LA-UR- LA-UR- 02-1257

Approved for public release;
distribution is unlimited.

Title: A Knowledge Typology for Imprecise Probabilities

Author(s): Gregory D. Wilson, D-1
Lauren J. Zucker, D-1

Submitted to: THIRD INTERNATIONAL CONFERENCE ON MATHEMATICAL METHODS IN RELIABILITY

A Knowledge Typology for Imprecise Probabilities

Greg Wilson
Statistical Sciences Group
Los Alamos National Laboratory
Los Alamos, NM 87544
United States
Gdwilson@lanl.gov

Lauren J. Zucker
Statistical Sciences Group
Los Alamos National Laboratory
Los Alamos, NM 87544
United States
lzucker@lanl.gov

Abstract

When characterising the reliability of a complex system there are often gaps in the data available for specific subsystems or other factors influencing total system reliability. At Los Alamos National Laboratory we employ ethnographic methods to elicit expert knowledge when traditional data is scarce. Typically, we elicit expert knowledge in probabilistic terms. This paper will explore how we might approach elicitation if methods other than probability (i.e., Dempster-Shafer, or fuzzy sets) prove more useful for quantifying certain types of expert knowledge. Specifically, we will consider if experts have different types of knowledge that may be better characterized in ways other than standard probability theory.

1. Introduction

At Los Alamos National Laboratory (LANL), Information Integration Technology (IIT) is one new method now being developed for characterising the reliability of large complex. Part of the IIT approach involves mapping out the elements of the complex system and the relationships between these elements. Elements may be physical parts of the system or other factors, such as ageing and temperature. Once we have a diagram that captures the complex dependencies of the system, we map in the available data about the performance or impact of each of the elements. For large systems the available data is often diverse: numerical test data, historical data on a similar element, computer model output, or expert judgment. We often have to elicit expert judgment in areas where the three other types of data are not available. If the choice is between leaving the influence of that element out of the system map altogether or doing a structured elicitation to quantify what the experts know about the likely performance of the part, we choose to do the latter. We believe it is better to have the best available information, no matter what form it is in, than no information at all.

However, expert judgements often prove difficult to quantify. The foundational text on expert judgment elicitation (Mary Meyer and Jane Booker's *Eliciting and Analysing Expert Judgment: A Practical Guide*) discusses a multitude of quantification strategies for eliciting information from an expert to fit particular statistical models. But, for some types of expert knowledge, perhaps probabilities are not the best way to quantify the knowledge—maybe the expert thinks in terms that don't gracefully equate to probabilities. Since the goal is to make the best information available for analysis, we want to use methods that capture most accurately how the expert thinks. Towards this end, Meyer and Booker's list of quantification strategies can also serve as guide for an expert-driven, instead of a model-driven, approach. If the elicitor can determine what metrics and terms the expert uses to understand the problem, then he or she can elicit information in terms familiar to the expert, instead of forcing the expert to express information in artificial terms suited to a statistical model of the analyst's choosing.

Even though Meyer and Booker offer many quantification strategies, most expert judgment elicitation (at LANL and elsewhere) is conducted in the realm of probability. The convenient assumption being that all knowledge an expert might have about a topic can be expressed in terms of probabilities. Statisticians are most comfortable working with probabilities, so experts are asked either to provide

probabilistic information about the topic in question or their answers are taken and transmogrified into probabilities. This highlights (a) potential (translation) gap(s) between the language of the expert and the language of the statistician, and in this communication gap, important knowledge can be lost. Given the rising interest in non-probabilistic uncertainty quantification methods like Dempster-Shafer and fuzzy sets, this paper asks the question: Are there different types of expert knowledge that may be better quantified by different types of uncertainty quantification? That is, are there some types of knowledge that map well to probability, others that map well to Dempster-Shafer methods, and others that map well to fuzzy set methods?

2. Kant Always Get What You Want

The first part of answering our driving question is distinguishing the different types of expert knowledge. Many different knowledge typologies could be the starting point for such an investigation since philosophers through the centuries have studied and debated questions of epistemology. But in the eighteenth century, the German philosopher Immanuel Kant wrote a landmark text addressing the shortcomings of epistemological thought up to his time. The paradigm shift inspired by Kant's contributions had a profound influence on subsequent philosophers and is key to how we think about knowledge and understanding today. By drawing upon Kant's framework for understanding knowledge, we will attempt to differentiate the types of knowledge gained though expert judgment elicitation and explore their correspondence with different methods of uncertainty quantification.

Whereas empiricist philosophers, such as John Locke or David Hume, believed that knowledge originates from experience (the world is only knowable *a posteriori*) and rationalists, such as Rene Descartes, believed that all knowledge of the world could be constructed out of simple innate cognitive principles (the world is knowable *a priori*), Kant found both of these positions incomplete. Kant felt that a framework for understanding knowledge must be more complex than the *a posteriori/a priori* distinction in the debate between the rationalists and empiricists. He further divided knowledge into *analytic* and *synthetic* categories, charging that philosophers before him had ignored this necessary distinction. The subject of an *analytic* statement contains the predicate, and therefore, *analytic* statements are necessarily true. According to Kant, previous philosophers had conflated *analytic* with *a prior*, the former being a necessarily true statement and the latter a statement our minds implicitly know to be true. These categories do intersect in statements that are logically true or straightforward matters of definition such as, "The bachelor is unmarried." We know the statement is true without needing to appeal to our senses and experience because the predicate or notion of "unmarried" is contained within the subject, "bachelor." (Philosopher Pages on Kant)

Synthetic statements combine or synthesize two or more concepts to present new information. It makes sense that some philosophers believed the truth of *synthetic* claims was contingent on sensory experience or *a posteriori* knowledge. In many cases, this is true, such as in the statement "John is unmarried" or "The tree is 10 meters tall." Both of these statements contain more than one idea, such as "tree" and "10 meters," and the truth of the synthesized statement can only be ascertained by investigating or measuring the external world. But neither of these statements are necessarily true, the idea of "unmarried" is not contained in the idea of "John"; but both can be proven true through experience. Table 1 shows a matrix of examples.

	Analytic (The predicate is contained in the subject, necessarily true)	**Synthetic** (Two or more different concepts are synthesized to produce a new claim, contingently true)
a posteriori (Knowledge from experience)	***Not Existent***	John is unmarried. The tree is 10 meters tall.
a priori (Knowledge	The bachelor is unmarried.	Every event has a cause.

independent of experience)		2 + 3 = 5

Table 1. A Matrix of statements/types of knowledge based on Immanuel Kant's *Critique of Pure Reason.* (Content derived from McCormick web page)

We see in the table that the four potential pairings only result in three logically possible pairings. An *analytic-a posteriori* statement cannot exist, because a statement that is necessarily/explicatively true (*analytic*) cannot need to appeal to experience for truth (*a posterior*). Much more interesting to Kant, and most relevant to this paper, are *synthetic – a priori* statements, statements that provide new information that we know to be true, but for proof of which we cannot appeal to experience. An important example of a *synthetic-a priori* claim is "All events have a cause." Our minds know this to be true, since all of our physical actions, scientific research, beliefs, etc. are based on causation, therefore it is *a priori.* However, "cause" is not part of the definition of "event" and just by analysing the term "event" we do not derive that a "cause" is necessary. This statement synthesizes the ideas of "cause" and "event." The conundrum of how we can know a statement is true if it is not logically or definitionally true, and our senses cannot provide proof resulted in the Empiricist David Hume's problematic assertion that causation does not exist. Understanding that the idea of causation is necessary for daily life, however, Kant was driven to discover a different way of explaining how we know *synthetic- a priori* statements.

Kant's answer is that the mind imposes structure, such as cause and effect or a concept of time-space on the information gathered by the senses. This mental structure is what allows us to formulate *synthetic-a priori* judgments. Kant articulated that human knowledge is an interplay of the exterior world and the innate cognitive structures of the mind. The result is less comforting than the Empiricist view that we can truly know the external world, or the Rationalist position that all we need to understand the external world we contain within our minds. Instead, he suggests that what we can know about the external world is mediated through the way our minds are organized. The role that the mind's structure acts out on *synthetic* information, such as organizing events based on the principle of causation, transforming the information into knowledge so that we know it is true or *a prior*, has important implications for how we understand expert knowledge.

3. Kant and Expert Judgment

So far with Kant's breakdown of knowlegde types, we have statements that are logically or definitionally true (*analytic-a priori*), statements that can be proven true through experience or testing (*synthetic-a posteriori*), and statements that we know to be true because of the interplay of experience and innate cognitive principles (*synthetic-a priori*). When discussing *synthetic-a priori* knowledge Kant enumerated four categories of thought that he felt constituted the innate cognitive framework that brought structure to experiene, and we will discuss those categories in a later section. But for now, lets consider how Kant's matrix above fits in with the types of data and/or expert judgement one might encounter when working on a reliability problem for a large complex system.

Analytic-a priori statements may not seem of much interest in their true sense—in the true tautological sense of redundant commonly used terms like bachelor and unmarried. But perhaps there is a similarity here with the tacit knowledge and cultural/disciplinary assumptions that the system experts will possess. While the associations will not be something a lay person would understand, the system experts will certainly know that certain things in their system are "just so" because "that's the way things are." Perhaps these conclusions do not come from innate cognitive categories, but ones that have been instilled by disciplinary or cultural training in the same way that all of us have been trained to understand what a bachelor is and what unmarried means. The engineering weltanschauung, for example, can generally describe most things in the world in terms of a block diagram that represents a process of discrete events. While it doesn't seem to most people that "The bachelor is unmarried" and "A nuclear explosion is a process" are similar, within a particular set of disciplinary assumptions, the

latter statement is nearly tautological. But it is the case that tautology is not the right word here. The association of process and explosion is not a strict definitional one, so perhaps it is a relational one that is made possible by the innate cognitive structures Kant attributes to *synthetic-a priori* knowledge. If we take experience to mean things that one has learned, then an interpretive statement like "A nuclear explosion is a process" would be an interplay of engineering learning/experience and a cognitive structure that supplies the idea that dissimilar things can share some similarities and that some events happen in sequence. We often ask experts to make interpretive statements that synthesize past knowledge into new ideas. It would be difficult to say these statements are "true" outside the *synthetic-a priori* category.

Synthetic-a posteriori statements are known true or can be proven true through testing or experience. We can know the statement "This train stops in Boston" is true because we ride the train every week and it always stops in Boston. Likewise we could look at the train schedule, or we could ask the conductor (perhaps an expert) where the train stops, or we could just get on the train and find out. That is, we could collect data to help us know the truth of the statement. We needn't rely on our own experience, though, we could ask someone else with train experience or we could consult another person's data (i.e., quantified experience). In theory, even a statement like "One thousand people would die if anthrax spores were released in this building's air ducts" is *synthetic-a posteriori.* We "could" know this from experience. In reality though, it is not something that we would test, and if this were a statement elicited from a bio-terror expert, we would assume that the statement did not come from direct experience, but is extrapolated from related experience. Such an interpretive statement must be drawing upon deeper cognitive structuring principles that allow extrapolate from what has been learned/experienced and what is possible.

We should also consider statements like "there is an 85% chance that this train stops in Boston," which seems different from "this train stops in Boston." There is no difference in the statements if the person making the former statement is a frequentist statistician who assumes that there are true probabilities and has observed that 85 out of 100 trains stop in Boston, or that this particular train stops in Boston 85 out of 100 times. If this is a statement of subjective probability, though, we can assume that the person making the statement, like the bio-terror expert above, is extrapolating from some base of related *synthetic-a posteriori* knowledge, and is perhaps utilizing some deeper structuring principles: that it is possible this train stops in Boston and that there is a quantifiable possibility of that happening.

We certainly cannot discuss *synthetic-a priori* knowledge for very long without explaining the four categories I've alluded to above. Kant postulated this type of knowledge was most interesting and problematic for philosophers because it involved a type of knowing that represented an intersection of experience with innate cognitive structures. Matt McCormick writes,

> In the **Analytic of Concepts** section of the *Critique*, Kant argues that in order to think about input from sensibility, sensations must conform to the conceptual structure that the mind has available to it. By applying concepts, the understanding takes the particulars that are given in sensation and identifies what is common and general about them. A concept of "shelter" for instance, allows me to identify what is common in particular representations of a house, a tent, and a cave. (Citation)

McCormick goes on to explain that the concept of shelter certainly does come from experience, but the ability to understand that there are classes of things that can relate in similar ways is an innate cognitive structure. Kant suggested those structures could be represented by the following four categories.

Quantity	**Quality**
Unity (one)	Reality (A is B)
Plurality (many)	Negation (A is not B)
Totality (all)	Limitation (A is B if . . .)

Relation	**Modality**
Inherence (A, B, & C are a group)	Possibility (it could be)
Causality/Dependence (If A then B)	Existence (it is)
Community (A is part of C)	Necessity/Contingency (A must be given B)

When you think in terms of expert knowledge, an expert can know things like bachelors are unmarried and they can know that a certain tree is 10 feet tall, but when we get to types of knowledge that fit into the *synthetic-a priori* category, it seems we are talking about judgments more than knowledge. It is often difficult to logically tease out how something like 2 + 3 = 5 represents the synthesis of experience and innate structuring principles, but a statement like the anthrax statement above shows us that what an expert would be doing in that instance is offering judgments. If the expert actually knew how many people would die, it would be a *synthetic-a posteriori* statement instead. The disciplinary issue that is also raised above also boils down to judgments. Disciplinary training is another type of experience, and when interpretations are called for, that experience must be synthesized with cognitive principles to get to something that isn't knowable. If experience (and maybe *synthetic-a posteriori* and *analytic-a priori* knowledge is what experience is) plus cognitive principles equals interpretations, and the four categories above are a realistic representation of those cognitive principles, then I think we have a good typology to begin to understand what kinds of judgments experts can supply. Even though experts may be drawing upon a vast range of experience, they all will be drawing upon the same set of cognitive principles. At the time of the conference we will further look into how judgments structured by these principles map onto different uncertainty quantification methods.

References

Kant, I. (1998). Critique of Pure Reason. New York: Cambridge.

Kemerling, G. (1997-2002). "Kant: Synthetic A Priori Judgments." http://www.philosophypages.com/hy/5f.htm#jdgts. Accessed March 4, 2002.

McCormick, M. "Immanuel Kant (1724-1804)." http://www.csus.edu/indiv/m/mccormickm/IEPKantArt.htm. Accessed July 23, 2001.

Meyer, Mary A. ; Booker, Jane M. (2001). *Eliciting and analyzing expert judgment : a practical guide.* Philadelphia: Society for Industrial and Applied Mathematics and American Statistical Association.