion* decision points, when a player chooses to build a new resource-producing base (Table 7). Extra resources from expanding allowed a player to gain an economic advantage over their opponent because they could build more units:


Pair1-P2: “Of course, if you have a stronger economy you will likely win in the end.”

Moreover, because those that identified any *Expansion* found at least three, *Expansions* seemed to be considered important throughout the duration of the game.

Pair6-P11: “... the third base is important for the same reason the first one was, because it was just more production and map presence.”

Even so, they missed some of the cues pointing out expansion decisions. The event logs in the replay file reveal that new bases were constructed at roughly {1:00, 1:30, 2:00, 5:00,





**Figure 4.** (Top:) The *Scouting* decision points identified by our participant pairs (y-axis), with game time on the x-axis. (Bottom:) The *Fighting* decision points identified, plotted on the same axes. After *Fighting* events begin (red line), *Scouting* decision points are no longer noticed often — despite important *Scouting* actions continuing to occur.

6:30, 11:20, 12:00, and 13:45}, each of which is marked with a red line on Figure 3. Only Pair 3 identified decision points for all 8 of these, and 7 pairs omitted at least one<sup>5</sup>, with one example highlighted with a red box in Figure 3.

Since *Expansion* decisions were so important to our participants, why did they miss some? Some cues in the information environment led participants *away* from desirable prey<sup>6</sup>. We term such cues “distractor cues.” Participants were so distracted by cues that provided an alluring scent, albeit to low-value information, they did not notice the other cues that pointed toward the “*Expansion*” decisions.

Distractor cues led participants astray from *Expansion* in nine cases, and eight of them involved units in combat or potentially entering combat. (The ninth involved being distracted by a scouting unit.) For example, Pair 7 missed the expansion at the 13:45 minute mark, instead choosing to track various groups of army units, which turned out to be unimportant to them:

Pair7-P14: “*These zerglings are still just chilling.*”

Interestingly, participants had trouble with distractor cues even when the number of events competing for their attention was very low. For example, in the early stages of the game, players were focused on building economies and scouting. There was little to no fighting yet, so it was not the source of distracting cues. We were not surprised that the *Expansion* event at 13:45, when the game state had hundreds of objects and events, was the most often missed (5 instances). However, we were surprised that even when the game state was fairly simple — such as at 1:30 where the game had only 13 objects — participants missed the *Expansion* events. The extent of

<sup>5</sup> Table 7 shows Pair 4 also finding eight *Expansion* decision points, but one of those is about the *commitment* to expand, based on building other structures to protect the base, rather than the action of building the base itself.

<sup>6</sup>Reminder: Cues are signposts in the *environment* that the predator observes, such as rabbit tracks. Scent is what predators make of cues in their *heads*, such as thinking that rabbit tracks will lead to rabbits.

distractibility the participants showed even when so little was going on was beyond what we expected.

So if decision points went unnoticed in simple game states, what did they notice in complex ones? *Fighting*. All participants agreed *Fighting* was key, identifying at least one decision point of that type (Table 7). The ubiquity of *Fighting* codes is consistent with Kim et al. [16], who found that combat ratings were the most important to the participant’s perception score. *Fighting* provided such a strong scent that it was able to mask most other sources of scent, even those which participants prioritized very highly.

*Scouting* offers an example of *Fighting* leading participants away from other important patches. *Scouting* decision points occurred in the first half of the game, but died out once *Fighting* decision points started to occur in the second half of the game. As Figure 4 shows, the start of *Fighting* decision points coincides with the time that *Scouting* decision points vanish — despite the fact that scouting occurred throughout the game, and that participants believed scouting information mattered:

Pair4-P8: “*But it’s important just to know what they’re up to and good scouting is critical to know who you are going to fight.*”

#### Implications for a Future Interactive Explanation System

Participants had a tendency toward following cues that were interesting or eye-catching, at the expense of those that were important but more mundane. In this domain, the “eye-catching” cues were combat-oriented, whereas the “mundane” cues were scouting oriented. Other domains may have similar phenomena, wherein certain aspects of the agent’s behaviors distract from other important views due to triggering an emotional response in the viewer. Thus, supporting users’ attending to actions that are important but mundane is a design challenge for future interactive explanation systems.

#### THREATS TO VALIDITY

Every study has threats to validity [42]. This paper presents the first study of information foraging either in the area of explainable AI or in the domain of RTS games, so its results cannot yet be compared or validated by other studies by other researchers. Thus, we must be conscious of its limitations.

Aspects of our study may have influenced our participants to ask less questions in general, such as not asking a question of their partner if they did not expect their partner to be able to answer it. Also, participants took different amounts of time to do the task, ranging from 20 minutes to an hour. Thus, certain participant pairs talked more than others in the main task, creating a form of sampling bias. Threats like these can be addressed only by additional empirical studies across a spectrum of study designs, types of intelligent interfaces, and intelligent agents.

#### DISCUSSION: WHAT IFT CAN OFFER EXPLAINABLE AI

The Implications sections along the way have discussed detailed implications of IFT on XAI explanation design. Here, we turn to how IFT allows us to “connect the dots” between our work and other work done from an IFT perspective.

IFT allows abstracting above game objects like “assimilators” to the IFT constructs of *prey* (RQ1), foraging *paths* (RQ2), and following *cues* (RQ3). The IFT lens revealed that participants faced difficult foraging problems – some of which are new to IFT research – and faced high foraging costs. For example, failure to follow the “right” paths resulted in a high *information cost* being paid, but finding a reasonable path needed to be done quickly due to the ever-changing game environment (at a high *cognitive cost*). Although the user could relax the real-time pressure by pausing the replay, excessive rewinding incurred not only a high *navigation cost* for rewind-positioning and pausing, but also an additional *cognitive cost* of remembering more context.

Participants had to make trade offs between two types of these costs, navigational and informational, and their triaging to manage such trade-offs led to even more cognitive cost. Each path participants followed incurred a navigational cost, so following more paths led to higher cumulative costs. However, reducing the number of paths they followed incurred the information cost of missing out on potentially important information. Worse, the information cost paid by adhering to a single path compounded over time. For example, if participants made a bad path choice early in the game, and repeated that mistake throughout the game, then later in the game they may be confused by an event that they did not expect — due to lacking appropriate context. One cause for making a bad choice was “distractor cues,” where, in order to curtail their current navigation direction and move to another, participants paid the high information cost of missing information important to them, often unwittingly.

The IFT perspective also connects some of the problems our participants faced to known problems of foraging in other domains. One open problem in IFT is the Prey in Pieces problem<sup>7</sup> [33]. Our participants encountered this problem because they had to piece together bits of evidence of the agent’s decisions in order to assess the agent. In doing so, participants were sometimes uncertain about what each of these decisions meant about the competencies and strategies of the agent. When aggregating multiple sources of uncertain data like this, prior research has shown that computational assistance can increase user confidence, although manual comparison is still preferred in high-stakes situations [11]. This seems to suggest that a recommender may be helpful to help users select and aggregate agent actions for explanation, though manual comparison may still be necessary at times.

Another open problem in IFT is the Scaling Up problem [33]. This problem was revealed in the domain of IDEs, in which foragers (developers) had great difficulty accurately predicting the cost and value of going to patches more than one link away. The problem that the developers faced was a *depth* problem (recall Figure 2). In contrast, in our domain, participants faced a *breadth* Scaling Up foraging problem: constantly having to choose which of many paths to follow. The Scaling Up

problem as a depth problem is still open; so too is the breadth version of it identified here.

In both cases, users foraging for information want to maximize value per cost. In the IDE case, this is accomplished by pruning low value paths unrelated to the bug. For example, if a developer is fixing a UI bug, they can potentially ignore database code. However, in RTS, any action could be important, so many paths need to stay on the table. Further, the rapid rate of change in the environment limits the user’s planning depth, which decreases accuracy of predictions of cost/value. Thus, the Scaling Up problem is different in the RTS domain. In depth domains like IDEs, the problem is predicting cost/value in far-distant patches, whereas in the RTS domain, the difficulty is rapidly choosing at the top level among the many, many available paths.

## CONCLUSION

In this paper, we presented the first theory-based investigation into how people forage for information about an intelligent agent in an RTS environment and the implications for Explainable AI. Our results suggest that people’s information seeking in this domain is far from straightforward. We saw evidence of this from multiple perspectives:

**RQ1 *The Prey*:** Participants favored *What* information over the *Whys* reported by most previous research, and their *Whats* were nuanced, complex, and sometimes expensive.

**RQ2 *The Paths*:** The dynamically changing RTS environment and the breadth-oriented structure of its information paths caused unique information foraging problems in deciding which paths to traverse. These problems led not only to navigation costs, but also to information and cognitive costs.

**RQ3 *The Decisions and the Cues*:** These costs rendered it infeasible for participants to investigate all of the decision points they wanted. This problem was exacerbated by “distractor cues,” which drew participants’ attention elsewhere with interesting-looking cues (like signs of fighting), at the expense of information that was often important to participants (like scouting or expansion).

Perhaps most importantly, our results point to the benefits of investigating humans’ understanding of intelligent agents through the lens of Information Foraging Theory. For example, the IFT lens enabled us to abstract beyond StarCraft, to reveal phenomena – such as the frequent need to trade off cognitive foraging costs against navigation foraging costs against information costs – that are widely relevant to the RTS domain. As we have noted in the “Implications” sections along the way, these theory-based results reveal opportunities for future Explainable AI systems to enable domain experts to find the information they need to understand, assess, and ultimately decide how much to trust their intelligent agents.

## ACKNOWLEDGMENTS

This work was supported by DARPA #N66001-17-2-4030 and NSF #1314384. Any opinions, findings and conclusions or recommendations expressed are those of the authors and do not necessarily reflect the views of NSF, DARPA, the Army Research Office, or the US government.

<sup>7</sup>“Prey in Pieces” is analogous to buying a coffeemaker one piece at a time [33]. The cost of going to every store must be paid plus the cost of piecing things together at the end, rather than the cost of going to one store that has a preassembled coffeemaker.

## REFERENCES

1. Juan Felipe Beltran, Ziqi Huang, Azza Abouzied, and Arnab Nandi. 2017. Don't just swipe left, tell me why: Enhancing gesture-based feedback with reason bins. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*. ACM, 469–480.
2. Sourav S Bhowmick, Aixin Sun, and Ba Quan Truong. 2013. Why Not, WINE?: Towards answering why-not questions in social image search. In *Proceedings of the 21st ACM International Conference on Multimedia*. ACM, 917–926.
3. Svetlin Bostandjiev, John O'Donovan, and Tobias Höllerer. 2012. TasteWeights: a visual interactive hybrid recommender system. In *Proceedings of the Sixth ACM Conference on Recommender Systems*. ACM, 35–42.
4. Nico Castelli, Corinna Ogonowski, Timo Jakobi, Martin Stein, Gunnar Stevens, and Volker Wulf. 2017. What Happened in my home? An end-user development approach for smart home data visualization. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 853–866.
5. Gifford Cheung and Jeff Huang. 2011. Starcraft from the Stands: Understanding the Game Spectator. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 763–772. DOI: <http://dx.doi.org/10.1145/1978942.1979053>
6. Ed H Chi, Peter Pirolli, Kim Chen, and James Pitkow. 2001. Using information scent to model user information needs and actions and the web. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 490–497.
7. Kelley Cotter, Janghee Cho, and Emilee Rader. 2017. Explaining the News Feed Algorithm: An Analysis of the “News Feed FYI” Blog. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 1553–1560.
8. Jonathan Dodge, Sean Penney, Claudia Hilderbrand, Andrew Anderson, and Margaret Burnett. 2018. How the experts do it: Assessing and explaining agent behaviors in real-time strategy games. In *ACM Conference on Human Factors in Computing Systems*. ACM.
9. Scott D. Fleming, Chris Scaffidi, David Piorkowski, Margaret Burnett, Rachel Bellamy, Joseph Lawrance, and Irwin Kwan. 2013. An information foraging theory perspective on tools for debugging, refactoring, and reuse tasks. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 22, 2 (2013), 14.
10. Wai-Tat Fu and Peter Pirolli. 2007. SNIF-ACT: A cognitive model of user navigation on the world wide web. *Human-Computer Interaction* 22, 4 (2007), 355–412.
11. Miriam Greis, Emre Avci, Albrecht Schmidt, and Tonja Machulla. 2017. Increasing users' confidence in uncertain data by aggregating data from multiple sources. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 828–840. DOI: <http://dx.doi.org/10.1145/3025453.3025998>
12. Valentina I Grigoreanu, Margaret M Burnett, and George G Robertson. 2010. A strategy-centric approach to the design of end-user debugging tools. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 713–722.
13. Bradley Hayes and Julie A Shah. 2017. Improving robot controller transparency through autonomous policy explanation. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 303–312.
14. Zhian He and Eric Lo. 2014. Answering why-not questions on top-k queries. *IEEE Transactions on Knowledge and Data Engineering* 26, 6 (2014), 1300–1315.
15. Ashish Kapoor, Bongshin Lee, Desney Tan, and Eric Horvitz. 2010. Interactive optimization for steering machine classification. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1343–1352.
16. Man-Je Kim, Kyung-Joong Kim, SeungJun Kim, and Anind K Dey. 2016. Evaluation of StarCraft Artificial Intelligence Competition Bots by Experienced Human Players. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 1915–1921.
17. Todd Kulesza, Margaret Burnett, Weng-Keen Wong, and Simone Stumpf. 2015. Principles of explanatory debugging to personalize interactive machine learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces*. ACM, 126–137.
18. Todd Kulesza, Simone Stumpf, Margaret Burnett, and Irwin Kwan. 2012. Tell me more? The effects of mental model soundness on personalizing an intelligent agent. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1–10.
19. Todd Kulesza, Simone Stumpf, Margaret Burnett, Weng-Keen Wong, Yann Riche, Travis Moore, Ian Oberst, Amber Shinsel, and Kevin McIntosh. 2010. Explanatory debugging: Supporting end-user debugging of machine-learned programs. In *Visual Languages and Human-Centric Computing (VL/HCC), 2010 IEEE Symposium on*. IEEE, 41–48.
20. Todd Kulesza, Simone Stumpf, Weng-Keen Wong, Margaret M Burnett, Stephen Perona, Andrew Ko, and Ian Oberst. 2011. Why-oriented end-user debugging of naive Bayes text classification. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 1, 1 (2011), 2.
21. Sandeep Kaur Kuttal, Anita Sarma, and Gregg Rothermel. 2013. Predator behavior in the wild web world of bugs: An information foraging theory perspective. In *Visual Languages and Human-Centric Computing (VL/HCC), 2013 IEEE Symposium on*. IEEE, 59–66.

22. Brian Y Lim and Anind K Dey. 2009. Assessing demand for intelligibility in context-aware applications. In *Proceedings of the 11th International Conference on Ubiquitous Computing*. ACM, 195–204.
23. Brian Y. Lim, Anind K. Dey, and Daniel Avrahami. 2009. Why and why not explanations improve the intelligibility of context-aware intelligent systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2119–2128.
24. M. Lomas, R. Chevalier, E. V. Cross, R. C. Garrett, J. Hoare, and M. Kopack. 2012. Explaining robot actions. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 187–188. DOI: <http://dx.doi.org/10.1145/2157689.2157748>
25. S. McGregor, H. Buckingham, T. G. Dietterich, R. Houtman, C. Montgomery, and R. Metoyer. 2015. Facilitating testing and debugging of Markov Decision Processes with interactive visualization. In *2015 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 53–61. DOI: <http://dx.doi.org/10.1109/VLHCC.2015.7357198>
26. Ronald Metoyer, Simone Stumpf, Christoph Neumann, Jonathan Dodge, Jill Cao, and Aaron Schnabel. 2010. Explaining how to play real-time strategy games. *Knowledge-Based Systems* 23, 4 (2010), 295–301.
27. Tim Miller. 2017. Explanation in Artificial Intelligence: Insights from the Social Sciences. *CoRR* abs/1706.07269 (2017). <http://arxiv.org/abs/1706.07269>
28. Nan Niu, Anas Mahmoud, Zhangji Chen, and Gary Bradshaw. 2013. Departures from optimality: Understanding human analyst’s information foraging in assisted requirements tracing. In *Proceedings of the 2013 International Conference on Software Engineering*. IEEE Press, 572–581.
29. Donald A Norman. 1983. Some observations on mental models. *Mental Models* 7, 112 (1983), 7–14.
30. S. Ontañón, G. Synnaeve, A. Uriarte, F. Richoux, D. Churchill, and M. Preuss. 2013. A Survey of Real-Time Strategy Game AI Research and Competition in StarCraft. *IEEE Transactions on Computational Intelligence and AI in Games* 5, 4 (Dec 2013), 293–311. DOI: <http://dx.doi.org/10.1109/TCAIG.2013.2286295>
31. Alexandre Perez and Rui Abreu. 2014. A diagnosis-based approach to software comprehension. In *Proceedings of the 22nd International Conference on Program Comprehension*. ACM, 37–47.
32. David Piorkowski, Scott D. Fleming, Christopher Scaffidi, Margaret Burnett, Irwin Kwan, Austin Z Henley, Jamie Macbeth, Charles Hill, and Amber Horvath. 2015. To fix or to learn? How production bias affects developers’ information foraging during debugging. In *Software Maintenance and Evolution (ICSME), 2015 IEEE International Conference on*. IEEE, 11–20.
33. David Piorkowski, Austin Z Henley, Tahmid Nabi, Scott D Fleming, Christopher Scaffidi, and Margaret Burnett. 2016. Foraging and navigations, fundamentally: developers’ predictions of value and cost. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. ACM, 97–108.
34. Peter Pirolli. 2007. *Information Foraging Theory: Adaptive Interaction with Information*. Oxford University Press.
35. Stephanie Rosenthal, Sai P. Selvaraj, and Manuela Veloso. 2016. Verbalization: Narration of autonomous robot experience. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI’16)*. AAAI Press, 862–868. <http://dl.acm.org/citation.cfm?id=3060621.3060741>
36. Sruti Srinivasa Ragavan, Sandeep Kaur Kuttal, Charles Hill, Anita Sarma, David Piorkowski, and Margaret Burnett. 2016. Foraging among an overabundance of similar variants. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 3509–3521.
37. Simone Stumpf, Vidya Rajaram, Lida Li, Margaret Burnett, Thomas Dietterich, Erin Sullivan, Russell Drummond, and Jonathan Herlocker. 2007. Toward harnessing user feedback for machine learning. In *Proceedings of the 12th International Conference on Intelligent user interfaces*. ACM, 82–91.
38. Katia Sycara, Christian Lebiere, Yulong Pei, Donald Morrison, and Michael Lewis. 2015. Abstraction of analytical models from cognitive models of human control of robotic swarms. In *International Conference on Cognitive Modeling*. University of Pittsburgh.
39. Joe Tullio, Anind K Dey, Jason Chalecki, and James Fogarty. 2007. How it works: A field study of non-technical users interacting with an intelligent system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 31–40.
40. Jo Vermeulen, Geert Vanderhulst, Kris Luyten, and Karin Coninx. 2010. PervasiveCrystal: Asking and answering why and why not questions about pervasive computing applications. In *Intelligent Environments (IE), 2010 Sixth International Conference on*. IEEE, 271–276.
41. Oriol Vinyals et al. 2017. StarCraft II: A New Challenge for Reinforcement Learning. Tech Report. (2017). Retrieved December 22, 2017 from <https://deepmind.com/documents/110/sc21e.pdf>.
42. Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, Björn Regnell, and Anders Wesslén. 2000. *Experimentation in Software Engineering: An Introduction*. Kluwer Academic Publishers, Norwell, MA, USA.
43. Kevin Wong. 2016. StarCraft 2 and the quest for the highest APM. (Jul 2016). <https://www.engadget.com/2014/10/24/starcraft-2-and-the-quest-for-the-highest-apm/>

44. Tom Zahavy, Nir Ben Zrihem, and Shie Mannor. 2016. Graying the black box: Understanding DQNs. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48 (ICML'16)*. JMLR.org, 1899–1908.  
<http://dl.acm.org/citation.cfm?id=3045390.3045591>
45. Matthew D. Zeiler and Rob Fergus. 2014. *Visualizing and understanding convolutional networks*. Springer International Publishing, Cham, 818–833. DOI:  
[http://dx.doi.org/10.1007/978-3-319-10590-1\\_53](http://dx.doi.org/10.1007/978-3-319-10590-1_53)