

Probabilistic Quantifier Logic for General Intelligence: An Indefinite Probabilities Approach

Matthew IKLÉ^a, Ben GOERTZEL^b

^a*Adams State College, Alamosa, Colorado
and Novamente LLC.*

^b*Novamente LLC*

Abstract: Indefinite probabilities are a novel technique for quantifying uncertainty, which were created as part of the PLN (Probabilistic Logic Networks) logical inference engine, which is a key component of the Novamente Cognition Engine (NCE), an integrative AGI system. Previous papers have discussed the use of indefinite probabilities in the context of a variety of logical inference rules, but have omitted discussion of quantification. Here this gap is filled, and a mathematical procedure is provided allowing the propagation of indefinite probabilities through universal and existential quantifiers, and also through a variety of fuzzy quantifiers corresponding to natural language quantifiers (such as “few”, “man”, “a lot”, “hardly any”, etc.). Proper probabilistic handling of various quantifier transformation rules is also discussed. Together with the ideas in prior publications, and a forthcoming sequel paper on indefinite probabilities for intensional inference, these results allow probabilistic logic based on indefinite probabilities to be utilized for the full scope of inferences involved in intelligent reasoning. To illustrate the ideas and algorithms involved, we give two concrete examples: Halpern’s “crooked lottery” predicate, and a commonsense syllogism that uses fuzzy quantifiers together with the standard PLN term logic deduction rule.

1. Introduction

Probability theory plays an increasingly large role in AI research, spanning the divide between AGI and narrow AI research, and infiltrating every subdomain of the AI field. Probabilistic methods hold a prominent role in linguistics ([1]), robotics ([2]) speech processing ([3], [4]) and data analysis and modeling ([5]), to name only a handful of the many areas that could be cited. Our own venture into the domain of Artificial General Intelligence, the Novamente Cognition Engine (NCE) ([6]), utilizes probability theory as a common methodology and language binding together diverse cognitive components.

Probability theory in itself, however, is a very general tool, and may be applied in a variety of different ways, in conjunction with a diversity of other formalisms. One of the key foci of recent research regarding the utilization of probability theory for AI is the unification of probability and logic ([7], [8]). Of course, the simple unification of probability and logic is trivial, as logic is a general tool that can be used to reason about anything; but unifying the two formalisms in a way that is helpful in terms of

constructing AI system is another story. One approach to carrying out this unification is the PLN (Probabilistic Logic Networks) framework ([9]), which is a key component of the Novamente Cognition Engine (NCE; [6]), an integrative AGI system developed by the authors and their colleagues during the period since 2001. PLN has been specifically constructed to unify probability and logic in a manner that

- supports the full scope of inferences required within an intelligent system, including e.g. first and higher order logic, intensional and extensional reasoning, and so forth
- lends itself naturally to methods of inference control that are computationally tractable, and able to make use of the inputs provided by non-logical cognitive mechanisms (such as dynamic attention allocation and evolutionary learning, both of which play roles in the NCE)

PLN as a general framework supports multiple measures of uncertainty, but the primary measure utilized currently is “indefinite probabilities,” a novel measure formed by hybridizing Walley’s imprecise probabilities with the Bayesian notion of credible intervals ([9], [10]). In order to utilize indefinite probabilities with PLN, each of PLN’s logical inference rules must be associated with an “indefinite truth value formula or procedure,” which tells you, given the indefinite truth values associated with the premises of the inference rule, what is the indefinite truth value associated with the conclusion of the rule. Prior papers have given indefinite truth value formulas for a number of PLN inference rules, but have not dealt with quantifiers. In this paper we remedy this deficit, and discuss the propagation of indefinite probabilities through universal and existential quantifiers as well as fuzzy quantifiers.

The overall goals of our approach to quantifiers, as presented here, are as follows:

- 1) logical and conceptual consistency
- 2) agreement with standard quantifier logic for the crisp case (for all expressions to which standard quantifier logic assigns truth values)
- 3) gives intuitively reasonable answers in practical cases
- 4) compatibility with probability theory in general and PLN in particular
- 5) handles fuzzy quantifiers as well as standard universal and existential quantifiers

The application to fuzzy quantifiers is particularly conceptually satisfying, in that it places these slippery notions, with their close connection to ambiguous natural language, on a firm probabilistic foundation, in a way that (we argue) is more conceptually coherent and pragmatically useful than prior attempts in similar directions. Zadeh refers to fuzzy quantifiers characterizing absolute cardinality (such as several, few, many etc.), as fuzzy quantifiers of the first kind; those characterizing relative cardinality (such as most, almost all, many) as fuzzy quantifiers of the second kind; and he also introduces a third-kind of fuzzy quantifiers which are ratios of quantifiers of the second kind ([11], [12]). Here we explicitly demonstrate how to treat analogs of fuzzy quantifiers of the first kind within the indefinite probabilities framework. Straightforward generalizations of the ideas contained here, however, can also generate analogs of fuzzy quantifiers of the second and third kinds.

There is no objective standard via which to compare different approaches to the combination of uncertain inference and quantifier inference. However, our own feeling

is that the approach articulated here is more satisfactory from a pragmatic AI perspective than prior attempts. As we are committed to a probabilistic foundation for AI, we are not fully satisfied with nonprobabilistic fuzzy approaches; and the classical probabilistic approaches appear to us to have conceptual problems. For instance, Halpern's approach [7] involves a counterintuitive mixture of two different kinds of probabilities (subjective and frequentist). Our approach does not involve these sorts of interpretive subtleties, avoiding them due to the introduction of third-order probabilities. On the other hand, the introduction of quantifiers into Markov Logic Networks as done in [13], while interesting and surely useful for certain applications, does not comprise a suitably general framework for uncertain quantifier logic: it deals with evaluation of expressions involving quantifiers across databases, but not with abstract quantifier manipulations. The approach presented here is comprehensive, conceptually coherent, probabilistically grounded and computationally tractable, and for these reasons we suggest it to be an adequate formulation of uncertain quantifier logic for AGI purposes.

2. Review of Indefinite Probabilities

As described in ([9], [10]), indefinite probabilities were motivated in part by Walley's imprecise probabilities ([14]). A truth-value for an indefinite probability takes the form of a quadruple $([L, U], b, k)$. The meaning of such a truth-value, attached to a statement S is, roughly: There is a probability b that, after k more observations, the truth value assigned to the statement S will lie in the interval $[L, U]$. Just as with Walley's imprecise probabilities, we interpret an interval $[L,U]$ by assuming some particular family of distributions (usually Beta) whose means lie in $[L, U]$.

The inclusion of the credibility level b , not present in Walley's interval truth values, allows our intervals to remain narrower than those produced by Walley's and other interval probabilities in the literature. As we showed in ([10]), our approach reduces to Walley's imprecise probabilities by setting $b = 1$. In that prior paper we also discuss, in detail, the philosophical and conceptual underpinnings of indefinite probabilities and so will not repeat those here.

2.1. Inference with Indefinite Probabilities

We now review the basic methodology by which indefinite probability formulas may be derived corresponding to inference rules. More detail on this methodology is given in ([9], [10]).

We assume that the inference rules in question already come with truth value formulas that apply in the case the probability value attached to each premise is exactly known (an ideal limiting case that will essentially never occur in reality). An example inference rule that has been treated this way is the term logic deduction rule

Inh A B
 Inh B C
 |-
 Inh A C

where Inh is shorthand notation for ExtensionalInheritance, a specific Novamente/PLN link type as defined in ([9]), which refers to a probabilistic subset relationship.

The probabilistic truth value formula for the deduction rule is given by

$$S_{AC} = S_{AB} S_{BC} + (1-S_{AB}) (S_C - S_B S_{BC}) / (1- S_B).$$

To execute inference formulas (corresponding to inference rules such as the deduction rule given just above) using indefinite probabilities, we make heuristic distributional assumptions. We assume a “second-order” distribution that has $[L, U]$ as a $(100b)\%$ credible interval. We then assume “first-order” distributions whose means are drawn from the second-order distribution. These distributions are to be viewed as heuristic approximations intended to estimate unknown probability values existing in hypothetical future situations. While the accuracy and utility of the indefinite probability method may depend on the appropriateness of the distributional assumptions, we have found that the beta and bimodal families seem adequate for most cases arising in real-world inference problems.

Given a logical inference rule, the indefinite probability based inference process proceeds in three basic steps.

1. Given intervals, $[L_i, U_i]$, of mean premise probabilities, we first find a distribution from the “second-order distribution family” supported on $[L1_i, U1_i] \supset [L_i, U_i]$, so that these means have $[L_i, U_i]$ as $(100 \cdot b_i)\%$ credible intervals.
2. For each premise, we use Monte-Carlo methods to generate samples for each of the “first-order” distributions with means given by samples of the “second-order” distributions. We then apply the inference rules to the set of premises for each sample point, and calculate the mean of each of these distributions.

Find a $(100 \cdot b_i)\%$ credible interval, $[L_f, U_f]$, for this distribution of means; e.g. by assuming a symmetric credible interval about the mean.

As an example, note the results obtained by doing indefinite probability calculations for the following deduction:

Women are beautiful.
 Beautiful things bring happiness.
 |-
 Women bring happiness

We summarize the truth value premises and conclusion in the following table:

Premises	Truth Value
Women	<[0.45, 0.55], 0.9, 10>
Women are beautiful	<[0.8, 0.95], 0.9, 10>
Beautiful things	<[0.4, 0.8], 0.9, 10>
Beautiful things bring happiness	<[0.8, 0.95], 0.9, 10>

Happiness	<[0.4, 0.9], 0.9, 10>
Conclusion	Truth Value
Women bring happiness	<[0.76939, 0.87028], 0.9, 10>

3. Quantifiers in Indefinite Probabilities

The above approach works perfectly well for many inference rules, but is inadequate to handle universal, existential or fuzzy quantifiers. The best way we have found to handle quantifiers within the indefinite probabilities framework is to introduce another level of complexity and utilize third-order probabilities. To understand this, we first consider the problem of “direct evaluation” of the indefinite truth values of universally and existentially quantified expressions.

We solve this problem via taking a semantic approach that is considerably conceptually different from the one standardly taken in formal logic. Normally, in logic, expressions with unbound variables are not assigned truth values; truth value assignment comes only with quantification. In our approach, however, we assign truth values to expressions with unbound variables, yet without in doing so binding the variables. This is unusual but not contradictory in any way: an expression with unbound variables, as a mathematical entity, may certainly be mapped into a truth value without introducing any mathematical or conceptual inconsistency. This allows one to define the truth value of a quantified expression as a mathematical transform of the truth value of the corresponding expression with unbound variables, a notion that is key to our approach.

This unusual semantic approach adds a minor twist to the notion that our approach to uncertain inference on quantified expressions reduces to standard crisp inference on quantified expressions as a special case. The twist is that our approach reduces to the standard crisp approach in terms of truth value assignation for all expressions for which the standard crisp approach assigns a truth value. However, our approach also assigns truth values to some expressions (formulas with unbound variables) to which the standard crisp approach assigns no truth value.

Following up on this semantic approach, we will now explain how, if we have an indefinite probability for an expression $F(t)$ with unbound variable t , summarizing an envelope E of probability distributions corresponding to $F(t)$, we may derive from this an indefinite probability for the expression “ForAll x , $F(x)$ ”? (Having carried out the transform in this direction, it will then be straightforwardly possible to carry out a corresponding transform in reverse.) The approach we take here is to consider the envelope E to be part of a higher-level envelope $E1$, which is an envelope of envelopes. The question is then: given that we have observed E , what is the chance (according to $E1$) that the true envelope describing the world actually is almost entirely supported within $[1-e, 1]$, where the latter interval is interpreted to constitute “essentially 1” (i.e., e is the margin of error accepted in assessing ForAll-ness), and the phrase “almost entirely supported” is defined in terms of a threshold parameter?

Similarly, in the case of existential quantification, we want to know the indefinite probability corresponding to “ThereExists x , $F(x)$.” The question is then: given that we have observed E , what is the chance (according to $E1$) that the true envelope describing the world actually is *not* entirely supported within $[0, e]$, where the latter interval is

interpreted to constitute “essentially zero” (i.e., ϵ is the margin of error accepted in assessing ThereExists-ness)?

The point conceptually is that quantified statements require you to go one level higher than ordinary statements. So if ordinary statements get second-order probabilities, quantified statements must get third-order probabilities. And, the same line of reasoning that holds for “crisp” universal and existential quantifiers, turns out to hold for fuzzy quantifiers as well. In fact, in the approach presented here, crisp quantifiers are innately considered as an extreme case of fuzzy quantifiers, so that handling fuzzy quantifiers doesn’t really require anything extra, just some parameter-tuning.

The following subsections elaborate the above points more rigorously.

3.1. Direct Evaluation of Universally Quantified Expressions

We first consider the case of the direct evaluation of universally quantified statements, an inference rule for which the idea is as follows: Given an indefinite truth value for $F(x)$, we want to get an indefinite TV for $G = \text{ForAll } x, F(x)$.

The roles of the three levels of distributions are roughly as follows. The first- and second-order levels play the role, with some modifications, of standard indefinite probabilities. The third-order distribution then plays the role of “perturbing” the second-order distribution. The idea is that the second-order distribution represents the mean for the statement $F(x)$. The third-order distribution then gives various values for x , and the first-order distribution gives the sub-distributions for each of the second-order distributions.

The process proceeds as follows:

Step 1: Calculate [lf1,uf1] Interval for the third-order distribution.

This step proceeds as usual for indefinite probabilities: see ([9], [10]). Given L , U , k , and b , set $s = 0.5$. We want to find a value for the variable **diff** so that the probability density function defined by

$$f(x) = \frac{(x - L1)^{ks} (U1 - x)^{k(1-s)}}{\int_{L1}^{U1} (x - L1)^{ks} (U1 - x)^{k(1-s)} dx}$$

where $L1=L-\text{diff}$ and $U1=U+\text{diff}$ is such that

$$\frac{\int_{L1}^L (x - L1)^{ks} (U1 - x)^{k(1-s)} dx}{\int_{L1}^{U1} (x - L1)^{ks} (U1 - x)^{k(1-s)} dx} = \frac{1 - b}{2}$$

and

$$\frac{\int_{U1}^U (x - L1)^{ks} (U1 - x)^{k(1-s)} dx}{\int_{L1}^{U1} (x - L1)^{ks} (U1 - x)^{k(1-s)} dx} = \frac{1 - b}{2}.$$

Once one of these last two integrals is satisfied, they both should be. Alternatively, one can find **diff** for which

$$\frac{\int_{L1}^U (x - L1)^{ks} (U1 - x)^{k(1-s)} dx}{\int_{L1}^U (x - L1)^{ks} (U1 - x)^{k(1-s)} dx} = b.$$

Step 2: Generate vectors of means for perturbed F(x) values.

Generate vectors of mean values for each premise.

Step 2.1: Generate values from desired “third-order” distribution family.

At present we are using only beta distributions. Generate a vector of length n1 of random values chosen from a standard beta distribution.

Step 2.2: Scale random means to interval [lf1,uf1]

Scale the vector from step 2.1 to [lf1,uf1] using a linear transformation.

Step 3: Generate symmetric intervals [lf2[i], if2[i]] for each of the means found in step 2.

These intervals are now the desired [L1, U1] intervals for the third-order distributions.

Step 4: Generate the second-order distributions.

For each mean for the third-order distributions, generate a sub-distribution. These sub-distributions represent the second-order distributions.

Step 5: Generate first-order distributions with means chosen from the second-order distributions.

Step 6: Determine the percentage of elements in each first-order distribution that lie within the interval [1-e, 1].

Recall that we are using the interval [1-e, 1] as a “proxy” for the probability 1. The goal here is to determine the fraction of the first-order distributions that are “almost entirely contained” in the interval [1-e, 1]. By almost entirely contained, we mean that the fraction contained is at least proxy_confidence_level (PCL).

Step 7: Find Conclusion ([L,U],b) Interval

For each of the third-order means, we calculate the average of all of the second-order distributions that are almost entirely contained in [1-e, 1], giving a list of n1 elements, **probs**, of probabilities. We finally find the elements of **probs** corresponding to quantiles

$$n1 * (1-b) / 2 \text{ .round for L, and } n1 * (0.5+b) / 2 \text{ .round for U.}$$

Unlike the case for crisp predicate logic, however, this rule is not, in general, true using indefinite probabilities. For example, consider the following set of premises with parameter settings $e=0.5$ and $PCL=0.7$: truth value for $G = \langle [0.45, 0.46], 0.9, 10 \rangle$ and truth value for $F(x) = \langle [0.71, 0.72], 0.9, 10 \rangle$. Then the result for $\forall X: G \text{ AND } F(x)$ becomes $\langle [0.0, 0.04913], 0.9, 10 \rangle$, while that for $G \text{ AND } \forall X: F(x)$ is $\langle [0.23046, 0.28926], 0.9, 10 \rangle$. On the other hand, we note that a different set of premises can yield similar results from the two approaches. Assuming the same parameter values for e and PCL , and truth values for both $F(x)$ and G of $\langle [0.99, 1.0], 0.9, 10 \rangle$ gives a result of $\langle [0.98331, 0.99626], 0.9, 10 \rangle$ using $\forall X: G \text{ AND } F(x)$, and a similar result of $\langle [0.98344, 0.99620], 0.9, 10 \rangle$ using $G \text{ AND } \forall X: F(x)$.

For insight into what is happening here, we view $H(F)(t) = G \text{ AND } F(t)$ as a distortion of the distribution of F . In addition, if $J(F) = \forall x J(x)$, then $J(F)$ is a nonlinear distortion of F , so that even though $H(F)$ is a linear distortion, it need not commute with J . An obvious and interesting question is then: Under what combination of premise values and parameter settings do the operators H and J ‘almost’ commute? Due to space considerations, we defer a thorough study of that question to a future paper. It does appear, however, that premise values near 1 lead to better commutativity than do values farther from 1.

5. Fuzzy Quantifiers

Analyzing the indefinite probabilities approach to the quantifiers *ForAll* and *ThereExists*, it should be readily apparent that indefinite probabilities provide a natural method for “fuzzy” quantifiers such as *AlmostAll* and *AFew*.

In our discussion of the *ForAll* rule above, for example, the interval $[PCL, 1]$ represents the fraction of bottom-level distributions completely contained in the interval $[1-e, 1]$. Recall that the interval $[1-e, 1]$ represents a proxy for probability 1.

In analogy with the interval $[PCL, 1]$ representing the *ForAll* rule, we can introduce the parameters `lower_proxy_confidence` (LPC) and `upper_proxy_confidence` (UPC) so that the interval $[LPC, UPC]$ represents an *AlmostAll* rule or *AFew* rule. More explicitly, by setting $[LPC, UPC] = [0.9, 0.99]$, the interval could now naturally represent *AlmostAll*.

Similarly, the same interval could represent *AFew* by setting LPC to a value such as 0.05 and UPC to, say 0.1.

Through simple adjustments of these two proxy confidence parameters, we can thus introduce a sliding scale for all sorts of fuzzy quantifiers. Moreover, each of these fuzzy quantifiers is now firmly grounded in probability theory through the indefinite probabilities formalism.

6. Examples


To further elucidate the above formalism, we now consider two examples. For our first example, we consider an example drawn from [15], which is there called the “crooked lottery” and extensively discussed:

$$\begin{aligned} & [\sim \text{ThereExists } x : \text{Winner}(x) \rightarrow \text{false}] \& \\ & [\text{ThereExists } y : \text{ForAll } x : (\text{Winner}(x) \parallel \text{Winner}(y)) \rightarrow \text{Winner}(y)] \end{aligned}$$


The first clause is intended to represent the idea that everyone has a nonzero chance to win the lottery; the second clause is intended to represent the idea that there is one guy, y , who has a higher chance of winning than everybody else. In [15] Halpern examines various formalisms for quantifying uncertainty in a logical reasoning context and assesses which ones can provide a consistent and sensible truth-value evaluation for this expression.

To evaluate the truth value of this expression using indefinite probabilities, suppose we assume the truth-value for $\text{Winner}(x)$ is $\langle [0.05, 0.1], 0.9, 10 \rangle$. For the second clause we also assume that the truth value for $\text{Winner}(y)$ is $\langle [0.25, 0.5], 0.9, 10 \rangle$ and that the truth value for the implication $(\text{Winner}(x) \parallel \text{Winner}(y)) \rightarrow \text{Winner}(y)$ is $\langle [0.8, 0.9], 0.9, 10 \rangle$. With these assumptions, we then vary the values of the parameters e and PCL for the ThereExists rule to generate the following graphs of the resulting truth value intervals. Note that the parameter values e and PCL used in the ForAll rule were the complements, $1-e$ and $1-\text{PCL}$, of the values used in ThereExists .



Conclusion Lower Interval Limit




Conclusion Upper Interval Limit



The intermediate results for each of the two main clauses provide insight into the interaction between these two clauses. First the results for clause 1: $[\sim\text{ThereExists } x : \text{Winner}(x) \rightarrow \text{false}]$.



For clause 2, $[\text{ThereExists } y : \text{ForAll } x : (\text{Winner}(x) \parallel \text{Winner}(y)) \rightarrow \text{Winner}(y)]$, all the lower limits are 0. The graph of the upper limit is:



For our second example, we consider an extension of the example we used for the standard PLN deduction rule:

Many women are beautiful.
 Almost all beautiful things bring happiness.
 ┆
 Many women bring happiness

In order to use the indefinite probabilities formalism, we first need to determine appropriate values for the parameters LPC and UPC to represent the fuzzy concepts “many” and “almost all.” In practice, in the case where these rules are used within an integrative AGI system such as the NCE, appropriate values for these fuzzy concepts will be determined by the context in which they appear. In one context, for example, the interval $[0.8, 0.9]$ might represent the idea “many,” but in a different situation, we may wish for $[0.6, 0.95]$ to represent “many.”

For our example we set $e=0.1$. Let us suppose that “many” is represented by the interval $[\text{LPC}, \text{UPC}] = [0.4, 0.95]$, and “almost all” by the interval $[0.9, 0.99]$. We will also assume identical truth-values to those in the previous example. The sequence of conclusions is then illustrated in the following tables.

Premises	Truth Value
Women	<[0.45, 0.55], 0.9, 10>
An individual woman is beautiful	<[0.8, 0.95], 0.9, 10>
Conclusion	Truth Value
Many women are beautiful	<[0.35451, 0.63574], 0.9, 10>

Premises	Truth Value
Beautiful things	<[0.4, 0.8], 0.9, 10>
A beautiful thing brings happiness	<[0.8, 0.95], 0.9, 10>
Conclusion	Truth Value
Almost all beautiful things bring happiness	<[0.03906, 0.37464], 0.9, 10>

Premises	Truth Value
Women	<[0.45, 0.55], 0.9, 10>
Many women are beautiful	<[0.35451, 0.63574], 0.9, 10>
Beautiful things	<[0.4, 0.8], 0.9, 10>
Almost all beautiful things bring happiness	<[0.03906, 0.37464], 0.9, 10>
Happiness	<[0.4, 0.9], 0.9, 10>
Conclusion	Truth Value
Many women bring happiness	<[0.41308, 0.53068], 0.9, 10>

7. Conclusions

The issues we have considered here are specific and technical, yet they address a problem that is key to the overall project of creating powerful artificial intelligence. If one wishes to create an AI system that carries out explicit probabilistic estimations, and also explicit logical reasoning, then one is faced with the problem of unifying probability and logic in an elegant and easily controllable way. The PLN framework represents our general solution to this problem; and, in order to be applied in a general way, PLN's truth value formulas must handle quantifiers both crisp and fuzzy. The PLN mathematics must make it possible to propagate uncertain truth values (in whatever representation is chosen, currently indefinite probabilities) through every logical inference rule utilized, including those involving abstract constructs such as quantifiers. Here we have explained how this may be accomplished.

By incorporating a third level of distributions, we have extended the PLN indefinite probabilities method to handle a wide variety of both crisp and fuzzy quantifiers, including the quantifiers that occur in natural language. This is interesting from a purely theoretical perspective, and also from a pragmatic AGI perspective. During the next year or two we expect to put the formulas presented here to work in carrying out inferences on logical relationships derived from natural language understanding and perceptual data analysis within the NCE. Due to the harmony of

these formulas with the overall PLN inference framework, we don't expect inference control to be a major issue: bringing quantifier-based truth value estimation within the scope of PLN, means bringing it within the scope of PLN's powerful inference-control methodology, which uses the history of prior inferences conducted to adaptively prune backwards and forwards chaining inference trees.

Finally, a note on computational tractability may be worthwhile. Although mathematically abstract, the formulas given here are actually not among the more computationally intensive of the indefinite probabilities formulas. The term logic deduction formula presents much greater computational challenges. On modern computers the implementation of the procedures given here yields code that is far from being a bottleneck in the context of a system such as PLN, which has fairly expensive inference control mechanisms going on as well as the execution of inference steps. By all indications the results described here have the potential to be a valuable part of pragmatic probabilistic logical inference in the Novamente Cognition Engine and potentially other AI approaches as well.

References

- [1] Christopher D. Manning, Hinrich Schuetze, Foundations of Statistical Natural Language Processing, Cambridge, MA, MIT Press, 1999.
- [2] Sebastian Thrun, Wolfram Burgard and Dieter Fox, Probabilistic Robotics, Cambridge, MA, MIT Press, 2005.
- [3] John N. Holmes, Wendy J. Holmes, W.J. Holmes, Speech Synthesis and Recognition, Reading, UK, Taylor & Francis Ltd, 2002.
- [4] John Coleman, Introducing Speech and Language Processing, Cambridge, UK, Cambridge University Press, 2005.
- [5] Judea Pearl, Probabilistic Reasoning in Intelligent Systems, San Mateo, Morgan Kaufmann, 1988.
- [6] Ben Goertzel, Moshe Looks, Cassio Pennachin, "Novamente: An Integrative Architecture for Artificial General Intelligence," Proceedings of AAAI Symposium on Achieving Human-Level Intelligence through Integrated Systems and Research, Washington DC, August 2004.
- [7] Joseph Y. Halpern, Reasoning About Uncertainty, Cambridge, MA, MIT Press, 2003.
- [8] progic07: The Third Workshop on Combining Probability and Logic
- [9] Ben Goertzel, Matthew Iklé, Izabela Goertzel, Ari Heljakka, Probabilistic Logic Networks: A Comprehensive Conceptual, Mathematical and Computational Framework for Uncertain Inference, Berlin; New York, Springer Verlag, 2008.
- [10] Matthew Iklé, Ben Goertzel, Izabela Goertzel, "Indefinite Probabilities for General Intelligence," Advances in Artificial General Intelligence, IOS Press, 2007.
- [11] Hans-Jürgen Zimmerman, Fuzzy Set Theory - and its Applications, 4th ed., Berlin; New York, Springer Verlag, 2007.
- [12] Lofti A. Zadeh, George J. Klir, Bo Yuan (eds.), Fuzzy Sets, Fuzzy Logic, and Fuzzy Systems: Selected Papers by Lotfi A. Zadeh, Singapore; River Edge, N.J.: World Scientific Press, 1996.
- [13] Aron Culotta, Andrew McCallum, Practical Markov Logic Containing First-order Quantifiers with Application to Identity Uncertainty.. HLT Workshop on Computationally Hard Problems and Joint Inference in Speech and Language Processing, 2006.

- [14] Peter Walley, Statistical reasoning with imprecise probabilities, London ; New York : Chapman and Hall, 1991.
- [15] Nir Friedman, Joseph Y. Halpern, and Daphne Koller, "First-order Conditional Logic Revisited," AAAI-96, 1996.