

A process model of immediate inferences

Sangeet Khemlani¹, J. Gregory Trafton¹, Max Lotstein², and P. N. Johnson-Laird²
khemlani@aic.nrl.navy.mil, trafton@itd.nrl.navy.mil, lotstein@princeton.edu, phil@princeton.edu

¹NCARAI, Code 5515, Naval Research Laboratory, Washington, DC 20375 USA

²Department of Psychology, Princeton University, Princeton, NJ 08540 USA

Abstract

Individuals can make inferences from a single quantified premise. For instance, if you know that *all of the Virginians are students*, you might infer that *some of the students are Virginians*. We describe a computational system, *mReasoner*, of the cognitive processes that underlie these so-called ‘immediate’ inferences. The account is based on the assumption that when individuals understand discourse, they construct discrete mental simulations, i.e., mental models, of the assertions in the discourse. To draw conclusions, reasoners describe the relation between the individuals in the models and, if they are prudent, they search for alternative models to corroborate or refute a conclusion. We describe an experiment in which participants’ carried out a series of immediate inferences, and present a simulation that predicts the accuracy and latency of their responses.

Keywords: quantifiers, mental models, *mReasoner*, immediate inferences

Introduction

Reasoners can make immediate, rapid inferences from a single quantified assertion such as, *None of the Xs are Ys*. For instance, if they know that *none of the lawyers in the room are men*, they might refrain from asking any of the men in the room for legal advice, because they can infer:

1. None of the lawyers are men.
2. Therefore, none of the men are lawyers.

The inference is *valid* because its conclusion must be true given that its premise is true (Jeffrey, 1981, p. 1). It is also easy to make in comparison with more complex reasoning problems, such as syllogisms based on two quantified premises (for a review, see Khemlani & Johnson-Laird, in press a). Psychologists have investigated immediate inferences for many years (e.g., Begg & Harris, 1982; Newstead & Griggs, 1983; Wilkins, 1928), but have yet to resolve how logically untrained individuals make them. We have accordingly formulated a theory based on mental models and implemented it computationally in a unified model-based reasoning system called *mReasoner*.

In the present paper, we outline the theory and derive some novel predictions from it. We then report the results of a study that tested these predictions, and we show how the theory provides an satisfactory process model of individual performance.

Immediate inferences

In immediate inferences based on singly-quantified assertions, we studied 4 different *moods* for the premise:

All the Xs are Ys
Some of the Xs are Ys
None of the Xs are Ys
Some of the Xs are not Ys

and 8 different sorts of conclusion (4 moods by 2 *figures*, i.e., arrangements of terms ‘X’ and ‘Y’). Therefore, there are 32 possibly immediate inference problems based on these premises. The reasoner’s task was to assess a given conclusion’s status with respect to the premise, i.e., whether the conclusion *must* be true or whether it *might possibly* be true. Hence, it must be the case that some of the students are Virginians given than All the Virginians are students. And, it is possible but not necessary, that some of the students are not Virginians.

A robust theory of immediate inference should specify an algorithm that accounts for how individuals represent quantified assertions, how they assess whether a conclusion is at least possible, and how they decide whether it holds of necessity. It should also explain the relative difficulty of the various sorts of immediate inference, i.e., both the accuracy of participants’ conclusions and the latency of the correct conclusions. We developed such a theory as part of a general account of quantificational reasoning, and describe its assumptions below.

Mental models and quantifiers

Quantified assertions such as *None of the Xs are Ys* are can be treated as referring to relations among the set of *Xs* and the set of *Ys* (see, e.g., Cohen & Nagel, 1934, p. 124-5). Psychological theories of how quantifiers are interpreted follow suit (but cf. Braine & O’Brien, 1998; Rips, 1994), though they differ in the way they treat the relations between the sets. For instance, some theorists make use of diagrammatic representations to handle relations (Ceraso & Provitera, 1971; Erickson, 1974; Ford, 1995; Newell, 1981), others rely on formal rules of inference (Geurts, 2003; Guyote & Sternberg, 1981; Politzer, van der Henst, Luche, & Noveck, 2006; Stenning & Yule, 1997); and yet others analyze sets in terms of simulated possibilities, i.e., mental models (Bucciarelli & Johnson-Laird, 1999; Johnson-Laird & Byrne, 1991; Polk & Newell, 1995). The psychological systems can all account for how individuals make valid deductions, however few of them can account for the differences in relative difficulty between various reasoning problems. The present theory relies on mental models to explain the processes that give rise to valid and invalid

responses, as well as the differences in difficulty between various sorts of immediate inference.

The mental model theory proposes that when individuals comprehend discourse, they construct simulations of the possibilities consistent with the contents of the discourse (Johnson-Laird, 2006). The theory depends on three main principles: 1) Individuals use a representation of the meaning of a premise, an *intension*, and their knowledge, to construct mental models of the various possibilities to which the premises refer. 2) The structure of a model corresponds to the structure of what it represents (see Peirce, 1931-1958, Vol. 4), and so mental models are iconic insofar as possible. 3) The more models a reasoner has to keep in mind, the harder an inference is. On a model-based account, a conclusion is *necessary* if it holds in all the models of the premises and *possible* if it holds in at least one model of the premises.

Inferential processes

The theory proposes that the reasoning system can carry out three processes whenever individuals reason about whether a given conclusion follows from premises. First, the parser constructs a representation of the meaning of each premise (an *intensional* representation) based on a linguistic analysis. Second, the system uses the intension to construct a model of the situation to which the assertion refers (an *extensional* representation). Third, the system checks whether the given conclusion holds in the model. These three processes are carried out in “system 1” (see, e.g., Evans, 2003, 2007, 2008; Johnson-Laird, 1983, Ch. 6; Kahneman, 2011; Sloman, 1996; Stanovich, 1999; Verschueren, Schaeken, & d’Ydewalle, 2005) because they yield rapid responses from a single mental model.

The theory also postulates an advanced set of “system 2” processes, which are used to evaluate and, if necessary, to correct initial inferences. They search for alternative models of the premises in which both the premise and the conclusion are true. We explain how these processes work in the following section.

mReasoner: A computational theory of reasoning

mReasoner v0.9 is a new, unified computational implementation of the mental model theory of reasoning. It implements three general systems:

- a) A linguistic system for parsing premises and building up intensional representations. This system’s purpose is to map an assertion’s syntax to an underlying semantic form (the intension).
- b) An intuitive system (1) for building an initial extensional representation, and drawing rapid inferences from that representation.
- c) A deliberative system (2) for more powerful recursive processes that search for alternative models. This system can manipulate and update the representations created in System 1, and it can modify conclusions, but

it too can fall prey to systematic errors (Johnson-Laird & Savary, 1999; Khemlani & Johnson-Laird, 2009).

System 1 is faster than system 2, and as a result it is more prone to errors. Below, we describe the various processes that each system implements.

The linguistic system

Parsing statements into intensions. An *intensional* representation is composed from the meanings of words and the grammatical relations amongst them. The first process in mReasoner is a shift-and-reduce parser (Hopcroft & Ullman, 1979) that uses a context-free grammar and a lexicon to compose the intensional representations of a sentence. We make no claims about the psychological reality of a shift-and-reduce parser or a context-free grammar, which we adopt for convenience. The grammatical rules and lexical entries consist of a word (such as “all”), its part of speech (“determiner”), and a specification of its semantics. The parser applies the appropriate semantic rule, which matches the syntactic rule it has used, to construct an intension of the current constituent of the sentence. A key assumption is that the semantics of a quantified assertion sets the values of parameters that constrain the construction of models.

The present set of parameters is presented in Table 1 for the assertions that occur in the immediate inferences under investigation, and for two representative examples of other sorts of quantified assertions: *Neither A is a B*, and *Five A are B*. In sum, intensions are collections of parameters that, as a whole, specify the semantics of an assertion.

Table 1: A summary of mReasoner’s parameters in the intensions of different sorts of singly-quantified assertion. The parameters are as follows: i) the cardinality of the overall set of entities and its boundary conditions; ii) the cardinality of the set referred to by the quantifier (e.g., the As); iii) the boundary on the two sorts of cardinality as specified in (i) and (ii); iv) the polarity of the determiner; and v) the universality of the quantifier. Where relevant, ‘?’ signifies a default value that can be modified provided it meets the boundary conditions in (i) and (iii). A sixth parameter specifies the set-theoretic relation between the As and the Bs, and for all the examples below, the parameter is set to the inclusion relation, which specifies that the As are included within the set of Bs in the manner described by the other parameters.

Assertion	The five parameters in singly-quantified intensions				
	i	ii	iii	iv	v
All As are Bs	?4 ≥ 1	?4	iii = i	positive	universal
Some As are Bs	?4 ≥ 1	?2	0 < iii ≤ i	positive	particular
No As are Bs	?4 ≥ 1	?4	iii = i	negative	universal
Most As are Bs	?4 ≥ 1	?3	½ i < iii < i	positive	particular
Neither A is a B	2	2	iii = i	positive	universal
Five As are Bs	5	5	iii = i	positive	universal

System 1

Model building. The system uses the intension of a premise to build an initial model, and it updates that initial model if additional premises occur. The model is built in accordance with the parameters of the intension. The system begins by building a model with a small set of individuals. For example, the model of *All the artists are bohemians* is built by first constructing a set of artists:

artist
artist
artist
artist

In the diagram above, each row represents an individual with the property of being an artist, and so the model as a whole represents a finite number of individuals. Mental models are representations of real individuals, not letters or words, which we use here for convenience. The inferential system in mReasoner is able to treat the model above as representing “all artists” and not, say, “exactly four artists” because it has access to the intension of the premise, which constrains the possible interpretations of models and therefore the possible modifications to models.

The intension of *all the artists are bohemians* also specifies the number of artists that are also bohemians. The model is updated accordingly:

artist bohemian
artist bohemian
artist bohemian
artist bohemian

At this point, the premise has been represented, and so the system assesses whether the given conclusion is true in the initial model.

Assessing initial conclusions. When reasoners have to assess a given conclusion, the system inspects the initial model to verify that the given conclusion holds or does not hold. For instance, suppose that reasoners are asked to decide whether it is possible that *some bohemians are not artists* given the previous premise. From the model above, the system initially responds in the negative, i.e., the putative conclusion is impossible. The process is simple, and the response is rapid. But, as we show in the next section, it is also fallible.

In many experiments and in daily life, reasoners have to draw their own conclusions. mReasoner accounts for this process too. The model above appears to support any of the following conclusions:

All the artists are bohemians
All bohemians are artists
Some artists are bohemians
Some bohemians are artists
Four artists are bohemians

But, the theory assumes that reasoners prefer to scan their initial model in systematic ways, and the computational system implements several heuristics that explain the general biases reasoners exhibit when they draw conclusions from inspecting a model of two quantified premises. Researchers often place heuristics at the forefront of theories of reasoning (e.g., Chater & Oaksford, 1999), but until now proponents of the model theory have downplayed their application. To unify model-based accounts of reasoning with heuristic-based systems, heuristics play a central role in inferring an initial conclusion (see Khemlani, Lostein, & Johnson-Laird, in press b, for an extended discussion).

The system’s ability to assess and generate initial conclusions is fallible. For instance, one can indeed show that *some of the bohemians are not artists* is possible, though the system infers initially that it is impossible. To revise its initial conclusion, the system needs to find an alternative model in which both the premise and conclusion hold. We turn to the final process in mReasoner to explain how such a model is found.

System 2

Searching for alternative models. In the preceding section, we focused on how mReasoner assesses conclusions based on an initial model. However, the conclusions it draws can be invalid. System 2 attempts to revise initial conclusions by searching for alternative models. To do so, it uses three separate operations: *adding* properties to individuals, *breaking* one individual into two separate individuals, and *moving* properties from individual to another (see Khemlani & Johnson-Laird, under review). These operations correspond to those that naïve participants spontaneously adopt when they reason about syllogisms (as evidenced by their manipulations of external models, see Bucciarelli & Johnson-Laird, 1999). In the earlier example, the system finds an alternative model by adding a new individual to the initial model to yield:

artist bohemian
artist bohemian
artist bohemian
artist bohemian
bohemian

The new individual, who is bohemian but not an artist, and the resulting model refutes the conclusion, *All the bohemians are artists*. But, the conclusion, *Some of the bohemians are artists*, still holds, and no model refutes it.

Predictions. The theory and its computational implementation distinguish between the relative difficulty of three sorts of immediate inference:

- a) zero-model inferences
- b) one-model inferences
- c) multiple-model inferences

Zero-model inferences are those in which the conclusion is identical to the premise, and so individuals needn't even build a model to be able to solve the problem. For instance, consider the following problem:

All the aldermen are barbers.
Is it possible that all the aldermen are barbers?

The reasoner should realize that the answer is true immediately; however, individuals should nevertheless need to build intensions out of the assertions, and they need to establish a set of goals in order to infer a conclusion.

One-model inferences are those in which the conclusion holds in the initial model of the premise, and so individuals can rapidly determine that an assertion is possible. For example:

All the aldermen are barbers.
Is it possible that some of the barbers are aldermen?

Reasoners have to construct intensions, use them to build a model, and to evaluate the truth of the conclusion in the model.

Finally, multiple-model inferences are those in which the conclusion holds in an alternative model of the premise. For instance:

All of the aldermen are barbers.
Is it possible that some of the barbers are not aldermen?

The theory predicts that zero-model inferences should be easier than one-model inferences, and one-model inferences should be easier than multiple-model inferences. Likewise, the computational model predicts that individuals should respond faster to zero-model than one-model than multiple-model inferences. The predictions are unique to mReasoner and the model theory, because the theory proposes that one of the most important factors in judging the relative difficulty of different inferences is the number of models people have to construct. We ran an experiment to test these two rank-order predictions.

Experiment

A typical trial in the experiment was:

All the artists are bakers.
Is it possible that all of the bakers are artists?

The experiment examined all 32 possible sorts of problem, but we focused our analysis on only the 22 logically valid inferences. The inferences comprise 4 zero-model problems, 12 one-model problems, and 6 multiple-model problems.

Method

Participants. 26 participants completed the study for monetary compensation on Mechanical Turk, an online platform hosted by Amazon.com. None of the participants

reported having had any training in logic, and they were all native speakers of English.

Design and materials. The participants carried out all 32 problems (4 sorts of premise x 8 sorts of conclusion), and they responded "yes" or "no" to a conclusion about a possible conclusion to each problem. The contents of the problems were based on nouns referring to common vocations. We devised a list of 32 pairs of such vocations, which we assigned at random to the problems to make two separate lists. The problems were presented to each participant in a different random order.

Procedure. The study was administered using an interface written in PHP, Javascript, and HTML. On each trial, participants read the premise, and, when ready, they pressed a button marked "Next", which replaced the premise with a question concerning the immediate inference, e.g., "Is it possible that all the bakers are artists?" They responded by pressing one of two buttons labeled, "Yes, it's possible" and "No, it's impossible". The program recorded whether or not their response was correct, and its latency (in s). The instructions stated that the task was to respond to questions about a series of assertions concerning what was possible given the truth of the assertion. The participants carried out three practice trials in order to familiarize themselves with the task before they proceeded to the experiment proper.

Results and discussion

The results corroborated the theory's predictions of difficulty, and they yielded the following trend: reasoners were 98% correct for zero-model problems, 84% correct for one-model problems, and 70% correct for multiple-model problems (Page's trend test, $L = 340.0$, $z = 3.88$, $p < .0001$). mReasoner predicted the participants' accuracy well, $R^2 = .98$.

The mean latencies also corroborated the predicted trend: 4.30 s for zero-model problems, 5.17 s for one-model problems, and 5.41 s for multiple-model problems (Page's trend test, $L = 336.0$, $z = 3.33$, $p < .0005$). mReasoner's predictions of accuracy likewise explained a significant portion of the latency variance, $R^2 = .76$.

We found a good fit between mReasoner's predictions and the data from the 22 individual problems, where any significant correlation suggests a good fit. The number of models correlated with participants' accuracy, $R^2 = .36$. And the system's latency predictions adequately fit the latencies across the problems, $R^2 = .24$. The fit could be improved further, however, and we suggest several ways of proceeding in the General Discussion.

General Discussion

The computational theory, mReasoner, simulates the construction of mental models in order to draw immediate inferences from singly-quantified premises. The theory uniquely predicts that individuals should be faster and more accurate when an inference can be drawn from an identity in

intensions. They should be next fastest and accurate when an inference can be drawn from the initial model constructed in system 1. And they should be slowest and least accurate when an inference can be drawn only from the discovery of an alternative model constructed in system 2. The predictions are a result of assumption of the mental model theory of reasoning: the more models you need, the more difficult a problem becomes. These rank-order predictions were borne out in the data from an experiment that tested all 22 valid inferences about possible conclusions in the set of 32 inferences. The immediate inferences were easy: people answered correctly 83% of the time. Yet they also revealed subtle differences in the difficulty between the various sorts of problems along the lines predicted by the theory.

The system we describe is limited, however, and it can be improved to yield a more fine-grained processing account of the data. We suggest two separate ways of proceeding. One way to improve the fit of the system is to make the system sensitive to the direction in which it scans models. For instance, if reasoners read a particular premise, e.g., *all artists are bohemians*, they may be biased to scan the model in the opposite directions by considering bohemians before artists. This *figural* bias is widely documented in syllogistic reasoning (Khemlani & Johnson-Laird, in press a) and it is likely to make a difference when reasoning about immediate inferences as well.

mReasoner could also place differential costs on the underlying processes of each of its three systems. The linguistic system, system 1, and system 2 are groups of interrelated processes, and for simplicity, mReasoner treats each of the processes as though it should place the same temporal cost on the inference as a whole. The processes are likely to place different costs on the system, however, and future versions of the theory might investigate the low-level mechanisms that give rise to such costs (cf. Khemlani & Trafton, under review).

Immediate inferences have been restricted to the study of syllogistic assertions, e.g., those that make use of the determiners *all*, *some*, and *none*. However, one major advance of the system we describe is that it can be used to make predictions if a broader range of inferences. Consider the following inference:

Most politicians are wealthy.

Is it possible that most wealthy people are not politicians?

The answer, like many of the inferences above, is easy: given the first premise, it is indeed possible that most wealthy people are not politicians. However, the inference likely engaged a search for alternative models, and so it should take reasoners longer to make it than a problem in which the conclusion follows straight from the premise:

Most politicians are wealthy.

Is it possible that some wealthy people are politicians?

Indeed, the former inference might have required a little thought, whereas the latter one may have felt “obvious.” A viable theory of immediate inferences should be able to account for any difficulty between the two problems, and at present, mReasoner is the only system that can do so. In the same vein, the system can be used to make predictions about problems that make use of statements such as:

Neither of the Xs is a Y

At most five of the Xs are Ys

More than a third of the Xs are Ys

Five of the Xs are not Ys

and so it is more general than theories restricted to Aristotelian syllogisms.

In sum, mReasoner is a new, unified computational cognitive model of deductive reasoning. It is based on the mental model theory of human reasoning, and so its primary prediction is that problems that require multiple models are difficult and take longer than problems that require only one model.

Acknowledgments

This research was funded by a National Research Council Research Associateship awarded to SK; ONR Grant #s N0001412WX30002 and N0001411WX20516 to JGT; and NSF Grant No. SES 0844851 awarded to P.J.L. We are also grateful to Len Breslow, Magda Bugajska, Hua Gao, Tony Harrison, Cathy Haught, Laura Hiatt, Gorka Navarrete, Marco Ragni, and Tobias Sonntag for their helpful comments. The views and conclusions contained in this document do not represent the official policies of the US Navy.

References

- Begg, I., & Harris, G. (1982). On the interpretation of syllogisms. *Journal of Verbal Learning and Verbal Behavior*, 21, 595–620.
- Braine, M.D.S., & O'Brien, D.P., Eds. (1998). *Mental logic*. Mahwah, NJ: Erlbaum.
- Bucciarelli, M., & Johnson-Laird, P. N. (1999). Strategies in syllogistic reasoning. *Cognitive Science*, 23, 247-303.
- Ceraso, J., & Provitera, A. (1971). Sources of error in syllogistic reasoning. *Cognitive Psychology*, 2, 400–410.
- Cohen, M.R., & Nagel, E. (1934). *An introduction to logic and scientific method*. London: Routledge & Kegan Paul.
- Erickson, J.R. (1974). A set analysis theory of behavior in formal syllogistic reasoning tasks. In R. Solso (Ed.), *Loyola symposium on cognition, Vol. 2* (pp. 305–330). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Evans, J.St.B.T. (2003). In two minds: Dual process accounts of reasoning. *Trends in Cognitive Sciences*, 7, 454–459.
- Evans, J. St. B. T. (2007). *Hypothetical thinking: Dual processes in reasoning and judgement*. Hove, UK: Psychology Press.

- Evans, J. St. B. T. (2008). Dual-processing accounts of reasoning, judgment and social cognition. *Annual Review of Psychology*, *59*, 255-278.
- Ford, M. (1995). Two modes of mental representation and problem solution in syllogistic reasoning. *Cognition*, *54*, 1-71.
- Geurts, B. (2003). Reasoning with quantifiers. *Cognition*, *86*, 223-251.
- Guyote, M.J., & Sternberg, R.J. (1981). A transitive-chain theory of syllogistic reasoning. *Cognitive Psychology*, *13*, 461-525.
- Hopcroft, J.E., & Ullman, J.D. (1979). *Formal Languages and Their Relation to Automata*. Reading, MA: Addison-Wesley.
- Jeffrey, R. (1981). *Formal logic: Its scope and limits* (2nd Ed). New York: McGraw-Hill.
- Johnson-Laird, P.N. (1983). *Mental Models*. Cambridge: Cambridge University Press. Cambridge, MA: Harvard University Press.
- Johnson-Laird, P. N. (2006). *How we reason*. Oxford, UK: Oxford University Press.
- Johnson-Laird, P.N., & Byrne, R.M.J. (1991). *Deduction*. Hillsdale, NJ: Erlbaum.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Strauss, Giroux.
- Khemlani, S., & Johnson-Laird, P.N. (2009). Disjunctive illusory inferences and how to eliminate them. *Memory & Cognition*, *37*, 615-623.
- Khemlani, S., & Johnson-Laird, P.N. (in press a). Theories of the syllogism: A meta-analysis. *Psychological Bulletin*.
- Khemlani, S., & Johnson-Laird, P.N. (in press b). The processes of inference. Manuscript under submission. *Argument and Computation*.
- Khemlani, S., Lotstein, M., & Johnson-Laird (under review). A unified theory of syllogistic reasoning. Manuscript under submission.
- Khemlani, S., & Trafton, J.G. (under review). mReactr: A computational theory of deductive reasoning. Manuscript under review.
- Newell, A. (1981). Reasoning, problem solving and decision processes: the problem space as a fundamental category. In R. Nickerson (Ed.), *Attention and Performance*, Vol. 8 (pp. 693-718). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Newstead, S.E., & Griggs, R.A. (1983). Drawing inferences from quantified statements: A study of the square of opposition. *Journal of Verbal Learning and Verbal Behavior*, *22*, 535-546.
- Peirce, C.S. (1931-1958). *Collected papers of Charles Sanders Peirce*. 8 vols. Hartshorne, C., Weiss, P., & Burks, A. (Eds.) Cambridge, MA: Harvard University Press.
- Politzer, G., van der Henst, J.B., Luche, C.D., & Noveck, I.A. (2006). The interpretation of classically quantified sentences: A set-theoretic approach. *Cognitive Science*, *30*, 691-723.
- Polk, T.A., & Newell, A. (1995). Deduction as verbal reasoning. *Psychological Review*, *102*, 533-566.
- Sloman, S.A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*, 3-22.
- Stanovich, K.E. (1999). *Who is rational? Studies of individual differences in reasoning*. Mahwah, NJ: Erlbaum.
- Stenning, K., & Yule, P. (1997). Image and language in human reasoning: A syllogistic illustration. *Cognitive Psychology*, *34*, 109-159.
- Verschueren, N., Schaeken, W., & d'Ydewalle, G. (2005). A dual-process specification of causal conditional reasoning. *Thinking & Reasoning*, *11*, 278-293.
- Wilkins, M. C. (1928). The effect of changed material on the ability to do formal syllogistic reasoning. *Archives of Psychology*, *16*, No. 102.