

Reflective Cognition as a Secondary Task

Lisette Mol (LisetteMol@cmu.edu), Niels Taatgen (Taatgen@cmu.edu),
Department of Psychology, Carnegie Mellon University;
5000 Forbes Av., Pittsburgh PA 15213 USA

Rineke Verbrugge (L.C.Verbrugge@ai.rug.nl),
Institute of Artificial Intelligence, University of Groningen;
Grote Kruisstraat 2/1, 9712 TS Groningen, The Netherlands

Petra Hendriks (P.Hendriks@let.rug.nl)
Center for Language and Cognition Groningen, University of Groningen;
P.O. Box 716, 9700 AS Groningen, The Netherlands

Abstract

Our hypothesis is that reflective cognition is necessary to achieve expert level performance in certain skills, and that reflective cognition can be seen as a secondary task in skill acquisition. To investigate to what extent people use and acquire complex skills and strategies in the domains of reasoning about others and natural language use, an experiment was conducted in which it was beneficial for participants to have a mental model of their opponent, and to be aware of pragmatic inferences. Individual differences in the use of complex skills and strategies, that support our hypothesis, were found.

Introduction

In every day life, people frequently make use of their ability to reason about others and to infer the implicit meaning of sentences. Consider the following two situations:

Situation 1 You are called by a friend who asks you for a phone number. You know the number by heart, so you ask her whether she has pen and paper. She answers you with “No, I don’t”. Can you conclude that she also does not have a pencil and paper ready?

Situation 2 You are playing happy families and you are the first to pose a question. You ask your opponent for the ‘elephant’ of the family ‘mammals’. Your opponent replies with “No, I don’t have this card”. Can you conclude that he doesn’t have any member of the mammals family?

In the first case, you *know* that your friend *has the desire* to be cooperative and thus your reasoning would be something like, ‘She does not have a pencil, for if she did she would have told me so, since she knows it is relevant’. In the second case you *know* that your opponent *does not want* you to *know* which cards he has, since he has the desire to win the game. You therefore are aware that he would not tell you whether he has any other members of the family, unless he really had to, and thus you do not conclude that he does not have them.

These examples make it clear that people use their knowledge about the situation and about others to determine the meaning of a sentence. It would be interesting to know how humans use and acquire such skills. In the study described in this article, it has been investigated to what extent people use

and acquire complex skills in the domains of reasoning about others and language use.

In the next section, some theories relevant to the present study will be described. Then follows a description of the research question and the hypotheses that were stated. To test the hypotheses, an experiment was conducted in which participants played a game against each other. This game was a variant of the game Mastermind. The experimental setup and the predictions are described, followed by the results of the experiment. The last two sections contain a discussion of the results, the conclusions that can be drawn from this study, and some ideas for future work.

Background

Learning by Reflection

The classical theory of skill acquisition (Fitts, 1964), describes learning as a process of automation: one starts a new skill in the cognitive stage, in which controlled deliberate reasoning is needed to perform the task. This stage is characterized by slow performance and errors. By repeatedly performing the skill, eventually the autonomous stage is reached, where performance is fast and automatic, requiring little working memory capacity. During the process of automation, the control that one has over the process of performing the task decreases. Deliberate access to automated skills is therefore limited.

Although the classical theory can explain many phenomena, it is limited: Skills are usually considered in isolation, whereas in reality they build on one another. For example, the skill of multiplication is based on the skill of addition. However, according to the classical theory, mastered and hence automated skills cannot in themselves serve as a basis for more advanced skills, because deliberate access to automated skills is limited. Hence, it remains unclear how transfer of knowledge from one skill to another is possible.

In Karmiloff-Smith (1992), it is reported that children can only describe what they are doing after they have mastered a skill (e.g., in number conservation experiments). Thus, the capacity for deliberate reasoning sometimes increases rather than decreases with expertise. This cannot be explained by assuming skill acquisition to end in the autonomous stage.

We suggest that skill acquisition is a continuous interplay between deliberate and automatic processes. During the ini-

tial stages of acquiring a skill all deliberation is focused on basic performance of the task. Once performing the task starts to become more automatic, deliberate processes can shift to reflection of the task. Basic performance of the task and reflection on the task can be considered dual tasks, both competing for resources. Once performance becomes more automated, more resources become available for reflection. Reflection allows using the skill as a building block for more complex skills, looking ahead a few steps, or, as is the case in the present study, reasoning about others' knowledge (also see Taatgen (2005), for examples of how automation can improve flexibility in reasoning).

We assume that to reach expert level performance in domains such as reasoning about others, pragmatics, and learning from instruction, deliberate reasoning processes, such as self-monitoring, are crucial.

Theory of Mind Use

One of the advanced skills that we are interested in is the use of Theory of Mind (ToM). Although children from the age of six are able to distinguish between their own mental states and those of others, Keysar, Lin, and Barr (2003) argue that even adults do not reliably use this sophisticated ability to interpret the actions of others. They found a stark dissociation between the ability to reflectively distinguish one's own beliefs from others', and the routine deployment of this ability in interpreting the actions of others. The second didn't occur in their experiment.

To have a first order ToM is to assume that someone's beliefs, thoughts and desires influence one's behavior. A first-order thought could be: 'He does not know that his book is on the table'. In a second-order ToM it is also recognized that to predict others' behavior, the desires and beliefs that they have of one's self and the predictions of oneself by others must be taken into account. So, for example, you can realize that what someone expects you to do will affect his behavior. A second-order thought could be: 'He does not know that I know his book is on the table'. To have a third order ToM is to assume others to have a second order ToM, etc.

In defining the different orders two choices have been made. The first is that to increase the order, another agent must be involved. 'I know his book is on the table' and 'I know I know his book is on the table' are said to be of the same order. So for the order to increase, the agents the knowledge is about must be different.

The second choice made is to consider both 'I know *p*' and *p* to be zeroth order knowledge. This mainly is a matter of speech. The fact *p* in itself, which can be true or false, only becomes knowledge when it is known by someone. So only when someone knows that *p*, *p* can be considered zeroth order knowledge.

Pragmatic Inferences

Besides ToM reasoning, a second skill that has been investigated is language use, especially drawing pragmatic inferences. According to Grice (1989), people use the quantity maxim to infer the implicit meaning of a sentence. The quantity maxim states that interlocutors should be as informative as is required, yet not more informative than is necessary.

Using the quantity maxim it can be inferred that, for example, if a teacher says 'Some students passed the test', it

is the case that not all students passed the test. This is because if all students would have passed the test, the teacher would have used the more informative term *all* instead of the weaker term *some*, since otherwise the quantity maxim would have been violated.

Some and *all* are scalar terms. Scalar terms can be ordered on a scale of pragmatic strength. A term is said to be stronger if more possibilities are excluded. An example is (*a, some, most, all*) which is ordered from weak to strong. The above example is an example of a scalar implicature. In case of a scalar implicature, it is communicated by a weaker claim (using a scalar term) that a stronger claim (using a more informative term from the same scale) does not hold.

Feeney, Scafton, Duckworth, and Handley (2004), propose that there are three stages to people's understanding of *some*:

- (a) the logical (truth-conditional) interpretation which precedes children's sensitivity to scalar implicatures,
- (b) the pragmatic interpretation which results from drawing pragmatic inferences,
- (c) a logical interpretation that results from choice rather than from the incapability to make the pragmatic inference.

The first two stages are in line with the results in Noveck (2001) and Papafragou and Musolino (2003). Feeney et al. (2004) found evidence for a third stage, in which adults can choose a logical interpretation over a pragmatic interpretation, even though they can make the pragmatic inference that *some* implies *not all*. They conducted an experiment in which undergraduate students performed a computerized sentence verification task. They recorded the student's answers and reaction times. Here are two of the *some* sentences they used.

1. Some fish can swim.
2. Some cars are red.

Feeney et al. (2004) found that for participants who gave logical responses only, reaction times for responses to infelicitous *some* sentences such as 1 were longer than those for logically consistent responses to felicitous *some* sentences as 2. Notice that to both sentences the logical response is 'true'. The pragmatic response to 2 is 'true' as well. The pragmatic response to 1 is 'false'. So the sentences in which the logical and pragmatic response are in conflict resulted in longer reaction times. These results favor a theory that logical responses are due to inhibition of a response based on the pragmatic interpretation over a theory that logical responses result from failure to make the pragmatic inference. This suggests that a more logical language use in adults can be seen as an advanced skill.

Research Question and Hypotheses

The context described in the previous section leads to the following problem statement: *How do deliberate and automatic processes interact in the acquisition of complex skills?* The study described in this article is an exploratory study, for which the following research question is stated: *To what extent do people use and acquire complex skills and strategies,*

in the domains of reasoning about others and language use. This is narrowed down to the specific case of playing Master(s)Mind(s), a symmetric version of the game Mastermind, which is described by Kooi (2000). A variant of this game is used in the experiment described in the following section. To find an answer to the research question, three hypotheses are stated.

Hypothesis 1 *Performing a task and simultaneously reflecting upon this task can be seen as a form of dual tasking.*

This hypothesis states that when people perform a task which involves reasoning with incomplete information, or drawing pragmatic inferences, reflection can be considered a secondary task. The first task includes reasoning based on one's own knowledge and the truth-conditional (e.g., logical) meaning of utterances. The second task is more complex, and includes using reflection to reason about others, and to infer from pragmatically implicated meaning. These tasks compete for resources and their demands decrease with skill acquisition.

When playing Master(s)Mind(s) (see next section), the first task is to play the game according to its rules. This involves reasoning about the game rules and determining which sentences are true. The second task is to develop a winning strategy for the game. This involves reasoning about what the opponent thinks, is trying to make you think, or thinks that you are trying to make him think, etc., as well as determining what is pragmatically implicated by an utterance, or which utterances reveal the least information while still being true.

Hypothesis 2 *In an uncooperative conversation, people will shift their interpretation and production of quantifiers from a pragmatic (using Grice's quantity maxim) to a less pragmatic (not using Grice's quantity maxim) use.*

The idea behind hypothesis 2 is that in an uncooperative situation, people will be aware that others are trying to reveal little information (first order knowledge) and therefore will be aware that the quantity maxim does not hold. They will therefore not use the pragmatic inferences that they usually do in interpretation. In addition, people may develop more logical productions to be less informative themselves.

The reasoning necessary for this change in strategy is part of the secondary task of reflective cognition. Therefore, people will only be able to make this change when the first task is sufficiently automated.

Hypothesis 3 *In using quantifiers, people make use of an automated process, which results in a pragmatic use of the quantifier. This automated process can be 'overruled' by a deliberate reasoning process, which results in a logical use of the quantifier.*

Hypothesis 3 is on what kind of reasoning is involved in using quantifiers, especially to make the shift described in hypothesis 2. The theory of three stages that is proposed by Feeney et al. (2004) seems in line with the theory of skill acquisition we propose. If so, the process of making pragmatic inferences should be an automated process and the ability to overrule this pragmatic interpretation would result from reflective cognition. Once the demands of the first task have decreased sufficiently, this reflective cognition can take place, resulting in the change of strategy described in hypothesis 2.

Experimental Setup

Participants (native Dutch speakers) had to complete two sessions, each of about three hours, in which they played a symmetric head to head game via connected computers. In this game they had to correctly guess the secret code, consisting of four different, ordered colors, of their opponent. Players gave each other feedback by selecting Dutch sentences from a list. Although not explicitly told to participants, these sentences differed in pragmatic strength. The game was about gaining as much information as possible, while at the same time revealing as little information as possible. Because of this second aspect, the conversation is not fully cooperative and thus hypothesis 2 is relevant.

During the game, players had to submit their interpretation of the sentences they received as feedback. They had to submit all the 'worlds' that they thought to be possible given the feedback sentences, using a code: For each right color in the right position they had to select a black circle and for each color which was correct but in the wrong place, a white circle. To represent ambiguity, participants could submit more than one combination of black and white circles that they considered possible.

Let's look at an example. Imagine John having the secret code 1 = red, 2 = blue, 3 = green, 4 = yellow and Mary guessing 1 = red, 2 = orange, 3 = yellow, 4 = brown. The evaluation of this situation is that exactly one guessed color is right and in the right place (red) and exactly one guessed color is right, but in the wrong place (yellow). John has to choose two feedback sentences to send to Mary, one about color and one about position. He could say 'Some colors are right.' and 'There is a color which is in the right place.' This would indicate that John thinks that *some* can mean *exactly two* and that *a* can mean *exactly one*. This is a pragmatic production (in accordance with Grice's maxims). If he had chosen the sentence 'One color is right.', then he would allow *one* to mean *exactly two*. This would be a more logical production (in logic *one* is true in case of *at least one*).

Mary now has to give her interpretation of the sentences chosen by John. So if she thinks that, given the first two sentences, it could be the case that two colors are right, of which one is in the right position, she would submit (black, white) as a possible interpretation. If she considers the situation where three colors are right, of which two colors are in the right position, possible as well, she would also submit (black, black, white). If she would only submit the first possibility, her interpretation would be pragmatic. If she would also submit the second case, her interpretation would be more logical.

In the experiment Mary would have to give John feedback about her guess compared to her own secret code as well, and John would then submit his interpretation of those sentences. Each turn, one player can make a guess, in this example Mary.

During the experiment participants had to answer questions. The purpose of those questions was to get information on their strategy and the order of the theory of mind they were using. For the same purpose, participants completed a questionnaire after each session. More details on this experiment and the results can be found in (Mol, 2004).

Predictions

Since the game Master(s)Mind(s) involves many actions which need to be performed each turn, participants are expected to start with a very simple or no strategy. As they get more experienced in playing the game they will have enough resources left to develop a more complex strategy.

Grice’s maxims are best applied in situations where conversation is cooperative. Since a rational strategy for playing the game in the experiment is to be as uninformative as possible, communication will probably not be cooperative in the experimental conditions. So once the participants have mastered the game well enough to think about strategy and have become familiar with the uncooperative context, they are expected to develop a less pragmatic use of the sentences.

It is expected that while playing the game, the order of the theory of mind used by the participants increases. This will lead to the participant considering the amount of information that is revealed by the feedback sentences chosen, and the amount of information that will have to be revealed as a result of a guess made (first order ToM). The participant will also become aware that his opponent is trying to reveal little information (second order ToM). This will lead to a more logical interpretation. Eventually, the participant may use the knowledge that his opponent knows that he is trying to hide certain information (third order ToM).

Individual differences in what order of ToM will be used and how logical language use becomes are expected, as well as individual differences in the speed of developing a better strategy. Since the logical language use participants eventually reach results from a conscious reasoning process, participants are expected to be able to describe this part of their strategy.

Results

The participants are numbered from 1 to 12. Participants 10, 11 and 12 completed only one session.

Table 1: Highest Order of ToM used. This table shows the highest order of ToM that participants used during the experiment. The numbers represent the participants. The order used was determined from the answers participants gave to questions that were asked during the experiment.

1st order	possibly 2nd order	2nd order
3, 5, 6, 7, 12 8, 9, 10	4	1, 2, 11

Three out of twelve participants showed clear signs of the use of second order ToM (table 1). One additional participant probably used second order ToM as well, but in this case it was less clear. An example of second order ToM use in this game is that agent 1 assumes that the guesses made by agent 2 are evasive about agent 2’s own code, since agent 2 does not want agent 1 to know agent 2’s secret code. All of these four participants played in accordance with a strategy of being uninformative (table 2) and had a fairly to strict logical language use (table 3).

Table 2: Strategy. This table shows what kind of strategy participants used during the experiment, initially and finally. The strategy was determined from answers that participants gave to questions posed during the experiment. The numbers of the participants who made a shift are in *italic* in the row that represents the final strategy.

	being uninformative	being informative	other
initially	1, 2, 4, 5, 10, 11	3, 8, 9, 12	6, 7
finally	1, 2, 3, 4, 5, 11	9, 12	<i>2, 3, 6, 7, 8 10</i>

Table 3: Language use. This table shows the type of productions and interpretations (logical or pragmatic) of participants during the experiment, initially and finally. The numbers represent the participants. The numbers of the participants who made a shift are in *italic* in the row that represents the final language use.

	pragmatic	fairly pragmatic	fairly logical	logical
initially	8	5, 6, 7, 9, 10, 12	1, 2, 3, 4	11
finally	<i>6, 7, 8, 12</i>	9, 10	1, 2, 3, 4, 5	11

The remaining eight participants all used first order ToM. An example of first order ToM use in this game is that agent 1 takes into account what agent 2 already knows about agent 1’s secret code. Two of these participants had a strategy of being uninformative and a fairly logical language use, similar to the participants who used second order ToM. The other six used the strategy of being informative or a strategy which did not consider the amount of information being revealed and had a fairly to strict pragmatic language use.

All participants with a strategy of being uninformative and a fairly to strict logical language use showed a type of behavior which the others did not show (table 4). This behavior consists of preferring less informative sentences to more informative ones in production. For example, favoring sentence 1 over sentence 2 in a case where, from a logical perspective, they both hold.

1. ‘Some colors are right.’
2. ‘All colors are right.’

All participants who used second order ToM did so from the start. No shifts in order of ToM used were observed. Some shifts were measured in language use. One participant shifted from a fairly pragmatic to a fairly logical use. This participant had a strategy of being uninformative. Three participants shifted from a fairly pragmatic to a fully pragmatic use. They did not use a strategy of being uninformative. The other participants were constant in their language use.

Table 4: The preference for uninformative sentences. This table indicates which participants preferred to use less informative sentences in production. The numbers represent the participants. The numbers of the participants who made a shift are in *italics* in the row that represents the final behavior.

	preferred less informative sentences	did not prefer less informative sentences
initially	1, 3, 4, 5, 11	2, 6, 7, 8, 9, 10, 12
finally	1, 2, 3, 4, 5, 11	6, 7, 8, 9, 10, 12

One participant shifted from a strategy of being informative to a strategy of being uninformative. This participant had a fairly logical language use. One participant abandoned the strategy of being uninformative, to give the opponent a better chance of winning (!). This participant had a fairly pragmatic use of language.

The participants using more advanced strategies clearly had to put little effort into playing the game and understanding the computer program used. The people with the least advanced strategies made more mistakes in playing the game than others.

Most participants wrote down thoughts on the meaning of scalar terms, the terms they considered possible and their strategy in their answers to the questions posed during the experiment.

Discussion and Conclusion

The research question for this study was: *To what extent do people use and acquire complex skills and strategies, in the domains of reasoning about others and language use.* It was found that four out of twelve participants used all complex skills that were measured: a strategy of being uninformative, using logical interpretation and production, and using second order ToM reasoning. By the end of the experiment half of the participants had a logical interpretation and production and also half of the participants had the strategy of being uninformative. One participant developed this strategy during the experiment. Another participant developed a more logical language use during the experiment. A third participant developed a better strategy in production.

There clearly were individual differences. Some participants did not seem to use complex skills and strategies. Some participants developed a more pragmatic language use. This may be because the pragmatic meaning of some of the scalar terms used was dependent on the situation and therefore was not yet fully automated for the context of playing Master(s)Mind(s).

Hypothesis 1 stated that performing a task and simultaneously reflecting upon this task is a form of dual-tasking. Although participants showed little development during the experiment, the results are in line with this hypothesis. The four participants that showed the use of all complex skills did so from the start of the experiment. It could be the case that the first task, playing the game according to its rules, was relatively easy for these participants. They made few or no mis-

takes, and had relatively much experience in working with computers and playing strategic games. Also, three of them indicated to have a fair knowledge of logic.

There were two participants that had a strategy of being uninformative and a logical language use, but did not seem to use second order ToM. These were the participants that developed either their strategy or their language use during the experiment. These participants may represent an intermediate stage in which some resources are available for reflective cognition, but not yet enough to use all the complex skills.

Half of the participants did not show the use of any complex skills. Some of them developed strategies which they described as making things difficult for the opponent, but they did not relate this to the amount of information they revealed. It could be that these participants were too occupied with the first task to use reflective cognition, or that they did not think of logical language use as a possible way to use language. These participants made relatively many mistakes.

Hypothesis 2 stated that in an uncooperative situation, people will shift their interpretation and production of quantifiers from pragmatic (using Grice's quantity maxim) to less pragmatic (not using Grice's quantity maxim). None of the participants developed a more logical language use in the way that was meant in hypothesis 2. However, by the end of the experiment, half of the participants had a fairly to strict logical language use. As explained above, it could be the case that these participants had enough resources left for reflective cognition from the start of the experiment, so that they could use logical productions and interpretations.

Hypothesis 3 stated that in interpreting and producing quantifiers, people make use of an automated process, which results in pragmatic use of the quantifier, and that this automated process can be 'overruled' by a deliberate reasoning process, which results in logical use of the quantifier. It seems that pragmatic language use is not automated for all people in the situation of the experiment, since some participants developed pragmatic language use while repeatedly playing Master(s)Mind(s). This could be because of the context dependent meaning of some of the scalar terms used. It is unlikely that half of the participants were unable to use pragmatic language use, since they were all adults. Thus, it seems that people can indeed choose not to use pragmatic language use. Since most participants wrote down comments on their way of interpreting the scalar terms, it seems that changing this interpretation is indeed a deliberate reasoning process.

Future Work

In future work, more evidence for or against hypothesis 1 has to be found. The difficulty of the first task needs to be varied, to investigate whether this influences the use of skills that result from reflective cognition. In the Master(s)Mind(s)-experiment, there are several ways to do so. The interface of the computer program used could be made less user friendly, time pressure could be added, and the number of colors in a secret code could be varied. Also, a less well known game similar to Master(s)Mind(s) could be used, because most people know the regular version of the game Mastermind, which can be a benefit in playing Master(s)Mind(s).

An improvement in the experimental setup should be made to better be able to measure complex skills and strategies.

Participants with pragmatic language use had a disadvantage in strategy development. A strong strategy for this game is to reveal little information. The less informative sentences that logical language users could prefer often were regarded as false by pragmatic language users such that they could not use these sentences. By including more expressions, such as for example *niet alle* (*not all*), the possibilities for pragmatic language users can be increased. In addition, logical language use should be taught to participants prior to the experiment, to make sure that all participants are aware of the possibility of using language in this way.

An alternative for hypothesis 2 could be: *In an uncooperative conversation, some people will use quantifiers in a way that is not in accordance with Grice's quantity maxim, but more truth-conditional.* To test this hypothesis, it should be investigated whether the cooperativeness of the situation has an influence on language use. This could be done by observing the language use of the participants who had a logical language use during the Master(s)Mind(s)-experiment, while they play a fully cooperative game, in which a mutual goal has to be reached by two or more players.

To make it more clear whether or not logical language use can only result from overruling pragmatic language use, as stated in hypothesis 3, it would be interesting to let the participants to the Master(s)Mind(s) experiment do an experiment like the one that was conducted by Feeney et al. (2004). This could also be done for other scalar terms than *some*. Such an experiment could reveal whether the participants who had a logical language use from the start still need to overrule their pragmatic language use. If participants were to complete such an experiment before and after doing the Master(s)Mind(s)-experiment, people who have shifted to more logical use are expected to have increased reaction times, since they now have to overrule their automated interpretation process.

In addition to conducting more experiments, cognitive modeling could also be used to find answers to the remaining questions. This could be particularly helpful in determining what kind of reasoning processes, automated or deliberate, are involved in using scalar terms and theory of mind reasoning. Also, it could be investigated what parameters, such as for example working memory capacity, correlate with the use of a particular order of ToM reasoning and a particular type of language use.

References

- Feeney, A., Scafton, S., Duckworth, A., Handley, S. J. (2004). The story of some: Everyday pragmatic inference by children and adults. *Canadian Journal of Experimental Psychology*, 58:90–101.
- Fitts, P. M. (1964). Perceptual-motor skill learning. In *Categories of human learning*. New York: Academic Press.
- Grice, P. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Karmiloff-Smith, A. (1992). *Beyond modularity: A Developmental Perspective on Cognitive Science*. Cambridge, MA: MIT/ Bradford Books.
- Keysar, B., Lin, S., Barr, D. J. (2003). Limits on theory of mind use in adults. *Cognition*, 89:25–41.
- Kooi, B. (2000). Master(s)mind(s). See <http://www.philos.rug.nl/~barteld/>.
- Mol, L. (2004). Learning to reason about other people's minds. Master's thesis, University of Groningen.
- Noveck, I. (2001). When children are more logical than adults: experimental investigations of scalar implicature. *Cognition*, 78:165–188.
- Papafragou, A. & Musolino, J. (2003). Scalar implicatures: experiments at the semantics-pragmatics interface. *Cognition*, 86:253–282.
- Taatgen, N. A. (2005). Modeling parallelization and flexibility improvements in skill acquisition: From dual tasks to complex dynamic skills. *Cognitive Science*, 29:421–455.