A framework for modeling evidence-based, context-influenced reasoning

Keith Devlin
Stanford University

Abstract

The continuing process of improving intelligence product requires attention not only to collection but to methods of analysis as well. The analysis of broad-based movements, complex organizations, or simply the use of large data bases often make use of quantitative tools. We present a framework for representing, and an associated method for analyzing, evidence-based, context-influenced human reasoning. Our specific interest is intelligence analysis but the framework we present should be applicable to a broad range of analytical uses and to decision-making as well. The approach used in this paper is analytically formal but there are important differences between classical formal logic and our framework. Our primary scientific purpose in developing our model is to provide increased understanding of analytic reasoning, but possible other applications of our framework are:

1. the formulation of reasoning and reasoning-report protocols and methodologies to improve evidence-based analysis and decision making;
2. the development of computer reasoning tools to assist knowledge workers;
3. a basis for teaching and developing better reasoning skills.

Consequently, this paper is relevant not only to analysts in the production of actionable product but to consumers as well. We believe that analysts, commanders on the ground, and policy decision makers can all benefit from this discussion. The benefits for trainers, teachers, and students seems self-evident.

This paper is dedicated to my former colleague and good friend, the logician Kenneth Jon Barwise (1942–2000). The work presented here is very much in the spirit of his approach to logic, a theme I pick up in my closing remarks.

1 Introduction

Richards J. Heuer, Jr., in his classic book *Psychology of Intelligence Analysis* (re-published by the Central Intelligence Agency in 1999), writes:

> *"To the leaders and managers of intelligence who seek an improved intelligence product, these findings offer a reminder that this goal can be achieved by improving analysis as well as collection. There appear to be inherent practical limits on how much can be gained by efforts to improve collection. By contrast, an open and fertile field exists for imaginative efforts to improve analysis.*
>
> *These efforts should focus on improving the mental models employed by analysts to interpret information and the analytical processes used to evaluate it. While this will be difficult to achieve, it is so critical to effective intelligence analysis that even small improvements could have large benefits."*[1]

Surprisingly, although many organizations have examined the information sources on which they base crucial decisions, few have devoted much time and effort to a study of the process whereby they use that information to make decisions. The CIA is one of the few that have. They had good reason to promote this work. Heuer, himself an intelligence analyst for many years, also writes:

> *"Major intelligence failures are usually caused by failures of analysis, not failures of collection."* [2]

In this paper we develop a mathematical framework — a model — in which to represent, and then study, evidence-based, context-influence reasoning, of which intelligence analysis is a prime example. How can such a mathematical model be of benefit to analysts? We see three principal benefits, depending on how the mathematics is viewed:

Benefit 1. Gaining a quick perspective. The mathematics is viewed as a mathematical model. In this case, the mathematical framework we present — the model — provides a simple, top-level overview of what is involved in analysis and decision making. While we do not suggest that any analyst should seek to follow the model in a step-by-step fashion, we believe that familiarity with the model provides a valuable, theoretical understanding of and appreciation for the reasoning process that will improve performance. While any analyst will benefit from having a "God's eye view" of the analytic process that a mathematical model provides, this benefit is particularly strong for experienced analysts, who are able to interpret the model in terms of their own experience, and thereby reflect on that experience. As Heuer observes:

> *"Intelligence analysts should be self-conscious about their reasoning process. They should think about how they make judgments and reach conclusions, not just about the judgments and conclusions themselves."* [3]

To maximize this benefit, we need to identify what "interpretation" (of the model) means. To do this, we have to develop a method for analyzing human reasoning strategies.

Benefit 2. Flexibility and practice. The mathematics is viewed as a theory, i.e., as a theoretical analysis of analytic reasoning. In this case, novice analysts may gain considerable benefit both from studying the theory and from tackling simulations of intelligence analysis based on the theory.[4]To optimize this benefit, we would need to develop a set of technologies for stimulating creative and innovative applications of the theory in representative samples of real-life analyses.

Benefit 3. Avoiding stovepipes. The mathematics is viewed as a representational framework and an associated analytic process. In this case, our mathematics provides a valuable means for analyzing past analyses and learning from them. Good analysts constantly reflect on their past experiences, particularly those where the outcome has been less than optimal, and they share the results of their analyses with colleagues, so that experience is realized as shared memory. Identifying the step or steps in an analysis process where things went wrong can be an extremely difficult task. Human beings have a strong tendency to create "mental tunnels" that will repeatedly lead them astray. (See [17].) Breaking out of such mental traps to be able to identify the problem steps requires adopting a systematic reflection process, one that makes problematic that which was previously assumed to be unproblematic and obvious.

To be useful in all three of the above ways, but particularly for benefit 3, the mathematical framework needs to have two features not found in the majority of applications of mathematics:

1. The mathematics needs to be simple and as far as possible self-explanatory.
2. The mathematical framework needs to be sufficiently flexible to express (i.e., represent) the key features of an analytic process, without distorting those features.

In connection with condition 1, we note that most analysts do not have a background in mathematics or in a mathematics-based discipline such as science or engineering, and may even be afraid of or averse to mathematical analyses. Even analysts who are comfortable with mathematics may not have much time to devote to the acquisition of new techniques, especially if they are yet to be convinced of their value to what they do.

The second condition here is particularly important. Mathematics gains much of its strength by virtue of its simplistic approach; for example, complex entities become points in a space, complex relationships become surfaces or manifolds, etc. But this process of simplification can lead to a state of affairs where the mathematical results tell you a lot about the mathematical objects being studied and relatively little about the real-world entitites they were intended to model. The analysis beomes largely theory internal, that is, the mathematical properties of the model become more significant than the real world properties of what they were intended to model. It is a particularly prevalent danger with mathematical analyses because mathematicians, by inclination, training, and the culture of their discipline, constantly seek simplifying assumptions that lead to "better" and "more elegant" mathematics. An important virtue of the use of mathematics here is that *the assumptions must be stated and stated precisely.*

A third condition that our mathematical framework ought to have is that it should be possible to view existing mathematical models of reasoning — insofar as they do not violate condition 2 above excessively — as special cases of ours. While meeting this condition is not essential to the efficacy of our methods, it would certainly be advantageous if the mathematical framework we develop can be viewed as a generalization, or relaxation, of existing methods that have proved their worth in certain circumstances.

2 Beyond logic

Classical logic provides the following conception of logical reasoning: A *logical proof* consists of a finite sequence $\sigma_1, \sigma_2, \ldots, \sigma_n$ of statements, such that for each $i = 1, \ldots, n$, σ_i is either an assumption for the argument (possibly an axiom), or else follows from one or more of $\sigma_1, \ldots, \sigma_{i-1}$ by a rule of logic.

The importance of formal logic in mathematics is not that mathematicians write proofs in the system. To do so would in general be far too cumbersome. Rather, formal logic provides a framework for analyzing the notion of mathematical proof. This has led to several benefits. One is a deeper understanding of mathematical proof. Another is the development of techniques for proving that certain statements are in fact not provable. A third is the development of computer tools to carry out automated proof procedures and to assist the human user construct proofs. Still another benefit is that the study of formal logic has educational value for the apprentice mathematician.

Similarly, we do not propose that the formal calculus of reasoning we develop here be used by the human reasoner. As is the case with formal logic and mathematics, it would be far too cumbersome to try to conduct reasoning using our framework. But, just as a study of logic provides a beneficial component of a mathematical education, so too we believe that study of our calculus will lead to a better understanding of evidence-based reasoning — and perhaps as a result to better performance. In particular, we believe that *a study of the framework we develop offers educational and performance benefits to both expert and novice reasoners*.

One major difference between our framework and formal logic is that the latter is a precise mathematical theory, for which it is possible to write down axioms and prove theorems. That works with a framework developed to study idealized mathematical reasoning because mathematics itself is precise, axiomatizable, and amenable to theorems — indeed, theorems lie at the heart of mathematics. Evidence-based reasoning about real life situations is very different. This is why so many attempts to develop AI systems failed to meet their ambitious objectives.

Three important specific products we see coming out of this project are:

(i) the formulation of reasoning protocols and methodologies to improve analysis in various business, commercial, political and military domains;

(ii) the development of computer reasoning tools to assist analysts working in such areas;

(iii) a basis for teaching and developing reasoning skills.

However, none of these three possible applications constitutes our primary motivation, which we see as prior to all three. Namely, just as formal logic was developed (initially by the ancient Greeks) with the primary goal of understanding (context-free) logical reasoning,[5]so too we seek to understand the underlying logical structure of evidence-based human reasoning that takes place in a specific real-life context. This brings up a general issue about mathematical modeling that we feel needs to be addressed before we proceed further.

Although mathematicians never formulate their arguments completely in the language of formal logic, which is what would be required to adhere to the strictest form of this conception of proof, the formal definition nevertheless does capture the essense of real mathematical proofs. Hence the classical definition of proof (given above) does provide a good model for what constitutes logical reasoning *within mathematics*. But the very efficacy and success of the model within mathematics can lead — and has done so — to an expectation that it may be applied to logical reasoning outside mathematics. And this is where problems can arise. In our view, familiarity with the classical model of mathematical proof is a potentially dangerous thing to bring to other forms of reasoning in nonmathematical domains.

Having a theoretical, idealized model of an activity at the back of our minds is a powerful asset. Human intelligence depends in large part on our ability to formulate and utilize such models. Simplified models help us cope with the often otherwise unmanageable complexities of real life situations. Acquisition of the ability of the human mind to construct such models constituted perhaps the most pivotal step in evolutionary human development.[6]The individual who trades in evidence-based reasoning about real life domains needs to have a good abstract model of that activity at the back of his or her mind just as much as the mathematician or any other specialist.

However, classical logic is not only an inappropriate model, it is *misleadingly* inappropriate, and hence should be replaced by something else.

To give just one instance of how pervasive background metaphors can be, consider the study carried out at the Harvard University Commencement a few years ago, when members of the graduating senior class and some of their professors were asked, among other things, to explain why Boston is warmer in the summer. The majority replied that this was because the Earth is nearer the Sun in the summer. Because this is so obviously false — if true, it would imply that Australia is also warmer in July, whereas it is in the middle of the Australian winter — this popular response was touted by many commentators as a sign that even a Harvard education is sorely lacking in basic science. Now, that may be a valid conclusion to draw, but there is something far more interesting going on here. Why did so many obviously bright and "well educated" individuals give *the same* wrong answer? Especially since, when prompted by the questioner, many of them eventually did produce the correct explanation, namely that the seasonal temperature change is a consequence of the changing angle between the Earth's axis of rotation and the Sun as the Earth orbits the Sun. The answer is that the respondents were basing their explanation — which was often immediate and given with confidence — on a familiar pattern (or *model*) they had acquired through their early childhood experiences: namely, when you have a heat source, nearer means hotter, further away means colder. This is an extremely important rule that we all learn, sometimes painfully, at an early age. As a model, we use it frequently whenever there is a heat source. In the case of the Harvard graduates, however, the model led them astray.

Likewise, the model for a logical argument provided by the classic definition of a mathematical proof can lead even highly intelligent, well-educated people astray. We suggest that the model of evidence-based, contextual reasoning we develop in this paper provides a far more reliable background metaphor on which to base an intuitive sense for such reasoning.

We should perhaps add that we are aware that, with some effort, it is possible to represent almost any instance of real-life, context-influenced, evidence-based reasoning in classical mathematical logic. But, while doing so may and can have positive consequences, it does little to help us understand the structure of that reasoning; indeed, it usually obscures it. To be a useful mental aid, an abstract model needs to represent the domain being modeled in a way that is faithful to the key issues ('key" here meaning the issues the model is designed to represent).

3 Modeling real-life reasoning

Some characteristic features of real-life evidence-based reasoning that we seek to represent in our model (all of which are ignored by classical logic) are:

1. It is not always linear.

2. It is often holistic.

3. The information on which the reasoning is based is often not known to be true. The reasoner must, as far as possible, ascertain and remember the source of the evidentiary information used and maintain an estimation of its likelihood of being reliable.[7]

4. Reasoning often involves searching for information to support a particular step. This may involve looking deeper at an existing source or searching for an alternative source.

5. Reasoners often have to make decisions based on incomplete information.

6. Reasoners sometimes encounter and must decide between conflicting information.

7. Reasoning often involves the formulation of a hypothesis followed by a search for information that either confirms or denies that hypothesis.

8. Reasoning often requires backtracking and examining your assumptions.

9. Reasoners often make unconscious use of tacit knowledge, which they may be unable to articulate.

The factors we have just outlined imply that real-life, evidence-based reasoning is rarely about establishing "the truth" about some state of affairs. Rather it is about marshalling evidence to arrive at a conclusion. If the reasoner wants to attach a reliable degree of confidence to the conclusion, she or he must keep track of the sources of all the evidence used, the nature and reliability of those sources, and the reliability of the reasoning steps used in the process.

The model we develop is it not intended to "mathematicize" human reasoning. Indeed, we believe that for all the factors just listed,[8]human reasoning cannot be captured in a mathematical model. Rather, our purpose in developing our model is to try to shine some light on the logical (*sic*) structure of human reasoning. This is what mathematics is supremely good for. The fact that on many occasions it is possible to take a mathematical model and implement it in some way should not seduce us into thinking that that is the only way a mathematical model may be of use. Simple (or relatively simple) mathematical models are above all useful because they provide a means to help us understand a given domain or phenomenon. They do so in large part because they are clean, lean, and simple (when they are), and the degree to which they possess those three characteristics generally correlates well to the effectiveness of the model in assisting our understanding. Complex mathematical models, while often of great value, are rarely a useful aid to *understanding* except for specialists who study and work with the model at length, such as physicists whose current understanding of the nature of the universe and of the matter in it depends on some highly complex mathematical models.

In order to meet the goals just outlined, our model should capture certain key features of the reasoning *process*, not just the result of that process. Moreover, in order to provide an analyst with a tool to understand a specific reasoning process, which it will have to do if it is to provide more than a purely theoretical understanding (not that this is without value, but we are going for more than that) the model should provide a means of "problematizing" various aspects of the process.

In developing a mathematical model of reasoning that takes into account the the features outlined above, we adopt a charitable view of the way analysts arrive at a decision, in that we assume they do the best job possible under the prevailing circumstances. This is often not the case. Alexander George[9]has identified a number of less-than-optimal common strategies people adopt in making decisions in the face of incomplete information and multiple, competing values and goals:

- Satisficing: Selecting the first identified alternative that appears "good enough," rather than examining all alternatives to determine which is "best."

- Incrementalism: Focusing on a narrow range of alternatives representing marginal change, without considering the need for dramatic change from an existing position.

- Consensus: Opting for the alternative that will elicit the greatest agreement and support. Simply telling the boss what he or she wants to hear is one version of this.

- Reasoning by analogy: Choosing the alternative that appears most likely to avoid some previous error or to duplicate a previous success.

- Relying on a set of principles or maxims that distinguish a “good” from a “bad” alternative.

Massimo Piatelli-Palmerini [17] likewise described a number of human tendencies to follow “mental tunnels” that lead to erroneous or less-than-optimal conclusions.

Our model may be applied to any of these reasoning strategies, by replacing one or more of the reasoning rules we describe below by one or more of George’s or Piatelli-Palmerini’s less-than-optimal rules. Thus, one possible application of our framework is to highlight different reasoning strategies and compare the good from the bad — or perhaps we should say the better from the worse.

It should be noted that the framework we develop below, while it has explicit representations of context, is not an attempt to develop yet another “logic of context.” The idea of trying to handle context in formal logic was first introduced as a research thread in artificial intelligence by John McCarthy in his 1971 Turing Award lecture, subsequently published as [15]. See also McCarthy’s more recent article on the subject [16]. Many other attempts have been made to incorporate features of context in formal logic, for example Attardi and Simi [1], Buvac and Mason [3], Farquhar and Buvac [9], Farquhar, Dappert, Fikes, and Pratt [10], Giunchiglia [12], Guha [13], and Shoham [18]. The microtheories employed by the Cyc system are also clearly an attempt to take account of context. In addition, Rich Thomason [19] is engaged in an ongoing project to formalize context.

Our approach is quite different. While we recognize the utility of computer reasoning systems that can reason in specific, defined “contexts” — we would prefer to call them *local domains* — we firmly believe that the influence of context cannot be fully captured this way. We articulate the evidence for this belief in [5] and in our forthcoming paper [7]. As we argue in [7], when a trained human analyst, working with one or more computer information systems, reasons about real world affairs in a domain in which she or he is expert, the human analyst is in general far better able to make key decisions about matters of context than the automated system. Thus, rather than reify contexts and integrate them into a formal system, we set out to build a framework that has two distinct parts, a formal system and a context-tracker. The contexts that the context-tracker represents are not reifications as mathematical objects but pointers to real-world contexts, with all that entails. Our framework sets out to capture, as far as possible, the linkages between, and influences of, context and the elements of the formal system. As a model of an activity, our “logic” attempts to model *the human reasoner* — perhaps aided by one or more information systems, possibly including automatic reasoners. It is not intended to form a blueprint for the design of an automated system. It may, of course, *inform* the design of such a system, in which case we would expect that system to make use of some of the “context logics” developed by the researchers referred to above. Our framework may also form a basis for various reasoning and reporting protocols to be followed by analysts. We investigate these issues more fully in [7].

4 The framework

Our framework views evidence-based reasoning as a temporal cognitive process that acts not on statements σ (as in the model of a mathematical proof) but on entities of the form

$$s \models_{\tau_1,\tau_2,\ldots} \sigma$$

where:

1. σ is a statement (or fact);

2. s is a *situation* (in the sense developed at CSLI during the period 1983–93 and described in Devlin [4]) which provides support or context of origin for σ; and

3. τ_1, τ_2, ... are the *indicators*[10] of σ, i.e., the specific items of information in s that the reasoner takes as justification of σ.

We call an entity of the form $s \models_{\tau_1,\tau_2,\ldots} \sigma$ a *basic reasoning element*.

Within our framework, a process of evidence-based reasoning to decide an issue $\mathcal{I}$ can be represented like this:

$$\begin{array}{ccc} & \mathcal{I} & \\ \hline s_1 & \models_{\tau_1,\ldots} & \sigma_1 \\ s_2 & \models_{\tau_2,\ldots} & \sigma_2 \\ s_3 & \models_{\tau_3,\ldots} & \sigma_3 \\ & \vdots & \\ \hline s & \models_{\tau_1,\ldots,\tau_2,\ldots,\tau_3,\ldots} & \sigma \end{array}$$

where each basic reasoning element either supplies evidence for the reasoning or else follows from one or more previous elements by a logical deduction rule.

Analogous to the concept of a mathematical proof (sequence), we define (subject to some technical modifications) an *evidential reasoning process* as a finite sequence ρ_1, ρ_2, ..., ρ_n of entities of the above form such that each ρ_i is either *evidential* (i.e., an input to the reasoning process) or else the result of applying some logical rule of evidence-based reasoning to one or more of ρ_1, ..., ρ_{i-1}. Here is the formal development of this notion.

By an *evidential reasoning element* we mean a 1×3 matrix of the form

FACT	SUPPORT	INDIC(1), INDIC(2), ...

such that

$$\text{SUPPORT} \models_{\text{INDIC}(1),\,\text{INDIC}(2),\,\ldots} \text{FACT}$$

By an *evidential reasoning step* we shall mean a finitary array of the form

OPERATOR	FACT_1	SUPPORT_1	$\text{INDIC}_1(1), \text{INDIC}_1(2), \ldots$
	FACT_2	SUPPORT_2	$\text{INDIC}_2(1), \text{INDIC}_2(2), \ldots$
	...		
	FACT_k	SUPPORT_k	$\text{INDIC}_k(1), \text{INDIC}_k(2), \ldots$
OUTPUT	FACT_{k+1}	SUPPORT_{k+1}	$\text{INDIC}_{k+1}(1), \text{INDIC}_{k+1}(2), \ldots$

where each row

FACT_i	SUPPORT_i	$\text{INDIC}_i(1), \text{INDIC}_i(2), \ldots$

is an evidential reasoning element. The index k depends on the operator OPERATOR, and is called the *arity* of the operator.

The idea is that a basic evidential reasoning step consists of the application of the logical operator to one or more consituents of the evidential reasoning elements in its *scope* (the first k elements listed) to produce the *output element* in the final row.

An *evidential reasoning process* is a finite sequence $\rho_1, \ldots \rho_n$ of basic reasoning steps such that each element is either *evidential* (i.e., an input to the reasoning process) or else the output of some previous (in the sequence) evidential reasoning step, or else is the special element STOP, which is the final element in the process. (STOP is a failure condition; we describe it later.)

The sequence of elements in an evidential reasoning process are not intended to provide a temporal model of the actual steps carried out by a reasoner. Rather, an evidential reasoning process models the logical flow of the reasoning as it leads to the conclusion. As we mentioned earlier, much real-life reasoning is not linear. However, our model is such that any linear progression of steps in the actual reasoning a human carries out will be mapped to a linear ordering of the corresponding basic reasoning elements in the model.

The actual operators that arise in any particular instance of evidence-based reasoning will depend on the specific circumstances that pertain in that application. In this document we simply indicate the general form of some of the more generic operations that are likely to be used in any instance.

For example, among the operators are some that correspond to classical logic. Since classical logic ignores context, we have to exercise care in porting classical logic operators into our calculus. This means that our rules all have restrictions on when they may be applied. We start with the following two rules, each of which involves a binary reasoning operator:

Evidential Conjunction Rule

CONJOIN	σ	s	$\tau_1, \tau_2, \ldots$
	θ	t	$\gamma_1, \gamma_2, \ldots$
OUTPUT	$\sigma \wedge \theta$	$s \cup t \cup \{\delta\}$	$\delta, \tau_1, \tau_2, \ldots, \gamma_1, \gamma_2, \ldots$

where $\delta = \mathrm{Con}\{\tau_1, \tau_2, \ldots, \gamma_1, \gamma_2, \ldots\}$, the assertion that the set $\{\tau_1, \tau_2, \ldots, \gamma_1, \gamma_2, \ldots\}$ is logically consistent (i.e., has no internal contradictions), and where the rule may be applied only if δ is valid. The restriction that δ is called the *indicator consistency condition* for the rule. If this condition is not satisfied, the rule produces the output STOP. (We consider later what happens when the STOP element is generated.)

Evidential Modus Ponens Rule

MP	σ	s	$\tau_1, \tau_2, \ldots$
	$\sigma \rightarrow \theta$	t	$\gamma_1, \gamma_2, \ldots$
OUTPUT	θ	$s \cup t \cup \{\delta\}$	$\delta, \tau_1, \tau_2, \ldots, \gamma_1, \gamma_2, \ldots$

where $\delta = \mathrm{Con}\{\tau_1, \tau_2, \ldots, \gamma_1, \gamma_2, \ldots\}$, and where the rule may be applied only if δ. If this condition is not satisfied, the rule produces the output STOP.

We need to exercise care in using these two rules. If the supports s and t are identical, there is in general no problem, nor if one support is contained within the other. In either of these cases, the indicator consistency condition can generally be assumed to be automatically satisfied, since reasoning generally proceeds under the tacit assumption that each individual source is internally consistent. (If, however, the reasoner suspects — or comes to suspect — that one of the supports used in the reasoning is internally inconsistent, then resolving that inconsistency becomes part of the reasoning process. This is a particular case of the following general observation concerning evidence-based reasoning.)

The idea behind our approach is this. Coupling a fact σ with its support s in our framework does two things: (i) it acknowledges that σ does come from a particular source, and (ii) it provides a record of that source. Explicitly listing the indicators

$\tau_1, \tau_2, \ldots$ with σ and s puts on record the particular items of information in s that the reasoner believes are salient in supporting σ, and uses to justify making use of σ in the reasoning. When an unexpected or troublesome conclusion is reached, or when the reasoning fails to yield a conclusion, it may be necessary to re-examine the veracity of some of the facts used in the reasoning, and that may involve reconsideration of the indicator already identified, or a search for indicators hitherto ignored. In an extreme case, the reasoner may have to question an entire source, perhaps rejecting it and looking for evidence elsewhere.

There are two unary reasoning operators associated with the indicators in a reasoning element: EVAL-INDIC, which checks the indicators already identified for veracity, and FACTORIZE, which identifies new items of information in the support that are salient to the use of the fact in the reasoning process. The rules associated with these operators are:

Indicator Evaluation Rule

EVAL-INDIC	σ	s	$\tau_1, \tau_2, \ldots$
OUTPUT::	σ	s	$\tau_1, \tau_2, \ldots$
	STOP		

where the notation here (note the double-colon after OUTPUT) indicates that the output of the rule is exactly one of the two elements

σ	s	$\tau_1, \tau_2, \ldots$

and

STOP

The former output is obtained if the evaluation of τ_1, τ_2, $\ldots$ affirms their veracity; the output is STOP if the evaluation determines that one of these indicators is in fact not valid, or at least is in doubt.

Thus, the evidential reasoning step generated by an application of the Indicator Evaluation Rule is of one of the two forms:

EVAL-INDIC	σ	s	$\tau_1, \tau_2, \ldots$
OUTPUT	σ	s	$\tau_1, \tau_2, \ldots$

EVAL-INDIC	σ	s	$\tau_1, \tau_2, \ldots$
OUTPUT	STOP		

Indicators Extension Rule

EXTEND-INDICS	σ	s	$\tau_1, \tau_2, \ldots$
OUTPUT	σ	s	$\tau_1, \tau_2, \ldots, \gamma_1, \gamma_2, \ldots$

where $\gamma_1, \gamma_2, \ldots \in s$.

This rule implies that

$$s \models_{\tau_1, \tau_2, \ldots, \gamma_1, \gamma_2, \ldots} \sigma$$

The intuition is that the reasoner identifies additional information (additional indicators) that she or he judges to contribute to the acceptance of the fact σ under consideration.

Use of the following rule, which involves the unary operator EVAL-SUPPORT, indicates a suspicion that the reasoning process has a serious flaw.

Support Evaluation Rule

EVAL-SUPPORT	σ	s	$\tau_1, \tau_2, \ldots$
OUTPUT::	σ	s	$\tau_1, \tau_2, \ldots$
	STOP		

The former output is obtained if the evaluation of s affirms its internal consistency and reliability; the output is STOP if the evaluation determines that s is inconsistent or unreliable, or at least that the consistency or reliability of s is in serious doubt.

Thus, the evidential reasoning step generated by an application of the Support Evaluation Rule is of one of the two forms:

EVAL-SUPPORT	σ	s	$\tau_1, \tau_2, \ldots$
OUTPUT	σ	s	$\tau_1, \tau_2, \ldots$

EVAL-SUPPORT	σ	s	$\tau_1, \tau_2, \ldots$
OUTPUT	STOP		

When a reasoning step produces the output STOP, the reasoner has to backtrack and examine the process so far. If it is not possible to make any changes to any previous steps, then the reasoning process breaks down. In such a case, the available information is either contradictory or else simply not adequate to resolve the target issue.

A common step in evidence-based reasoning is to decide between two or more different possibilities, which may or may not be mutually exclusive. The exact mechanism by which the comparison is made will vary from case to case, but functionally such an operation produces the following basic reasoning step:

Selection Rule

SELECT	σ_1	s_1	$\tau_1(1), \tau_1(2), \ldots$
	σ_2	s_2	$\tau_2(1), \tau_2(2), \ldots$
	$\ldots$		
	σ_n	s_n	$\tau_n(1), \tau_n(2), \ldots$
OUTPUT	σ_i	$s_i \cup s$	$\gamma, \delta, \tau_i(1), \tau_i(2), \ldots$

for some i, $1 \leq i \leq n$, where s is the very reasoning process the agent is carrying out (and which we are capturing with our calculus), $\gamma \in s$ is the fact that this particular selection has been made, and $\delta \in s$ is the criterion for making the selection.

Note that the output of a selection step carries a record of the selection having been made and of how it was made.

In practice, making a selection may involve examination of the supports and the indicators associated with the facts being compared, possibly leading to additional factorization for some facts or other operations. Such factorizations, or other steps, will be captured in our model by being represented explicitly as earlier steps in the process sequence.

Sometimes during the course of reasoning, the reasoner believes it is necessary to expand the scope of the domain from which particular facts were obtained, perhaps with a view to finding additional indicators to strengthen confidence in the fact or to replace the fact with a stronger version. This is captured by the following rules, often used in successively in conjunction, together with the indicators extension rule.

Support Expansion Rule

EXPAND-SUPPORT	σ	s	$\tau_1, \tau_2, \ldots$
OUTPUT	σ	s'	$\tau_1, \tau_2, \ldots$

where $s \subseteq s'$.

Strengthen Fact Rule

STRENGTHEN-FACT	σ	s	$\tau_1, \tau_2, \ldots$
OUTPUT	σ'	s	$\tau_1, \tau_2, \ldots$

where $s \models_{\tau_1,\tau_2,\ldots} \sigma' \to \sigma$.

Multiple Views Uniformization Rule

Reasoners sometimes view more than one data source in order to use their experience and tacit knowledge to synthesize a conclusion that may not follow directly from the different sources by logical reasoning. To capture such actions, we could add an operator that provides some form of merge or unification for simultaneous views of information from different sources. However, the evidential conjunction rule that we already have will handle many cases of multiple views of data.

In circumstances where two views of a data item σ can be regarded as providing two indicator sets for the same fact within the same context:

$$s \models_{\tau_1,\tau_2,\ldots} \sigma \quad \text{and} \quad s \models_{\gamma_1,\gamma_2,\ldots} \sigma$$

we can apply the following operator:

MV UNIF	σ	s	$\tau_1, \tau_2, \ldots$
	σ	s	$\gamma_1, \gamma_2, \ldots$
OUTPUT	σ	s	$\delta, \tau_1, \tau_2, \ldots, \gamma_1, \gamma_2, \ldots$

where δ is the fact that this unification has taken place.

Subtasking

Reasoners often need to break a particular task into subtasks. Typically, this entails defining a set of subtasks that together will complete the given task, and then working on each subtask in turn. Alternatively, the reasoner may decide to abandon (perhaps just for the time being) the current goal and concentrate solely on some subtask, which then becomes the new goal.

The framework as described so far can handle the individual steps in each subtask analysis, and can track choices of subtasks as localized reasoning contexts. But we have not introduced an operator for subtask selection or for breaking a task into a sufficient group of subtasks. Instead, we have left this as a meta-level operation. We did so in order to avoid making our technical machinery more complicated than it already is. Since our primary aim is to provide a framework to aid human reasoners, not a blueprint for an automated reasoning system, we feel this is a reasonable choice. But before moving on let's take a brief look at what would be required to modify our framework to incorporate subtasking.

Within our current framework, a process of evidence-based reasoning to decide an issue $\mathcal{I}$ is represented like this:

$$\begin{array}{ccc} & \mathcal{I} & \\ \hline s_1 & \models_{\tau_1,\ldots} & \sigma_1 \\ s_2 & \models_{\tau_2,\ldots} & \sigma_2 \\ s_3 & \models_{\tau_3,\ldots} & \sigma_3 \\ & \vdots & \\ \hline s & \models_{\tau_1,\ldots,\tau_2,\ldots,\tau_3,\ldots} & \sigma \end{array}$$

The issue $\mathcal{I}$ is kept constant throughout our development. In order to incorporate subtask selection, we could introduce a mechanism to represent the selection of a subtask $\mathcal{J}$ of $\mathcal{I}$ or else the division of $\mathcal{I}$ into a collection of subtasks $\mathcal{J}_1$, $\mathcal{J}_2$, $\ldots$, $\mathcal{J}_n$. The framework would need to keep track of the supports and the indicators, both when the subtask(s) is (are) selected and when the completion of all the tasks in a subdivision results in the completion of the original task. This is all very straightforward.

5 Some special cases

To get a sense of how our framework operates, we show how it applies to some familiar special cases or models for reasoning

Mathematical reasoning. First of all, let's take the case of mathematics, where $\sigma_1, \ldots, \sigma_n$ are statements about some mathematical structure $\mathcal{M}$, say a group or a field. We may assume that $\sigma_1, \ldots, \sigma_n$ are written in the first-order language for $\mathcal{M}$. In that case, each of the expressions $s_i \models \sigma_i$ denotes a standard proposition of classical Tarski-based model theory. In this case, by the Completeness Theorem of first-order predicate logic, $s = \mathcal{M}$ and the deduction takes the form

$$\begin{array}{c} \mathcal{I} \\ \hline \mathcal{M} \models \sigma_1 \\ \mathcal{M} \models \sigma_2 \\ \vdots \\ \mathcal{M} \models \sigma_n \\ \hline \mathcal{M} \models \sigma \end{array}$$

If this reasoning is valid, then we must have

$$\mathcal{I} = ![\mathcal{M} \models \sigma]?$$

where an expression of the form $!P?$ for some proposition P denotes the goal "Determine whether P true or false." That is, the goal is to determine whether or not σ is true of $\mathcal{M}$.

The completeness theorem also tells us that (if the deduction is valid), σ follows from $\sigma_1, \ldots, \sigma_n$ by the rules of logic alone.

Reasoning from a common source. Another special case is where all of the information $\sigma_1, \ldots, \sigma_n$ comes from the same source, $\mathcal{S}$. In this case the conclusion support s is also $\mathcal{S}$, and the deduction takes the form:

$$\frac{\mathcal{I}}{\dfrac{\begin{array}{c}\mathcal{S} \models \sigma_1 \\ \mathcal{S} \models \sigma_2 \\ \vdots \\ \mathcal{S} \models \sigma_n\end{array}}{\mathcal{S} \models \sigma}}$$

For a valid process, we must have

$$\mathcal{I} = \,!\,\sigma?$$

(Determine whether to do σ, or else determine whether σ is true.)

Bayesian inference. In some cases, knowledge of the source of each data item σ_i may be converted into a numerical probability of the reliability of σ_i, i.e. the probability that σ_i is true. In such a situition, we may be able to apply Bayes' Theorem repeatedly in order to obtain a conclusion σ and assign a probability to σ. In this case, the function F is a numerical function based upon Bayes' Theorem and the function H is an instance of Bayesian inference. This kind of reasoning is quite common, particularly in intelligence gathering.

We may represent a Bayesian reasoning process using the original notation

$$\frac{\mathcal{I}}{\dfrac{\begin{array}{c}s_1 \models \sigma_1 \\ s_2 \models \sigma_2 \\ \vdots \\ s_n \models \sigma_n\end{array}}{s \models \sigma}}$$

with the understanding that each of $s_1,$, $\ldots$, s_n, s is a number between 0 and 1 inclusive, and each expression $s_i \models \sigma_i$ should be interpreted as a probability statement $p(\sigma_i) = s_i$, and similarly for $s \models \sigma$.

6 Summary and discussion

The basis for our method is to view evidence-based reasoning as a temporal cognitive process that acts on entities of the form

$$s \models_{\tau_1,\tau_2,\ldots} \sigma$$

where σ is a statement (or fact), s is its support or context of origin, and τ_1, τ_2, ... are its indicators, the specific items of information in s that the reasoner takes as justification of σ.

We analyze evidence-based reasoning so described in terms of a number of basic reasoning steps, an illustrative example being the Evidential Modus Ponens Rule:

MP	σ	s	$\tau_1, \tau_2, \ldots$
	$\sigma \to \theta$	t	$\gamma_1, \gamma_2, \ldots$
OUTPUT	θ	$s \cup t \cup \{\delta\}$	$\delta, \tau_1, \tau_2, \ldots, \gamma_1, \gamma_2, \ldots$

where $\delta = \mathrm{Con}\{\tau_1, \tau_2, \ldots, \gamma_1, \gamma_2, \ldots\}$, and where the rule may be applied only if δ.

We list a number of such rules, but acknowledge that many applications will involve rules not listed here. Our framework is designed to allow for such additional rules to be incorporated.

Readers familiar with situation theory will have observed that our present framework amounts to making explicit in the model the features of the context situation — our *indicators* — that provide direct support for the items of information considered in the reasoning — what we call the *facts*. Moreover, we model (aspects of) the process of reasoning, not just the sequence of facts and their situational supports. By making this additional salient information explicit in the model, we can obtain a finer grained analysis than is possible in situation theory, that requires much less ad hocing when we carry out an analysis of a specific reasoning process. In our framework, the Evidential Modus Ponens rule performs the task that was handled by constraints in situation theory. Our decision to ignore much of the machinery for handling situation-theoretic constraints was based on pragmatic grounds, with a view to the kinds of reasoning we are attempting to model.

Although our primary goal is to develop a framework that aids understanding, we are aware that any enterprise such as ours has the potential of forming the basis for the specification of reasoning protocols or the design of reasoning support tools. The model we have developed would result in protocols or support tools that:

1. Force explicit identification and tracking of sources.
2. Force explicit identification and tracking of supporting information (the indicators).
3. Force regular reconsideration of the reasoning process itself.
4. Allow for backtracking when a problem is encountered, without the necessity of starting over afresh.

Above all, our framework makes it clear that reasoning involves three components: facts, sources, and indicators. Real-life reasoning typically involves all three. Any protocol or tool developed in line with our model should provide the user with regular prompts to check all three components. Many examples of failures in human reasoning and analysis have resulted from a neglect of one or more of the three basic components.

Jon Barwise

I think it is appropriate to end with a quotation from my former friend and colleague Jon Barwise, whose untimely death in 2000 deprived the world of one of the most innovative logicians of the twentieth century. In his collected work *The Situation in Logic* [2], Barwise wrote [pp.xv–xvi]:

> Back in the days before I became interested in the situated aspects of logic, I sometimes used to wonder how logicians felt in the first quarter of this century. Did they *feel* confused. Reading the literature of that period, one senses the extent to which they were groping toward the view of logic that eventually emerged, but also the extent to which they were still in the dark about what was central and what was peripheral. One also realizes that they were just missing certain key distinctions. In other words, they were confused. It was only with the pioneering work of Gödel, Church, Turing, Tarski, and Kleene in the 1930's that the modern conception of logic really took hold.

> I now feel I have some idea of how logicians must have felt in that period before the really seminal work, since I feel we are in an analogous stage now ... As we try to let go of some of the simplifying idealizations made in standard logic, we too are groping for the key notions, and probably missing some key distinctions. In giving up these simplfying assumptions, there are many things to be rethought, many choices to be made, and many things to be tried. It is an exciting time, if you have the patience for that sort of thing, and a taste for the basic task of conceptual clarification. But it is also frustrating ...
>
> ... There is only one point about which I am really certain. That is that the view of language and logic as situated activities is an important one, and that situating logic is a task that must be carried out if we are to come to grips with some of the problems that currently vex the field.

I say Amen to that.

Acknowledgements

Many of the ideas presented in this paper were developed over several years, during which time my research was based at CSLI. Preparation of the paper was supported in part by an award from the Avanced Research Development Agency, under a subcontract to Veridian Systems, Inc. (now a division of General Dynamics), as part of ARDA's NIMD Program.

Notes

1. [14], Chapter 5, p6. The book is currently available only on the Web. Page numbers are internal to each chapter.

2. [14], Chapter 6, p1.

3. [14], Chapter 4, p1.

4. Use of theory based simulations in training are fairly common, and widely believed to be beneficial to the trainee, although we have to say that we are not aware of any concrete evidence to that effect.

5. At the time, little attention was paid to context. The ancient Greek logicians simply set out to understand logical reasoning. It is only with hindsight that we can look back at what they did and recognize that they ignored issues of context.

6. In our book *The Math Gene*, we argued that the capacity to construct abstract mental models is cognitively equivalent to having language, generally viewed as the key capacity that sets humans apart from all other species.

7. Heuer [14, Chapter 4, p1] observes: "Judgment is an integral part of all intelligence analysis. While the optimal goal of intelligence collection is complete knowledge, this goal is seldom reached in practice. Almost by definition of the intelligence mission, intelligence issues involve considerable uncertainty. Thus, the analyst is commonly working with incomplete, ambiguous, and often contradictory data. The intelligence analyst's function might be described as transcending the limits of incomplete information through the exercise of analytical judgment."

8. And for a great many more reasons we described at length in our book *Goodbye Descartes*, New York, NY: Wiley (1997).

9. *Presidential Decisionmaking in Foreign Policy: The Effective Use of Information and Advice*, Boulder, CO: Westview Press (1980), Chapter 2.

10. Our use of the term "indicators" with this meaning comes from social science.

References

[1] Attardi, G. and Simi, M. Proofs in context, in *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth Conference* (1994).

[2] Barwise, J. *The Situation in Logic*, CSLI Lecture Notes 17 (1989).

[3] Buvac, S and Mason, I. Propositional logic of context, in *Proceedings of the Eleventh National Conference on Artificial Intelligence, Washington, D.C.* (1993).

[4] Devlin, K. *Logic and Information*, Cambridge University Press (1991).

[5] Devlin, K. *Goodbye Descartes: The End of Logic and the Search for a New Cosmology of the Mind*, John Wiley (1997).

[6] Devlin, K. *Infosense: Turning Information into Knowledge*, W. H. Freeman (1999).

[7] Confronting context effects in intelligence analysis:
How can mathematics help?, in preparation

[8] Devlin, K. and Rosenberg, D. *Language at Work: Analyzing Communication Breakdown in the Workplace to Inform Systems Design*, Stanford University: CSLI Publications and Cambridge University Press (1996).

[9] Farquhar, A. and Buvac, S. Putting Context Logic into Practice, technical report of the Stanford Knowledge Systems Laboratory, January, 1997.

[10] Farquhar, A, Dappert, A, Fikes, A. and Pratt, W. Integrating Information Sources Using Context Logic, technical report of the Stanford Knowledge Systems Laboratory, January 1995.

[11] George, A. *Presidential Decisionmaking in Foreign Policy: The Effective Use of Information and Advice*, Boulder, CO: Westview Press (1980).

[12] Giunchiglia, F. Contextual reasoning, *Epistemologia*, XVI (1993), pp.345–364.

[13] Guha, R.V. *Contexts: A Formalization and Some Applications*, Ph.D. thesis, Computer Science Department, Stanford University (1991).

[14] Heuer, R. J. Jr., *Psychology of Intelligence Analysis*, Central Intelligence Agency (1999). Available on the Web at http://www.odci.gov/csi/books/19104/ .

[15] McCarthy, J. Generality in artificial intelligence, *Communications of the ACM*, 30 (12), 1987, pp.1010–1035.

[16] McCarthy, J. Notes on formalizing context, in *Proceedings of the Thirteenth International Joint Conference in Artificial Intelligence*, Chamberry (1993).

[17] Piatelli-Palmerini, M. *Inevitable Illusions: How Mistakes of Human Reason Rule Our Minds*, John Wiley (1996).

[18] Shoham, Y. Varieties of context, in Lifschitz, V. (ed) *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, Academic Press (1991), pp.393–408.

[19] Thomason, R. Type Theoretic Foundations for Context, Part 1: Contexts as Complex Type-Theoretic Objects (1999), preprint available for download at http://www.eecs.umich.edu/ rthomaso/documents/context/

[20] Thomason, R. Contextual Intensional Logic: Type-Theoretic and Dynamic Considerations (2001), preprint available for download at http://www.eecs.umich.edu/ rthomaso/documents/context/