Towards a Computationally Viable Framework for Semantic Representation*

Shalom Lappin

University of Gothenburg
shalom.lappin@gu.se

Abstract. Classical theories of formal semantics employ possible worlds to model intensions and modality. If worlds are construed as corresponding to ultrafilters in a lattice of propositions (maximal consistent sets of propositions), then they pose serious problems of tractable representability. Moreover, these theories cannot accommodate vagueness, which is a pervasive feature of predication in natural language. It is also unclear how they can be extended in a straightforward way to explain semantic learning. A cognitively plausible account of interpretation should generate computationally tractable representations of meaning. It must also deal with vagueness and semantic learning. A probabilistic Bayesian approach to natural language semantics provides a more promising approach to these issues. It can also cover epistemic states and inference, in a tractable way. This framework offers the promise of a robust, wide coverage treatment of natural language interpretation that integrates meaning and information update.

Keywords: possible worlds · tractable representations · Bayesian semantics · intensions · modality · epistemic states · information update

1 Introduction

Since [35] a mainstream view among formal semanticists has depended on possible worlds to model the meanings of natural language expressions. Montague imported possible worlds into his model theory through his use of Kripke frame

* The research reported in this paper was supported by grant 2014-39 from the Swedish Research Council, which funds the Centre for Linguistic Theory and Studies in Probability (CLASP) in the Department of Philosophy, Linguistics, and Theory of Science at the University of Gothenburg. Earlier versions were presented in the Syntax Reading Group of University College London and the Colloquium of the Centre for Logic, Language, and Mind at the University of Stockholm, both in October 2017. I am grateful to the participants in these forums for useful discussion and feedback. I would also like to thank Jean-Philippe Bernardy, Rasmus Blanck, Stergios Chartzikyriakidis, Robin Cooper, Simon Dobnik, Staffan Larsson, Per Martin-Löf, Peter Pagin, and Daniel Rothschild for helpful comments and suggestions on some of the ideas presented here. Of course I bear sole responsibility for any errors that may remain.

semantics ([27, 28]) for modal logic. This approach to intensions and modality is anticipated in [5]'s characterisation of intensions as functions from state descriptions to extensions.

Possible worlds have played a central role in the characterisation of belief ([45]) and the related field of epistemic reasoning (see, for example [19]). Dynamic semantics ([17, 18]), and, more recently, Inquisitive Semantics ([6, 7]) use possible worlds to incorporate epistemic elements into formal semantics. They characterise sentence meanings as functions from discourse contexts to discourse contexts. From this perspective speakers use sentences to communicate information by modifying their hearers' representation of a discourse context.[1]

There are, in fact, serious computational problems of representability for worlds. Moreover, specifying intensions as functions from worlds to extensions does not respect important fine-grained distinctions of meaning. I discuss these issues in detail in [30, 31]. In Section 2 I summarise the representability problems raised there, and argue that these must be solved in order to develop a cognitively viable semantics for natural language.

It is possible to "de-modalise" intensions by characterising them as computable functions. This involves borrowing the difference between operational and denotational meaning from the semantics of programming languages and applying it to the meanings of natural language expressions ([30, 31, 12]). In Section 3 I review this approach to representing intensions.

An operational treatment of intensions might solve the representability problem for (some) natural language meanings, and provide the basis for a fine-grained semantics. However, it leaves the treatment of modality and epistemic states untouched. At first glance it would seem that there is no alternative but to invoke worlds to model possibilities, knowledge, and beliefs. But if we are forced to reintroduce worlds to handle these concepts, then we have not solved the representability problem, and so we have not grounded semantics on cognitively viable foundations.

In Section 4 I suggest an approach to this question that avoids worlds. It involves characterising both modality and epistemic states through probability distributions over situations, rather than complete worlds. On his account probabilities are assigned to possible, as well as to actual situations. However, it is not necessary to represent or enumerate the complete class of possible situations, which, as we argue in Section 2, is even more problematic than representing a complete world, or the set of worlds. It is sufficient to specify those situations to which probabilities are assigned, and the situations expressed by the conditions on which a probability assignment depends. Large subclasses of probability models can be efficiently represented, and tractability problems with computing probability distributions or complex sampling spaces can frequently be overcome by estimation and approximation. The probabiistic accounts of modality

[1] It may be possible, at least in principle, to develop versions of Inquisitive Semantics and Dynamic Semantics which do not rely on possible worlds. However, they are an integral element of the current theories.

and epistemic states proposed in this Section constitute the main contribution of the paper.

I offer an overview of some current related work in Section 5, and I briefly compare these approaches to the framework proposed here.

Finally, in Section 6 I present some conclusions, and I briefly indicate several problems to be addressed in future work on these questions.

2 A Repesentability Problem with Worlds

In Kripke frame semantics a model $M = \langle D, W, F, R \rangle$, where D is a non-empty set of individuals, W is a non-empty set of worlds, F is an interpretation function that assigns intensions to the constants of a language, and R is an accessibility relation on W. Formal semanticists have expanded M to include additional indices representing elements of context, such as sets of points in time, and sets of speakers. The elements of W are points at which a maximal consistent set of propositions are satisfied.[2]

There is a one to one correspondence between the elements of W and the elements of the set of maximal consistent sets of propositions. [13, 14, 42] use this correspondence to formally represent worlds as the set U of ultrafilters in the prelattice of propositions. On this approach a proposition p holds at a world w_i iff $p \in u_i$, where $u_i \in U$. The question of how to represent W reduces to the representability of U

To simplify the problem considerably, I assume that the the prelattice on which the elements of U are defined encodes classical Boolean propositional logic. This system is complete and decidable, and so minimal in expressive power. To identify any $u_i \in U$ we need to specify all and only the propositions that hold at u_i. As u_i is an ultrafilter, for any $p_i \in u_i$, all of the propositions that p_i entails are also in u_i, and so it will be an infinite set. We can enumerate the elements of an infinite set if there is an effective procedure (a finite set of rules, an algorithm, a recursive definition, etc.) for recognising its members. It is not clear what an effective procedure for enumerating the propositions of u_i would consist in.

Simplifying further, let's assume that we are able to generate u_i from a finite set P_{u_i} of propositions, where each $p \in P_{u_i}$ is in *Conjunctive Normal Form* (CNF). A proposition in CNF is a conjunction of disjunctions of literals (elementary propositional variables or their negations). The propositions in P_{u_i} can be conjoined in a single formula p_{u_i} that is itself in CNF. For p_{u_i} to hold it is necessary to determine a distribution of truth-values for its literals that renders the entire formula true. Determining the complexity of this satisfaction problem is an instance of the kSAT problem, where k is the number of literals in p_{u_i}. If $3 \leq k$, then the satisfiability problem for p_{u_i} is, in the general case, NP-complete, and so intractable.[3]

[2] In fact [5, 23, 27] originally characterised worlds as maximal consistent sets of propositions.

[3] See [39] for a discussion of the complexity properties of kSAT classes.

Given that this formula is intended to express the finite core of propositions from which the entire ultrafilter u_i is derived, it is not plausible to limit it to two literals, and it is reasonable to allow it to contain a large number of distinct elementary propositional constituents, each corresponding to a "core" fact that holds in u_i. It will also be necessary to include law like statements expressing regular relations among events that hold in a world (such as the laws of physics). These will be expressed as conditionals $A \rightarrow B$, which are encoded in a CNF formula by disjunctions of the form $\neg A \vee B$.

Therefore, even given the generous simplifying assumptions that we have made concerning the enumeration of u_i, specifying the ultrafilter of propositions that corresponds to an individual world is, in general, a computationally intractable problem. It follows that it is not possible to compute the set of worlds W efficiently.[4]

There are (at least) three ways in which one might try to evade this problem. First, we could follow Montague in claiming that formal semantics is a branch of mathematics rather than psychology. It involves the application of model theory, or, on the perspective adopted here, algebraic, and specifically, lattice theoretic methods, to develop formal models of meaning in natural languages. If this is the case, questions of efficient computability and representability are not relevant to the theoretical constructions that it employs. This move raises the obvious question of what formal semantics is explaining. If it seeks to account for the way in which people interpret the expressions of a natural language, then one cannot simply discard issues of representation. To do so is to ignore the cognitive aspect of meaning, which risks eliminating the empirical basis for assessing semantic theories.

A weaker form of this approach acknowledges that using and interpreting natural language is indeed a cognitive process, but it invokes the competence-performance distinction to insulate formal semantic theory from computational and processing concerns. On this view formal semantics offers a theory of semantic competence, which underlies speakers' linguistic performance.

[40] seems to suggest a move of this kind in distinguishing between semantic and psychological facts. But this is simply a version of the competence-performance distinction applied to semantics. Interestingly, this distinction is not generally regarded as granting immunity from the requirement of tractable representation in other areas of linguistic representation. So, for example, if a class of grammars (more precisely, the languages that they generate) is shown to be intractable for the recognition/parsing task, it is generally regarded as unsuitable for encoding the syntax of a natural language. Consequently, the full

[4] One might seek to treat propositions as unstructured, and worlds as ontologically primitive. It is unclear how either move could alleviate the representability problem. Literals are unstructured in the sense that they are elementary propositional variable or their negations. To banish the additional logical structure necessary to construct propositions in CNF would deprive propositions of any content at all. Taking worlds as primitive begs the question of how we identify and distinguish them. The conclusion that there is a one-to-one correspondence between a world and the ultrafilter of propositions that hold in it seems inescapable.

class of Context Sensitive Grammars, which, in some cases, require exponential time to decide membership in a context sensitive language, is regarded as too powerful to model NL syntax. Instead, the weaker subclass of Mildly Context Sensitive Grammars, for which the recognition problem is polynomial, is preferred. Consistency requires that tractability of representation also apply to semantic theories, even when these are taken to be abstract models of linguistic competence.

The difficulty here is that unless one provides an explicit account of the way in which competence drives processing and behaviour, then the distinction becomes vacuous. The notion of competence remains devoid of explanatory content.[5] We cannot simply set aside questions of effective computability if we are interested in semantic theories that are grounded on sound cognitive foundations.

A second strategy for dealing with the representability problem for possible worlds is to invoke the method of stratification. This technique stratifies a class of intractable problems into subclasses in order to identify the largest subsets of tractable tasks within the larger set.[6] So, for example, work on the tractable subclasses of kSAT problems is an active area of research. Similarly, first-order logic is undecidable (the set of its theorems is recursively enumerable, but the set of its non-theorems is not). However, many efficient theorem provers have been developed for subsets of first-order logic that are tractably decidable. We could focus on identifying the largest subsets of each $u_i \in U$ that can be tractably specified.

The problem with using stratification here is that, by definition, a world is (corresponds to) a maximal set of consistent propositions, an ultrafilter in a prelattice. If we specify only a proper subset of such an ultrafilter (a non-maximal filter), then it is not a world in the intended sense. It is no longer identified by all and only the propositions that hold at that world. In fact, in principle, several distinct worlds could share the same set of efficiently representable subsets of propositions, in which case they would not be efficiently distinguishable.[7]

Note that one cannot avoid this problem by claiming that, in principle, a "clever" algorithm could be devised to identify the ultrafilter of propositions that corresponds to a world. Unless one specifies such a procedure and shows that it efficiently identifies the set of worlds needed for a semantic theory, asserting

[5] See [33] for a detailed critical discussion of the difficulties raised by using the competence-performance distinction to protect syntactic theories from responsibility for handling a wide range of observed phenomena concerning speakers' syntactic judgments.

[6] See [8] on stratification of classes of grammars as a way of dealing with complexity in the context of computational learning theory for natural language grammars.

[7] [44] seems to have partial worlds in mind when he characterises worlds as elements in a partition of logical space, where such partitions are dependent on context. The problem with Stalnaker's suggestion is that he does not provide procedures for identifying partitions in logical space or their elements. In the absence of these it is not clear how such worlds/possibilities are to be represented or enumerated. Therefore, it does not offer a solution to the representability problem.

the mere possibility that one might be devised adds nothing of substance to the discussion.

Finally, a third approach to the problem of representability is to substitute possible situations for possible worlds. As situations are partial worlds, one may think that they are easier to represent. This is indeed be the case for individual situations, which are non-maximal, and for certain sets of situations.[8] However, it is not the case for the complete set of possible situations.[9]

For any given u_i corresponding to a world w_i, a situation $s_i \subseteq u_i$. The set of situations S_i for u_i is $\mathcal{P}(u_i)$, the power set of u_i. As $|u_i| = \aleph_0$, by Cantor's theorem on the cardinality of power sets, $|S_i|$ is uncountably infinite. Therefore S_i is not recursively enumerable. The set of all possible situations $S = \bigcup S_i$, and S inherits non-recursive enumerability from its constituent S_is. The representability problem for the set of possible situations is, then, even more severe than the one that we encounter for the set of possible worlds.

It may be possible to avoid this difficulty if we do not invoke the entire set of possible situations, but limit ourselves to subsets that we can specify effectively as we require them for particular analyses. This is, in effect, a form of stratification. But as situations are not maximal in the way that worlds are, it might be a viable method when applied to situations. In order for this method to work, it is necessary to show that we do, in fact, have effective procedures for representing the situations that we need for our theories. I will explore this approach in greater detail in Section 4.

3 Operational and Denotational Meaning

In the formal characterisation of programming languages it is common to distinguish between the operational and the denotational semantics of a program.[10] Operational meaning corresponds (roughly) to the sequence of state transitions that occur when a program is executed. It can be identified with the computational process through which the program produces an output for a specified input. The denotational meaning of a program is the mathematical object that represents the output which it generates for a given input. The operational and denotational meanings of the constituents of a program can be understood compositionally in terms of the contributions that they make to determining the state transitions performed by the program, and the value that it yields, respectively.

We can illustrate this distinction with two simple examples. First, it is possible to construct a theorem prover for first-order logic using either semantic tableaux or resolution.[11] Both theorem provers use proof by contradiction, but

[8] See [2] for the basic ideas of situation semantics.

[9] [21, 29, 26], for example, use the set of possible situations instead of the set of possible worlds to develop intensional semantic analyses.

[10] See, for example, [46] on these two types of meaning for expressions in programming languages.

[11] See [4] for tableaux and resolution theorem provers implemented in Prolog, and applied as part of a computational semantic system for natural language.

they employ alternative formal methods, and they are implemented as different computational procedures. They exhibit distinct efficiency and complexity properties. Consider the two predicates $TheoremTableaux$, which is true of the elements of the set of classical first-order theorems that a tableaux theorem prover produces, and $TheoremResolution$ that is true of the members of the set of classical first-order theorems that a resolution prover identifies. The predicates are intensionally distinct, but they are provably equivalent in their extensions.

The second example involves two functions from fundamental frequencies to the letters indicating musical notes and half tones. The first takes as its arguments the pitch frequency waves of the electronic sensor in a chromatic tuner. The second has as its domain the pitch frequency graphs of a spectrogram. Assume that both functions can recognise notes and half tones in the same range of octaves, to the same level of accuracy. Again, their operational semantics are distinct, but they are denotationally equivalent. The pairs of corresponding classifier predicates for these functions, $\langle A_{ChromTuner}, A_{SpecGram}\rangle$, $\langle A\#_{ChromTuner}, A\#_{SpecGram}\rangle$, $\ldots$, $\langle G_{ChromTuner}, G_{SpecGram}\rangle$, are intensionally distinct but denotationally equivalent. Both classifiers in a pair select the same set of notes, each through a different method.

We can apply this distinction to natural languages by taking the operational meaning of an expression to be the computational process through which speakers compute its extension, and its denotational meaning to be the extension that it generates for a given argument. We identify the intension of an expression with its operational meaning. This view of intension avoids the intractability of representation problem that arises with possible worlds.

It is also allows us to solve the difficulty of fine-grained intensionality (sometimes referred to as hyperintensionality). This issue arises because logically equivalent expressions are not, in general, inter-substituable in all contexts in a way that preserves the truth-value of the matrix sentence in which the expressions are exchanged. But if logically equivalent expressions have the same denotations in all possible worlds and intensions are functions from worlds to denotations, then these expressions are identical in intension. The following example illustrates the problem.

(1) a. If $A \subseteq B$ and $B \subseteq A$, then $A = B$. $\Leftrightarrow$
    b. A prime number is divisible only by itself and 1.

(2) a. Mary believes that if $A \subseteq B$ and $B \subseteq A$, then $A = B$. $\not\Leftrightarrow$
    b. Mary believes that a prime number is divisible only by itself and 1.

(1)a and b are both mathematical truths, but they are not inter-substitutable in the complement of *Mary believes that* in (2). However, if we identify intensions with operational meaning, then (1)a and b are intensionally distinct. (1)a is a theorem of set theory, while (1)b is a theorem of number theory. Their proofs are entirely different, and so they encode distinct objects of belief. The operational notion of intension permits us to individuate objects of propositional attitude with the necessary degree of fine-grained meaning.[12]

This solution to the issue of hyperintensionality is a secondary consequence of the operational account of intensions. Its primary motivation is to avoid the representability problem posed by possible worlds. [38] and [36] suggest related solutions, which retain possible worlds. See [30] for discussion of these proposals.

We have eliminated the dependence of intensions on possible worlds, and with it the representability problem for meanings, to the extent that the interpretation of an expression can be expressed as a procedure for computing its denotation. However, this only takes us part of the way to solving the cognitive plausibility problem for natural language semantics. We still need to develop an approach to modality and epistemic states which does not require possible worlds.

4 Modality and Epistemic States

Consider the following modal statements.

(3) a. Necessarily if $A \subseteq B$ and $B \subseteq A$, then $A = B$.
  b. Possibly interest rates will rise in the next quarter.
  c. It is likely that the Social Democrats will win the next election in Sweden.

In possible worlds semantics modal operators are construed as generalised quantifiers (GQs) on worlds. Necessity is a universal quantifier, possibility an existential quantifier, while *likely* is a variant of the second-order GQ *most*.[13] Let α, β, γ be the propositions to which the modal adverbs *necessarily, possibly* and *likely* apply in (3)a-c, respectively. The truth conditions of the sentences in (3) would be given by (something like) the following.

(4) a. $\|\Box\alpha\|^{M,w_i} = t$ iff $\forall w_{\in W} \|\alpha\|^{M,w} = t$.
  b. $\|\Diamond\beta\|^{M,w_i} = t$ iff $\exists w_{\in W} \|\beta\|^{M,w} = t$.
  c. $\|Likely\ \gamma\|^{M,w_i} = t$ iff for an appropriately defined $W' \subseteq W$, $|\{w_j \in W' : \|\gamma\|^{M,w_j} = t\}| \geq \epsilon$, where ϵ is a parametric cardinality value that is greater than 50% of W'.

[12] [11], Chapter 6 proposes an account of modality and propositional attitudes which dispenses with possible worlds, within the framework of Type Theory with Records, an intensional theory of types as judgements classifying situations. Some of Cooper's suggestions run parallel to the account proposed here. However, it is not clear how TTR solves the problem of complexity in representing the full set of record types. Moreover, it is not obvious that type membership in TTR is decidable.

[13] [25] presents a treatment of modalised degree modifiers that posits an ordering of possible worlds for similarity to a normative world. [22] discuss problems with this account and offer an alternative, which uses epistemically possible worlds. Given their reliance on a classical notion of possible world, neither theory avoids the representability problem.

On an alternative approach, we can reformulate modal statements as types of probability judgments. As a prelude it will be useful to review some basic ideas of probability theory.[14] A probability model M consists of a sample space of events with all possible outcomes given, and a probability distribution over these outcomes, specified by a function p. So, for example, a model of the throws of a die assigns probabilities to each of its six sides landing up. If the die is not biased towards one or more sides, the probability function will assign equal probability to each of these outcomes, with the values of the sides summing to 1.

Probability theorists often refer to the set of possible outcomes in a sample space as possible worlds. In fact this is misleading. Unlike worlds in Kripke frame semantics, outcomes are non-maximal. They are more naturally described as situations, which can be as large or as small as required by the sample space of a model. Therefore, in specifying a sample space it is not necessary to distribute probability over the set of all possible situations. In fact one need not even represent all possible situations of a particular type. One estimates the likelihood of an event of a particular type on the basis of observed occurrences of events, either of this type, or of others that might condition it. If we are working with Bayesian models, then we compute the posterior probability of an event A (the hypothesis) given observed events B (the evidence) with Bayes' Rule, where $p(B) \neq 0$.

(5) $p(A|B) = \dfrac{p(B|A)p(A)}{p(B)}$

Computing the full set of such joint probability assignments is, in the general case, intractable. However, there are efficient ways of estimating or approximating them within a Bayesian network.[15] It is, then, possible to efficiently represent a large subset of probability models, and to compute probability distributions for the possible events in their sample spaces.

Returning to the modal statements in 3, we can construct the following alternatives to 4, where M is a probability model, and p is the probability function in M.

(6) a. $\|Necessarily\ \alpha\|^{M,p} = t$ iff for all models $M' \in R, p_{\in M'}(\alpha) = 1$, where R is a suitably restricted subset of probability models.
b. $\|Possibly\ \beta\|^{M,p} = t$ iff $p(\beta) > 0$.
c. $\|Likely\ \gamma\|^{M,p} = t$ iff $p(\gamma) > \epsilon$, where ϵ is a parametric probability value that is greater than 0.5.

(6)a expresses universal necessity. Notice that to demonstrate this necessity it is sufficient to prove that assuming a probability model $M' \in R$ in which $p(\alpha) \neq 1$ produces a contradiction. If we are limiting ourselves to an appropriate

[14] See [20] for a particularly clear introduction to probability theory, that is relevant to some of the issues discussed here.
[15] See [41, 37, 20, 24] on Bayesian networks.

class of probability assertions and models, an efficient theorem prover may be available for such a result.[16] (6)b identifies possibility in a model with non-nil probability of occurrence. (6)c characterises likelihood in a model with a high degree of probability. These probabilistic characterisations of the modal adverbs *necessarily, possibly* and *likely* do seem to identify core aspects of their meanings in many of their common uses.[17]

In general we may use stratification to identify classes of probability models that can be efficiently represented, and we might invoke approximation techniques to estimate at least some of the others which are not. This is in contrast to individual worlds and sets of worlds. The maximality of worlds and the absence of any apparent procedure for generating their representations seem to exclude the application of these methods to possible worlds of the kind that figure in the formal semantics of natural language.

Let's consider how we might extend the probability-based approach proposed here for modality to epistemic states. Within a possible worlds framework knowledge and belief have traditionally been characterised along the following lines. Let W_B be the set of worlds (understood as ultrafilters of propositions) compatible with an agent a's beliefs. Take F_B to be a possibly non-maximal filter such that $F_B \subseteq \bigcap W_B$, where for every proposition $\phi \in F_B$, a regards ϕ as true. Let w_{actual} be the actual world. a's knowledge is contained in $F_K \subseteq F_B \cap w_{actual}$.[18]

As an alternative to this account we can use a probability model to encode an agent's beliefs. The probability distribution that this model contains expresses the agent's epistemic commitments concerning the likelihood of situations and events. One way of articulating the structure of causal dependencies implicit in these beliefs is to use a Bayesian network as a model of belief.[19]

---

[16] I am grateful to Robin Cooper for correcting a mistake in an earlier version of (6)a. One might be tempted to think that (6)a expresses a metaphysical concept of necessity, while (6)b,c correspond to epistemic modalities. In fact this is not the case. (6)a characterises necessity as a generalised quantifier over a suitably restricted set of probability models, each of which specifies a probability distribution over a number of events. These distributions constitute an agents' perception of the likelihood of certain events in the world. Therefore (6)a is not less of an epistemic specification of modality than (6)b,c.

[17] In order for this approach to modality to succeed, it will be necessary to develop accounts of the full class of modal expressions, including auxiliary verbs, other modal adverbs, and a variety of modal modifiers within the framework presented here. This is an important task for future work, but it is well beyond the scope of this paper. My objective here is programatic. I wish to show the viability of a probabilistic view of modality as an alternative to the traditional possible worlds treatment. Therefore, I have limited myself to the modal expressions that have been highlighted in the classical theories.

[18] [19] presents a version of this view.

[19] [34] considers the connection between conditional statements of the form $A \rightarrow B$ and the conditional probability $p(B|A)$. While this is an important issue, it is tangential to my concerns here. I am seeking a way of characterising epistemic states that does not invoke possible worlds.

Formally a Bayesian network is a Directed Acyclic Graph (DAG) whose nodes are random variables, each of whose values is the probability of one of the set of possible states that the variable denotes. Its directed edges express dependency relations among the variables. When the values of all the variables are specified, the graph describes a complete joint probability distribution (JPD) for its random variables.

The Bayesian network given in Fig 1, from [43], contains only boolean random variables, whose values are T (true) and F (false). In general, a discrete random variable X may have values $X_1, \ldots, X_n$ for any $n > 1$. Random variables may also be continuous.

The values of the instances of a variable depend directly only on the value of the variable of its parent. The dependency of a variable V on a non-parent ancestor variable A is mediated through a sequence of dependencies on the variables in the the path from V to A.

The only observable event for the network in Fig 1 is if the weather is cloudy or not, and the variable whose probability value we seek to determine is the likelihood of the grass being wet. We do not know the values of the random variables corresponding to rain, and to the sprinkler being on. Both of these events depend on whether the weather is cloudy, and both will influence the probability of the grass being wet. Sample conditional probabilities are given for each variable at each node of the network. The probability of the event C (cloudy) corresponding to the variable at the root of the graph is not conditioned, and its T and F instances are given equal likelihood.

We can compute the marginal probability of the grass being wet (W = T) by marginalising out the probabilities of the other variables on which W conditionally depends, either directly, or through intermediate variables. As we have seen, this involves summing across all the joint probabilities of their instances.

(7) $p(W = T) = \sum_{s,r,c} p(W = T, S = s, R = r, C = c)$

As we have a complete JPD for the variables of this network, it is straightforward to compute $p(W = T)$ using the chain rule for joint probabilities, together with the independence assumptions encoded in the network, which gives us (8).

(8) $p(W = T) = \sum_{s,r,c} p(W = T|S = s, R = r)p(S = s|C = c)p(R = r|C = c)p(C = c)$

In principle we could model an agent's beliefs as a single integrated Bayesian network. This would be inefficient, as it would be problematic to determine the dependencies among all of the random variables representing event types that the agent has beliefs about, in a way that sustains consistency. Moreover, the complexity involved in determining the conditional probabilities for the instances of each variable in such a global network would be daunting. It is more computationally manageable, and more epistemically plausible to construct local Bayesian networks to encode an agent's a's beliefs about a particular domain of situations. A complete collection of beliefs for a will consist of a set of such local

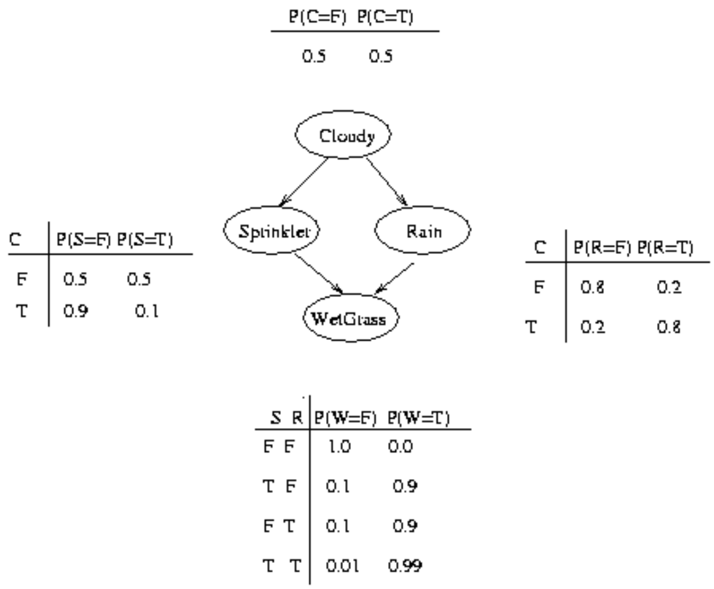

Fig. 1. Example of a Bayesian Network ([43])

networks, where each element of this set expresses a's beliefs about a specified class of events.

Two graphs G_i and G_j are *isomorphic* iff they contain the same number of vertices, and there is a bijection from the vertices of G_i to the vertices of G_j and vice versa, such that the same number of edges connect each vertex v_i to G_i and v_j to G_j, through identical corresponding paths.[20] For isomorphic DAGs this condition entails that the edges going into v_i and coming from it are of the same directionality as the edges going into and coming out of v_j, and vice versa.

Let's say that two subgraphs of two Bayesian networks *match* iff they are isomorphic, and the random variables at their corresponding vertices range over the same event instances, with the same probability values. Let BN_B be the

[20] [1] presents an algorithm for solving the graph isomorphism problem in quasi-polynomial time. An error was discovered in Babai's proof for this result. He subsequently repaired the proof in 2017, and posted the fix on his personal website at `http://people.cs.uchicago.edu/~laci/`.

Bayesian network that expresses a's beliefs about a given event domain. Take BN_R to be the Bayesian network that codifies the actual probabilities and causal dependencies that hold for these events.

We can identify a's knowledge for this domain as the maximal subgraph BN_K of BN_B that matches a subgraph in BN_R, and which satisfies additional conditions C. These conditions will enforce constraints like the requirement that the beliefs encoded in BN_B are warranted by appropriate evidence. Notice that on this characterisation of knowledge, if a knows ϕ, then a believes ϕ, but of course the converse does not hold. C can be formulated to permit justified true belief to count as knowledge, or it can be strengthened to block this implication.[21]

By characterising knowledge and belief in terms of Bayesian networks we avoid the representability problem that traditional analyses inherit from possible worlds. The proposed account offers two additional advantages. First, it exhibits the acquisition of beliefs as a dynamic process driven by continual updates in an epistemic agent's observations. This flow of new information generates revised probability distributions over the instances of the random variables in a network. Belief revision has to be handled by a task specific update function in a classical worlds based model of belief. It is intrinsic to Bayesian networks.

Second, a Bayesian network generates causal inferences directly, through the dependencies that it encodes in its paths. In a traditional worlds model of epistemic states, inference depends on an epistemic logic, whose rules are added to the model. By contrast, in a Bayesian network BN inference follows from the probability theory that BN instantiates. The network is both a dynamic model of belief, and a system that supports epistemic inference.

5 Related Work

[47] propose a theory in which probability is distributed over the set of possible worlds. The probability of a sentence is the sum of the probability values of the worlds in which it is true. If these worlds are construed as maximal in the sense discussed here, then this proposal runs into the representability problem for worlds.

[9, 10] develop a compositional semantics within a probabilistic type theory (ProbTTR). On their approach the probability of a sentence is a judgment on the likelihood that a given situation is of a particular type, specified in terms of ProbTTR. They also sketch a Bayesian treatment of semantic learning. It is not entirely clear how probabilities for sentences are computed in their system. They do not offer an explicit treatment of vagueness or probabilistic inference. It is also not obvious that their type theory is relevant to a viable compositional probabilistic semantics.

[16, 32] propose a probabilistic view of natural language semantics and pragmatics. They take probability to be distributed over partial worlds. They do not make entirely clear the relationship between partial and complete worlds.

[21] The claim that knowledge is justified true belief has been controversial at least since [15].

They also do not address the complexity issues involved in specifying worlds, partial or complete, as well as probability models. They implement probabilistic treatments of a scalar adjective, *tall*, and the sorities paradox for nouns like *heap* in the functional probabilistic programming language Church. Their analyses require a considerable amount of lexically specified content, and detailed information concerning speakers' and hearers' contextual knowledge. While their analyses offer thoughtful and promising suggestions on how to treat meaning in probabilistic terms, It is not obvious how their approach can be generalised to a robustly wide coverage model of combinatorial semantics and interpretation for natural language.

In addition, the Goodman-Lassiter account models vagueness by positing the existence of a univocal speaker's meaning that hearers estimate through distributing probability among alternative possible readings. They posit a boundary cut off point parameter for graded modifiers, where the value of this parameter is determined in context.

The approach that I am suggesting here is not forced to assume such an inaccessible boundary point for predicates. It allows us to interpret the probability value of a sentence as the likelihood that a competent speaker would endorse an assertion, given certain conditions (hypotheses). Therefore, predication remains intrinsically vague. It consists in applying a classifier to new instances on the basis of supervised training. We are not obliged to posit a contextually dependent cut off boundary for graded predicates.

[3] propose a compositional Bayesian semantics of natural language that implements this approach in a functional probabilistic programming language. It generates probability models that satisfy a set of specified constraints, and it uses Markov Chain Monte Carlo sampling to estimate the likelihood of a sentence being true in these models. It also sketches an account of semantic learning.

6 Conclusions and Future Work

I have argued that the tradition of formal semantics which uses possible worlds to model intensions, modality, and epistemic states is not built on cognitively viable foundations. Possible worlds of the kind posited in Kripke frame semantics are not tractably representable. Therefore, theories that rely on such a framework cannot explain the processes through which speakers actually interpret the expressions of a natural language. They also do not provide computationally manageable accounts of the ways in which epistemic agents reason about modality, knowledge and beliefs.

We have seen that by adapting the distinction between operational and denotation semantics from programming languages to natural language it is possible to develop a fine-grained treatment of intensions that dispenses with possible worlds. The intension of an expression is its operational meaning. Two expressions can have different intensions but provably equivalent denotations.

We replace Kripke frame semantics with probability models in order to interpret modal expressions, and we use Bayesian networks to encode knowledge, belief, and inference. While probability distributions, and Bayesian networks in particular, pose tractability problems, stratification, estimation, and approximation techniques allow us to effectively represent significant subclasses of these models. Therefore they offer a computationally realistic basis for handling epistemic states and inference.

If the approach that I have suggested here is to offer an interesting alternative to possible worlds semantics, then it will have to integrate the operational view of intensions into the probabilistic treatment of knowledge and belief. Specifically, it must explain how intensions are acquired by the sort of learning processes that are expressed in Bayesian networks.

In addition, it must develop a wide coverage system that combines a compositional semantics with a procedure for generating probability models in which it is possible to sample a large number of predicates. [3] provide an initial prototype for this system. Much work remains to be done on both the compositional semantics and the model testing components in order to create a robust Bayesian framework for natural language interpretation.

References

1. Babai, L.: Graph isomorphism in quasipolynomial time (2016), arXiv:1512.03547
2. Barwise, J., Perry, J.: Situations and Attitudes. MIT Press (Bradford Books), Cambridge, MA (1983)
3. Bernardy, J.P., Blanck, R., Chatzikyriakidis, S., Lappin, S.: A compositional bayesian semantics for natural language. In: Proceedings of the International Workshop on Language, Cognition and Computational Models, COLING 2018, Santa Fe, New Mexico (2018)
4. Blackburn, P., Bos, J.: Representation and Inference for Natural Language. CSLI, Stanford (2003)
5. Carnap, R.: Meaning and Necessity. University of Chicago Press, Chicago (1947)
6. Ciardelli, I., Groenendijk, J., Roelofsen, F.: Inquisitive semantics: A new notion of meaning. Compass **7(9)**, 459–476 (2013)
7. Ciardelli, I., Roelofsen, F., Theiler, N.: Composing alternatives. Linguistics and Philosophy **40**, 1–36 (2017)
8. Clark, A., Lappin, S.: Linguistic Nativism and the Poverty of the Stimulus. Wiley-Blackwell, Chichester, West Sussex, and Malden, MA (2011)
9. Cooper, R., Dobnik, S., Lappin, S., Larsson, S.: A probabilistic rich type theory for semantic interpretation. In: Proceedings of the EACL 2014 Workshop on Type Theory and Natural Language Semantics (TTNLS), pp. 72–79. Association of Computational Linguistics, Gothenburg, Sweden (2014)
10. Cooper, R., Dobnik, S., Lappin, S., Larsson, S.: Probabilistic type theory and natural language semantics. Linguistic Issues in Language Technology **10**, 1–43 (2015)
11. Cooper, R.: Type Theory and Language: From perception to linguistic communication. Unpublished ms, November 2016, University of Gothenburg (2016)
12. Fox, C., Lappin, S.: Type-theoretic logic with an operational account of intensionality. Synthèse **191(3)**, 563–584 (2015)
13. Fox, C., Lappin, S., Pollard, C.: A higher-order, fine-grained logic for intensional semantics. In: Alberti, G., Balough, K., Dekker, P. (eds.) Proceedings of the Seventh Symposium for Logic and Language. pp. 37–46. Pecs, Hungary (2002)

14. Fox, C., Lappin, S.: Formal Foundations of Intensional Semantics. Blackwell, Oxford (2005)
15. Gettier, E.L.: Is justified true belief knowledge? Analysis **23(6)**, 121–123 (1963)
16. Goodman, N., Lassiter, D.: Probabilistic semantics and pragmatics: Uncertainty in language and thought. In: Lappin, S., Fox, C. (eds.) The Handbook of Contemporary Semantic Theory, Second Edition, pp. 655–686. Wiley-Blackwell, Oxford (2015)
17. Groenendijk, J., Stokhof, M.: Dynamic Montague grammar. In: Kálmán, L., Polós, L. (eds.) Papers from the Second Symposium on Logic and Language. pp. 3–48. Akadémiai Kiadó, Budapest, Hungary (1990), also available as ITLI Prepublication Series LP–90–02
18. Groenendijk, J., Stokhof, M.: Dynamic predicate logic. Linguistics and Philosophy **14**(1), 39–100 (Feb 1991)
19. Halpern, J.: Should knowledge entail belief? Journal of Philosophical Logic **25**, 483–494 (1995)
20. Halpern, J.: Reasoning About Uncertainty. MIT Press, Cambridge MA (2003)
21. Heim, I.: E-type pronouns and donkey anaphora. Linguistics and Philosophy **13**, 137–177 (1990)
22. Holliday, W., Icard, T.: Measure semantics and qualitative semantics for epistemic modals. In: Proceedings of SALT 23. pp. 514–534 (2013)
23. Jonsson, B., Tarski, A.: Boolean algebras with operators. American Journal of Mathematics **73**, 891–939 (1951)
24. Koski, T., Noble, J.: Bayesian Networks: An Introductions. Wiley, Hoboken, NJ (2009)
25. Kratzer, A.: Modality. In: von Stechow, A., Wunderlich, D. (eds.) Semantics: An International Handbook of Contemporary Research, pp. 639–650. de Gruyter, Berlin (1991)
26. Kratzer, A.: Situations in natural language semantics. In: Stanford Encylopedia of Philosophy. Stanford University (2014)
27. Kripke, S.: A completeness theorem in modal logic. Journal of Symbolic Logic **24**, 1–14 (1959)
28. Kripke, S.: Semantical considerations on modal logic. Acta Philosophica Fennica **16**, 83–89 (1963)
29. Lappin, S.: An intensional parametric semantics for vague quantifiers. Linguistics and Philosophy **23**, 599–620 (2000)
30. Lappin, S.: Intensions as computable functions. Linguistic Issues in Language Technology **9**, 1–12 (2013)
31. Lappin, S.: Curry typing, polymorphism, and fine-grained intensionality. In: Lappin, S., Fox, C. (eds.) The Handbook of Contemporary Semantic Theory, Second Edition, pp. 408–428. Wiley-Blackwell, Malden, MA and Oxford (2015)
32. Lassiter, D., Goodman, N.: Adjectival vagueness in a bayesian model of interpretation. Synthese **194**, 3801–3836 (2017)
33. Lau, J.H., Clark, A., Lappin, S.: Grammaticality, acceptability, and probability: A probabilistic view of linguistic knowledge. Cognitive Science pp. 1–40 (2016)
34. Lewis, D.: Probabilities of conditionals and conditionals of probabilities II. The Philosophical Review Review **95**(4), 581–589 (1986)
35. Montague, R.: Formal Philosophy: Selected Papers of Richard Montague. Yale University Press, New Haven, CT/London, UK (1974), edited with an introduction by R. H. Thomason
36. Moschovakis, Y.N.: A logical calculus of meaning and synonymy. Linguistics and Philosophy **29**, 27–89 (2006)

37. Murphy, K.: An introduction to graphical models (2001), unpublished ms, University of British Columbia
38. Muskens, R.A.: Sense and the computation of reference. Linguistics and Philosophy **28**, 473–504 (2005)
39. Papadimitriou, C.: Computational Complexity. Addison-Wesley Publishing Co., Reading, MA (1995)
40. Partee, B.: Semantic facts and psychological facts. Mind and Language **2**(1), 43–52 (1988)
41. Pearl, J.: Bayesian decision methods. In: Shafer, G., Pearl, J. (eds.) Readings in Uncertain Reasoning, pp. 345–352. Morgan Kaufmann (1990)
42. Pollard, C.: Hyperintensions. Journal of Logic and Computation **18**, 257–282 (2008)
43. Russell, S., Norvig, P.: Artificial Intelligence: A Modern Approach. Prentice Hall Series in Artificial Intelligence, Englewood Cliffs, NJ (1995)
44. Stalnaker, R.: Indexical belief. Synthèse **49**(1), 129–151 (1981)
45. Stalnaker, R.: Inquiry. MIT Press/Bradford Books, Cambridge MA (1984)
46. Stump, A.: Programming Language Foundations. Wiley, Hoboken, NJ (2013)
47. van Eijck, J., Lappin, S.: Probabilistic semantics for natural language. In: Christoff, Z., Galeazzi, P., Gierasimszuk, N., Marcoci, A., Smets, S. (eds.) Logic and Interactive Rationality (LIRA), Volume 2. ILLC, University of Amsterdam (2012)