# Recursive thinking in adversarial decision making: A naturalistic study

Shuli Yu, Angela L.S. Tan, Andrea J. Chen, Kenny T.W. Choo, Boon Kee Soh

POC: Boon Kee Soh
DSO National Laboratories
27 Medical Drive #11-01 Singapore 117510
sboonkee@dso.org.sg

Being able to accurately predict and understand the actions of an adversary is integral to successful decision making in military operations. Often, this involves creating mental models of the adversary's thought processes and knowledge states. Such mental models can vary in complexity because thinking extensively about the adversary is a cognitively demanding task, especially in high-pressure situations. Although current research on recursive modeling algorithms and simulation aims to aid human decision makers via decision support systems, there is a general lack of descriptive research on how people make decisions in adversarial contexts under naturalistic settings. We examined the role of recursive thinking in adversarial decision making by interviewing junior and senior commanders as they played a real-time strategy game called Rise of Nations over multiple trials. This wargaming microworld provided an effective platform for participants to think freely and strategize about adversarial intent at different orders of recursion. Such a platform enabled us to qualitatively and quantitatively explore the extent to which people naturally thought about their adversary through recursive thinking under a situation with uncertainty, time-stress and limited resources. Our results suggest that participants generally think about the enemy in the first order most of the time, and think about what the enemy thinks of the participants (second order) only occasionally. Although the first two orders of recursion were intuitive to all participants, few participants reached the third order of recursion. We discuss the relation of recursive thinking to prediction accuracy.

**Keywords**: recursive thinking, naturalistic decision making, adversarial decision making

## 1. INTRODUCTION

Understanding adversarial intent involves comprehending the actions of an opponent by generating a mental model of the adversary's thinking processes and knowledge states. Being able to accurately assess the opponent is crucial in any conflict situation because it enables decision makers to solve adversarial problems effectively by anticipating, understanding and counteracting the actions of the opponent (Thagard, 1992).

In warfare and other adversarial situations, decision makers often need to make decisions in a dynamic and uncertain environment, against intelligent opponents who can also respond and make decisions of their own. This involves understanding the interactions of various factors (e.g. weather, terrain, troops, weapons) in the environment and their dynamic relationships with time. Concurrently, decision makers have to contemplate with unknowns or inaccurate information since they may not have all the necessary information required. These aspects heavily influence how the adversary deploys resources, decides on actions and stages deception. Hence, appreciating the motivations and intentions of an adversary in such a complex decision space is a challenging task that military officers have to undertake.

In recent years, many computational models and algorithms (Gilmer & Sullivan, 2003; Gilmour, Krause, Lehman, Santos, & Zhao, 2005; Kott & McEneaney, 2007; Pelta & Yager, 2009) have been developed to represent the process of adversarial intent inference. However, even though the bulk of these models are designed with the objective of creating decision support systems to help human decision makers predict adversarial intent, developing a computational tool that can truly match the adaptive human decision maker remains a challenging endeavor. This is especially so in an environment with incomplete information.

In addition, no two adversarial contexts are exactly alike, and the adversary can change his tactics and strategies. The adversary may even make irrational choices. Ultimately, the decision maker has to rely on his own deduction, and the computer tools of today (Kott & Ownby, 2005; Santos Jr & Zhao, 2007) can at best serve as additional analysis aids.

So how do human decision makers infer adversarial intent when the environment is new, or when the adversary has changed his rules? In general, humans can attribute mental states – beliefs, intents, desires, pretence, knowledge – to themselves as well as to others, and can understand that others have beliefs, desires, and intentions that may be different from their own. This is what psychologists term "theory of mind" (Premack & Woodruff, 1978) where individuals are able to form a theory of the mind of another person. This basic ability to attribute different beliefs and desires to the opponent is necessary for adversarial intent inference. In order to form mental states of the opponent, the decision maker may have to pick up cues in the environment, use historical data, obtain new information about the adversary, or "stand in the adversary's shoes" or "take the adversary's perspectives". The adversarial decision maker will need certain theory of mind processes or what Nichols and Stitch (2003) term as "mindreading".

Hence, Kott & Ownby (2005) argued that practical adversarial reasoning tools call for cross disciplinary approaches including artificial intelligence planning, control theory, machine learning, game theory and cognitive modeling. Although there is a growing interest to develop machine learning algorithms that predict adversarial intent, there is a lack of descriptive research on how humans derive adversarial intent from a naturalistic perspective.

In 1992, Thagard developed ECHO, a system that simulates explanatory coherence using connectionist models. ECHO can be used to describe and simulate human cognitive processes related to adversarial inference, ranging from business strategy to wargaming. Recursion is a central concept in this framework. MacInnes (2003) proposed that Thargard's framework of opponent modeling and deception can be demonstrated at zero to three levels of recursion: At Level 0, individuals have self awareness (What do I intend to do?). At Level 1, individuals have a basic concern about adversary (What does my opponent intend to do?). At Level 2, individuals have a meta-awareness of the adversary (What does my opponent think I intend to do?). At Level 3, individuals have a meta-meta-awareness (What does my opponent think I think he intends to do?)

We extend Thagard's and MacInnes' research by investigating the extent to which human decision makers think recursively about the enemy as they play a real-time strategy game -- a complex adversarial setting that can feature a dynamic environment with uncertainties. Specifically, we elicited mental models people have about the adversary by interviewing decision makers at semi-regular intervals during game play. By classifying the interview responses, we examined the frequency at which people naturally made predictions about the enemy at the various levels of recursion, as well as the distribution of accuracy achieved at each of the levels. We also qualitatively identified other supporting processes that aided adversarial intent inference.

Understanding the human processes used for understanding and predicting adversarial intent enables us to inform computational researchers on the processes to consider in their algorithms. This can potentially inspire new approaches to develop new inference engines and decision aids.

## 2. METHOD

Ten male military officers who had gone through Officer Cadet Training were recruited for the study: five were senior commanders aged 35-57 and five were junior commanders aged 22-26. All but three senior commanders had some previous experience playing real-time strategy games.

We employed a real-time strategy game Rise of Nations (RON) by Microsoft Games Studios to simulate a war gaming environment (Figure 2.1). Participants played the game against a confederate experimenter (Gamer) in another room via a local area network connection. They commanded ground and air forces by giving instructions to another experimenter (Cyber Warrior), who acted as a proxy to enact their plan in the game. This allowed participants to concentrate on the overall game strategy and served to mitigate differences between skills in controlling the game interface.

All participants first completed a series of interactive game tutorials, practice game sessions and a quiz to familiarize them with the rules, features, and controls in the Rise of Nations scenario. They then played seven game trials.

In these trials, participants were tasked to destroy and capture the enemy's city by occupying it with at least three infantry or armor units. They were presented with a map of the scenario and were given a list of the type and number of units and installations available in the game. The Gamer simulated a consistent adversary by following several opening moves based on a predetermined set of enemy's intentions. In the first four games, the Gamer would always send armored forces to attack the participant's northern defenses while launching a simultaneous air attack on the participant's airbase. From the fifth game onwards, the Gamer executed another script by sending armored forces from the south to attack the participant's airbase. This was designed to capture how people reacted to changes and how they calibrated themselves to the change.

Figure 2.1: Game scenario and interface in Rise of Nations, © Microsoft Corporation.

Participants were allocated ten minutes for planning before every game and were interviewed on the first, third, fifth, and seventh games. On those trials, they were interviewed after the planning session, and after every five to seven minutes of play once the game began. The variance of five to seven minutes were to avoid interrupting participants from the execution of series of his planned actions. A concluding interview was conducted when the game ended.

The interviews were semi-structured and involved asking participants about their understanding of past and current events in the game and their expectations of what would happen in the future. If the enemy featured in any aspect of their initial responses, participants were probed further to give a detailed account about what they thought about the enemy and their rationale for doing so. The interview methodology and questions were reported at the

International Command and Control Research and Technology Symposium (Tan, Yu & Soh, 2010). Such an interview protocol provided a naturalistic means of eliciting when participants thought about the enemy and the contents of their thoughts. This allowed us to examine the extent to which people naturally thought about their adversary through recursive thinking as they captured participants' initial experience with the enemy, learning and conditioning to the enemy's strategies, reaction to the change in script, and re-learning after the change in script.

The game play and interview sessions were recorded and transcribed using an event logging and analysis software, Morae, © Techsmith. Statements made about the enemy from the transcripts were first identified then classified based on their order of recursion, and were ascertained for their accuracy. This was accomplished by cross-referencing participant's predictions with the intentions and actions of the Gamer. The content of the predictions and the corresponding reasons given was also qualitatively analyzed.

## 3. RESULTS

A total of 661 inference statements about the enemy were identified from 40 interview transcripts, which can be organized into three broad categories:

*(1) Predicting future enemy goals or actions (Predictions).* This includes predicting what actions the enemy might perform next, including mentally simulating how the enemy might react to a proposed plan. For example, a participant might infer that the enemy's armored force would launch an attack on the participant's airbase soon, or how the enemy would respond when the participant sends special forces to attack an enemy surface-to-air missile site before the events manifested.

*(2) Observing and making sense of past enemy actions (Realizations).* For example, when the enemy attacked both the participant's bunker and airbase simultaneously, a participant might realize that the enemy's main objective was the airbase after the attack was complete, because he noticed that the enemy forces at the bunker retreated quickly, while the attack on the airbase was persistent.

*(3) Considering the enemy's capabilities, preferences and goals in their plan.* For instance, thinking about what type and how many enemy troops were left and how that would affect the subsequent battles, or how the playing style of the enemy could affect game play.

Subsequent analysis focused on the first category, which consists of 463 predictions that participants made about the enemy's future goals or actions based on incomplete information. Analyzing these predictions reveal that there was a wide variety of inferences made, and these differed in the order of recursion in thinking about the enemy and the supporting processes that aid prediction generation. There were differences noted in the level of detail and the certainty of predictions, though these aspects were not considered in detail.

### 3.1 Orders of recursion in adversarial intent prediction

Participants thought about the enemy at different orders of recursion. Participants thought about the enemy by making predictions and reasoning about the predictions in the first order most of the time (383 or 82%), where they directly thought about the enemy's actions, mental states and doctrine. All of the participants, at minimum, made predictions about the direction of the enemy's attack (north, south, center) by considering terrain features and unit capabilities. The participants also quickly formed impressions of the enemy's behavioral patterns across the initial games and used this information to make predictions about the enemy in subsequent games. However, there was only a select group of participants that thought about the enemy's personality (7 participants) and attention (1 participant).

Several predictions and their accompanying reasoning (77 or 17%) were made in the second order, where participants thought about what the enemy was thinking of them. This was exhibited by all of the participants, although in a much smaller quantity compared to the first order. For example, many participants reasoned that "the enemy knows I'm playing defensive and thinks I'll likely continue to be defensive", which led them to change their strategy and adopt an aggressive stance.

Finally, there were only four instances (1%) where predictions and reasoning reached a third order of recursion, where participants thought about how the enemy would think they thought about the enemy. These were expressed by only three participants. An example would be one participant who noticed that the enemy consistently attacked him from the north coast in the first three scripts, and anticipated that the Gamer would change strategy soon. He said, "I think in the next mission, the enemy will think very hard if he wants to go by the coast again. I think that enemy thinks that he has conditioned me to go the coast to counter him, so he will then go by the south."

On the whole, participants generally made more correct than incorrect predictions (see Table 3.1), although the ratio of correct: incorrect predictions decreases as the order of recursion increases.

Table 3.1: Accuracy of predictions by recursion order

| Recursion order | Correct | Incorrect | Partially Correct | Inconclusive | Total | Correct: Incorrect |
|---|---|---|---|---|---|---|
| First | 260 (70%) | 99 (27%) | 8 (2%) | 6 (2%) | 373 | 2.63 |
| Second | 51 (60%) | 30 (35%) | 2 (2%) | 2 (2%) | 85 | 1.70 |
| Third | 1 (20%) | 4 (80%) | 0 (0%) | 0 (0%) | 5 | 0.25 |
| Total | 312 (67%) | 133 (29%) | 10 (2%) | 8 (2%) | 463 | 2.35 |

Analyzing the frequency of predictions by trial games reveals that the mean number of predictions and reasoning made in the first order was similar in trial 1 (8.7) and trial 3 (7.8), however, the number increased drastically in trial 5 (12.6) which coincides with the Gamer changing script. The number then dropped back to 8.2 in trial 7 (Figure 3.1a). The number of predictions and reasoning made in the second order remained fairly constant between 2.3 and 2.4 from trials 1 to 3, and dipped slightly to 1.5 in trial 7.

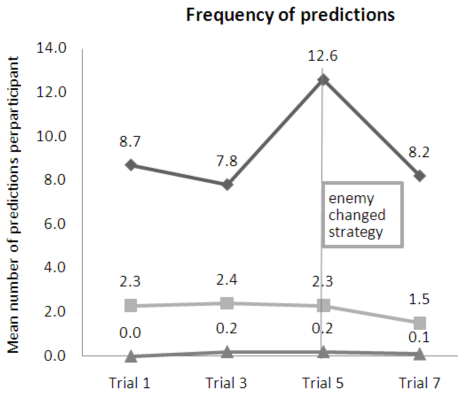


Figure 3.1a: Mean number of predictions per participant by trial game

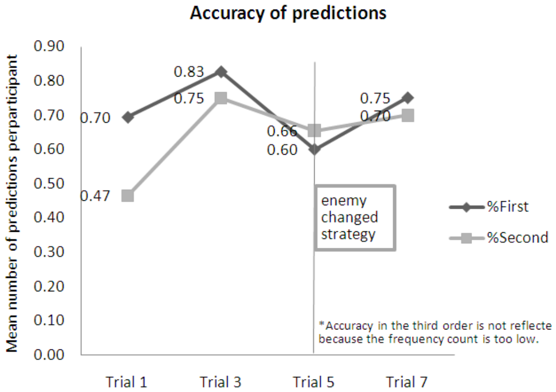


Figure 3.1b: Mean accuracy of predictions by trial game

The accuracy of predictions increases from trial 1 to trial 3 in both the first (0.70 to 0.83) and second (0.47 to 0.75) orders of recursion (Figure 3.1b). There was a subsequent dip in accuracy in trial 3 for both the first (0.66) and second (0.60) orders, which reflects the change in script by the gamer, followed by an increase in trial 7 in the first order (0.75) and a slight increase in trial 7 in the second order (0.60). Accuracy in the third order is not reflected in Figure 3.1b because the frequency count was extremely low.

### 3.2 Supporting processes that aid adversarial intent prediction

In addition to analyzing the predictions made by the participants, we also qualitatively analyzed the interview transcripts to draw themes that describe the supporting processes that participants used to derive their predictions. From the transcripts, it was discovered that because of the uncertain nature of the game scenario, participants relied heavily on gathering information to improve their situational awareness in the game.

This initial step of information seeking was essential for prediction generation, and was exhibited across all participants in a variety of ways:

*(1) Scouting.* All the participants positioned scouts at strategic locations so they had a better sense of where enemy troops were located. These scouts could be stationary, patrolling, or actively exploring specific areas in the map.

*(2) Probing reactions.* Some participants sent units near the enemy to probe for a response. (E.g. Sending decoys out to see how enemy reacts.)

*(3) Hypothesis testing.* When some participants had a hunch about what the enemy was doing or where the enemy might have been, they searched for evidence to verify their hypothesis. (E.g. Sending around the rear of the base if the participant suspects that enemy troops might launch a sneak attack from the back.)

Besides seeking information about the enemy by scouting, probing reactions and hypothesis testing, participants also gained a better understanding of enemy capabilities by actively tracking various units, buildings and situations as the game progresses. Participants varied in the type of information they tracked, ranging from the number of units lost or gained (through bribery) during battle, to the location and progress of certain task forces as they went about their mission, to the health status of particular units or buildings. Actively tracking such information enabled participants to refine their plans since they were acutely aware of their own and the enemy's current progress and capabilities. The information was then used for the reasoning of the current and future intents of the adversary.

## 4. DISCUSSION AND FUTURE WORK

Our results imply that participants seek information about the enemy in order to make predictions about what the enemy would do in the future. While making these predictions, participants generally think about the enemy in the first order most of the time, and think about what the enemy thinks of themselves only occasionally. Although the first two orders of recursion were used by all of the participants, few participants reached the third order of recursion. This leads us to wonder how and when people choose what orders of recursion to think at.

Results of the study suggest that familiarity with the enemy is a key aspect that influences the order of recursion because patterns in prediction frequency and accuracy vary according to the trial number and are consistent with the changes in the script executed by the Gamer. In trial 1, participants were unfamiliar with the enemy, and made a base number of predictions that was slightly more accurate than chance. By trial 3, however, participants became familiar with the enemy because the enemy used a similar strategy repeatedly, so it was natural for participants to predict more accurately at both the first and second orders of recursion compared to trial 1.

Participants encountered a drastic change in trial 5 as the Gamer deviated from his original script and attacked from a different direction. This is associated with a sudden spike in predictions, possibly because participants realized their initial predictions were incorrect as they were inconsistent with the enemy's behavior observed in the game. Hence, participants derived new insights from the Gamer's change in strategy which led to a new set of predictions that were more appropriate to the current game situation in trial 5.

Trial 7 was similar to trial 3 in that the participants were familiar with the enemy script because they had already encountered the Gamer using the strategy across multiple occasions. The familiarity with the new script led participants to make a similar number of predictions in the first order as in trials 1 and 3. These predictions were also highly accurate. However, there was a slight dip in the number of predictions made in the second order in trial 7, and these predictions was also less accurate. The drop in frequency and accuracy was rather surprising because we expected participants to be most familiar with the enemy as trial 7 was the last trial. Nevertheless, after a thorough qualitative analysis of the participant's plans and actions in trial 7, we discovered that many participants perceived the last trial to be special and wanted to end the experiment by surprising the enemy and trying out entirely new strategies. For instance, many participants who were usually defensive changed their approach and launched an aggressive attack on the enemy instead. Since their plans to aggressively attack the enemy were based on shaping what the enemy would do rather than anticipating and countering the enemy's actions, participants might have felt that thinking at the second order was more difficult and less relevant.

Therefore, apart from the slight deviation in frequency and accuracy in trial 7, the results show that participants made more predictions with greater accuracy as they became more familiar with the enemy.

The frequency distribution across the orders of recursion also indicates that people experience an effort vs speed tradeoff in the gaming environment. Thinking extensively about the adversary is a cognitively demanding task, and it is likely that the cognitive effort to generate predictions and their accompanying reasoning increases at higher orders of recursion. The third order of recursion is especially complex as it necessitates maintaining an awareness of multiple knowledge states in working memory. Yet participants do not have the luxury of time. As the game scenario forces participants to make decisions under time pressure, participants must quickly assess the environment to make predictions about the enemy so that they can act on the predictions and determine their next moves.

Furthermore, it appears that people are poor at making predictions at higher orders of recursion, as the ratio of correct: incorrect predictions decreases from 2.63, 1.70 and 0.25 when participants predictions at the first, second and third orders respectively. Predicting at higher orders can thus be seen as more difficult because of the increased complexity.

In order to explore the link between effort and order of recursion further, we conducted a subsequent pilot study that introduced four participants to the definitions of recursion and how it relates to making predictions about the enemy. The pilot participants then paired up and played the game scenario against each other. Informally questioning these pilot participants revealed that most participants felt that thinking at the first and second orders of recursion was intuitive to them, and they did so naturally. However, they felt that thinking at the third order of recursion was challenging and very tiring. Such feedback corroborates with the idea that thinking about the enemy at the third order of recursion is substantially more challenging and cognitively effortful compared to thinking about the enemy at the first and second orders.

Since people need to expand greater cognitive effort to think about the enemy at higher orders of recursion, and this greater cognitive effort is associated with lower accuracy, it is no wonder that people often stop at the first or second orders of recursion. So is thinking at higher orders of recursion even useful? More specifically, how might it possible to determine which order of recursion is optimal?

This study served as an initial attempt to describe how humans make predictions about the enemy in a complex environment based on the frequency and accuracy of the orders of recursion. Although it appears that thinking at higher orders is difficult and effortful, people do think at these orders, especially at the second order, albeit only occasionally. This implies that the order of recursion required may be situationally dependent.

Thagard (1992) illustrates that at minimum, people must think at the second order of recursion when planning to deceive an opponent. For a person to mislead his opponent, the person must have a rich model of the opponent, as well as a model of the opponent's model of himself, because the person's actions must be such that the opponent with either interpret them in the way that the person desires, or fails to interpret them in the way that the person does not desire (a double bluff). As an example, he states that a poker player who is not bluffing may display nervous behavior like chewing on knuckles so that his opponent might explain his behavior by thinking he is bluffing. As a result, the opponent stays in the pot, allowing the poker player to earn more money.

A similar thought process using the second order of recursion could apply in the game scenario when our participants try to execute a plan of deception. Moreover, there could be other situations where the second and third orders of recursion are necessary. One means of investigating this involves a detailed qualitative analysis of the participant's predictions and reasoning. We intend to accomplish this in the future by deriving other processes that are involved in adversarial intent prediction and relating these processes used to the order of recursion used, prediction frequency and prediction accuracy. We will also examine the content of the predictions, especially in the second and third orders, to study if there are any consistent themes that span across the examples. Being able to identify and relate specific themes and processes to human adversarial intent prediction is a crucial next step that will provide greater direction for human training and computational modeling in complex decision making.

## 5. REFERENCES


Gilmer, J. B. Jr. & Sullivan, J. F. (2003). The use of recursive simulation to support decision making. Paper presented at the 2003 Winter Simulation Conference.

Gilmour, D. A., Krause, L. S., Lehman, L. A., Santos, E., Jr., & Zhao, Q. (2005). *Intent driven adversarial modeling*. Paper presented at the 10th International Command and Control Research and Technology Symposium: The Future of C2.

Kott, A., & Ownby, M. (2005). *Tools for real-time anticipation of enemy actions in tactical ground operations*. Paper presented at the International Command and Control Research and Technology Symposium.

MacInnes, W. J., & Gilin, D. A. (2005). *What I think you think I am going to do next: Perspective-taking and recursive modeling in computer mediated conflict*. Paper presented at the International Association for Conflict Management 18th Annual Conference.

Nichols, S., & Stitch, S. P. (2003). *Mindreading : An Integrated Account of Pretence, Self-awareness, and Understanding Other Minds*. Oxford : Clarendon ; New York :: Oxford University Press.

Pelta, A. D., & Yager, R. D. (2009). *Dynamic vs. Static decision strategies in adversarial reasoning*. Paper presented at the Joint 2009 International Fuzzy Systems Association World Congress and 2009 European Society of Fuzzy Logic and Technology Conference.

Premack, D., & Woodruff, G. (1978). *Does the chimpanzee have a theory of mind? The behavioral and brain sciences*, 4, 515-526.

Santos Jr, E., & Zhao, Q. (2007). Adversarial Models for Opponent Intent Inferencing. In A. W. Kott & W. M. McEneaney (Eds.), *Adversarial Reasoning: Computational Approaches to Reading the Opponent's Mind* (pp. 1-22). Boca Raton: Taylor & Francis.

Tan, L. S. A., Yu, S., & Soh, B. K. (2010) Interruption Interview: An approach to elicit situation assessment for an ill-defined task. Paper presented at the $15^{th}$ International Command and Control Research and Technology Symposium.

Thagard, P. (1992). Adversarial problem solving: Modeling an opponent using explanatory coherence. *Cognitive Science*, *16*(1), 123-149.