

Uniform Coherence

Scott Garrabrant^{1,2}, Benya Fallenstein¹, Abram Demski^{1,3}, and Nate Soares¹

¹Machine Intelligence Research Institute

²University of California, Los Angeles

³University of Southern California

Abstract

While probability theory is normally applied to external environments, there has been some recent interest in probabilistic modeling of the outputs of computations that are too expensive to run. Since mathematical logic is a powerful tool for reasoning about computer programs, we consider this problem from the perspective of integrating probability and logic. Recent work on assigning probabilities to mathematical statements has used the concept of *coherent* distributions, which satisfy logical constraints such as the probability of a sentence and its negation summing to one. Although there are algorithms which converge to a coherent probability distribution in the limit, this yields only weak guarantees about finite approximations of these distributions. In our setting, this is a significant limitation: Coherent distributions assign probability one to all statements provable in a specific logical theory, such as Peano Arithmetic, which can prove what the output of any terminating computation is; thus, a coherent distribution must assign probability one to the output of any terminating computation. To model uncertainty about computations, we propose to work with *approximations* to coherent distributions. We introduce *uniform coherence*, a strengthening of coherence that provides appropriate constraints on finite approximations, and propose an algorithm which satisfies this criterion.

1 Introduction

Recently there has been some interest in the problem of assigning probabilities to the outputs of computations which are too expensive to run. For example, [1] call for the development of probabilistic numerical methods that are uncertain about their calculations; [2] study metareasoning procedures for controlling Monte Carlo simulations by estimating which simulations are likely

Research supported by the Machine Intelligence Research Institute (intelligence.org).

to be effective; and [3] give methods for estimating probabilistic program variables via Bayesian optimization.

Formal logic is a tool that is particularly well-suited for making claims about computations, such as claims of the form “this computation will halt and produce a number larger than 7” or “this computation will use less memory than that computation” or “this operating system’s scheduler will not deadlock” [4, 5].

When developing methods for handling uncertainty about the results of computations, then, it is natural to approach the problem from the angle of combining logic with probability theory, a topic which has received a lot of attention; see [6]. Since we are using logic to reason about mathematical facts, as opposed to reasoning about an uncertain external world, the approach taken by [7], [8], [9], and others is particularly relevant. This approach involves assigning probabilities to logical sentences in a formal theory powerful enough to represent claims about computations, such as Peano Arithmetic (PA) or Zermelo-Fraenkel set theory (ZFC).

It is standard to study distributions P of this form which are *coherent*, in that they obey some simple logical constraints such as $P(\perp) = 0$ and $P(\phi) + P(\neg\phi) = 1$. Coherence can be thought of as a generalization of the notion of “consistency” to probability distributions over sentences in logic. However, coherent distributions are uncomputable—they assign probability 1 to all theorems, and thus, they must assign probability 1 to the statement “computation f outputs x ” whenever $f() = x$. For this reason, coherent probability distributions cannot represent a state of uncertainty about the outputs of computations. Rather, they represent a *final* state of knowledge about logical facts that a reasoner might obtain if they could think forever [10, 11].

[8] has proposed instead investigating algorithms that *computably approximate* a coherent probability distribution, that is, algorithms that output a series of probabilities for each sentence such that those probabilities converge in the limit, and such that the distribution the approximation converges to is coherent. Unfortunately, coherence in the limit is too weak for our purposes: It does not impose constraints on any individual finite approximation of the coherent distribution. For example, an approximation to a coherent distribution

might assign arbitrary probabilities to some theorem ϕ right up until it proves it, then assigns probability 1 thereafter. If ϕ was a statement about a computation, this means the approximation might only start assigning reasonable probabilities to ϕ after actually running the computation, which defeats the purpose. To get approximations of coherent distributions that assign reasonable probabilities to ϕ *before* running the computation, we require some notion like coherence that we can impose on *approximations* to the final distribution.

In this paper, we propose one such property, *uniform coherence*, argue that it is a desirable generalization of coherence to approximations of coherent distributions, and show that a uniformly coherent approximation scheme exists. Roughly speaking, uniform coherence demands that, for any pattern in what is provable that can be identified by a Turing machine in polynomial time, the approximation must recognize and exploit that pattern “not much later” than that Turing machine. We define this more formally in Section 2. In Section 3 we propose a modification of Demski’s algorithm [8] and show that it is uniformly coherent. We conclude with a discussion of open problems and future work.

Other Related Work The study of assigning probabilities to sentences in mathematical logic dates back to Łoś [12] and [7]; see [13] for a more thorough history. Since then, the idea has been extended to, e.g., infinitary logic [14], databases with uncertain data [15], and higher-order logic [9]. Computing probability assignments for logical statements can be seen as an extension of these approaches to the case where the reasoner’s beliefs may be incoherent; see also the work of [16] and [17], who study methods for measuring and handling inconsistency in knowledge bases.

Another method for unifying logic with probability is inductive logic programming [18, 19]. For example, [20] have described techniques for learning from proofs and program traces, and [21] proposes combining first-order logic and probabilistic graphical models using a method known as “Markov logic networks.” For a review of recent work, refer to [6].

Methods for reasoning probabilistically about the outputs of computations are a type learning scheme for probabilistic logic programs. For more on this subject, refer to [22], [23], [24], [25], [26], [27], [28], [29].

2 Uniform Coherence

In this paper we study probability distributions over sentences of first-order logic, with the goal of describing computable distributions that assign reasonable probabilities to claims about computations that are too expensive to be run. (For example, imagine a reasoner that wants to know whether a particular $O(n^2)$ computation will outperform a particular $O(n \log n)$ computation on a particular dataset, without taking the time

to run both computations.) We will fix a language \mathcal{L} and a theory T that is powerful enough to make claims about computations, such as PA.

“Coherence” can be seen as the natural generalization of consistency to probability distributions. Roughly, coherence demands that theorems be certain, contradictions be impossible, and the probabilities of mutually exclusive sentences add:

Definition 1 (Coherence). *A probability distribution over sentences in \mathcal{L} is a function $P : \mathcal{L} \rightarrow [0, 1]$ from sentences to probabilities. It is called coherent with respect to the \mathcal{L} -theory T if the following three conditions hold:*

1. If ϕ is a theorem of T , $P(\phi) = 1$.
2. If ϕ is contradictory, $P(\phi) = 0$.
3. For each ϕ and ψ , $P(\phi) = P(\phi \wedge \psi) + P(\phi \wedge \neg\psi)$.

It is not hard to see that coherence ensures P obeys other obvious logical constraints, such as $P(\phi) = 1 - P(\neg\phi)$ and $P(\phi \rightarrow \psi) = 1 \implies P(\psi) \geq P(\phi)$. We say that coherence is a generalization of consistency because P agrees with T on all theorems and contradictions, but can assign probabilities to undecidable sentences so long as those probabilities obey logical constraints. [7] has shown that any coherent P is isomorphic to a measure μ on complete consistent extensions of T —in other words, a coherent P assigns probability to undecidable sentences in a fashion that corresponds to choosing some weighted mixture of all possible consistent ways to assign truth values to sentences.

Coherence is a *very* strong constraint. For example, if T is PA, then P must assign probability 1 to all true statements about the behavior of computations. One way to think of a coherent distribution is that it represents the state of knowledge a reasoner could achieve after thinking *forever*, after proving everything provable, refuting everything refutable, and assigning consistent probabilities to everything undecidable. It seems reasonable to ask that the *limit* of a good reasoner’s beliefs about logical statements should be coherent, but no computable distribution can achieve coherence. This motivates the study of computable *approximation schemes* to coherent distributions, that is, algorithms which output a sequence of probabilities for each sentence such that the sequences converge, and the resulting probability assignments are coherent.

Definition 2 (Approximation scheme). *An approximation scheme is a Turing machine M which takes as input a natural number n and an encoding of a sentence ϕ , and outputs a rational number $M_n(\phi) \in \mathbb{Q} \cap [0, 1]$.*

We can think of M as a machine that runs for longer and longer as n gets larger, producing better and better estimates of the probability of ϕ each iteration. We concern ourselves with approximation schemes that converge to a coherent distribution in the limit.

Definition 3 (Approximation). An approximation scheme M approximates $P : \mathcal{L} \rightarrow [0, 1]$ if

$$\lim_{n \rightarrow \infty} M_n(\phi) = P(\phi) \quad (1)$$

for all ϕ . If there exists an M which approximates P , then we say that P is “approximable.”

Coherence in the limit does not impose strong constraints on an approximation scheme. Given a sentence ϕ describing a claim about a computation, M could simply assign it arbitrary probabilities right up until it can run the computation, and then assign it probability 1 or 0 thereafter. M could then be coherent in the limit, but it would never assign reasonable non-extreme probabilities to claims about computations too expensive to run. To get an M that accurately estimates the results of computations before it can run them, we need to impose stronger constraints on the approximations.

Clearly, we cannot demand that the approximate distributions M_n be completely coherent. We could try weakening coherence by demanding that M_n assign probability 1 to all theorems *that have a proof of length n or less*, and indeed, this is the approach taken by [8]. However, in our setting, this runs into the same problem as above: M could still assign arbitrary probabilities to a decidable ϕ right up until it proves ϕ one way or the other, and might not have anything reasonable to say about the behavior of computations before running them. We require some other weakening of coherence that ensures M places reasonable probabilities on ϕ *before* it can run the computation.

Intuitively, we want M to start noticing *patterns* in computations before it’s able to actually run them. For example, consider some computation f that takes an input and either outputs 1, outputs 0, or loops. Let ϕ_n be the claim $f(n) = 0$ and ψ_n be the claim $f(n) = 1$. We want M to start assigning probabilities to ϕ_n and ψ_n that sum to at most 1, and we want it to start doing so before it can simply compute $f(n)$.

More generally, we want M to recognize patterns such as “these claims are mutually exclusive” and “these claims are equivalent.” We formalize this idea as follows. Take any method for quickly (in polynomial time) identifying patterns of this form. A good approximation scheme should eventually be able to recognize and exploit that pattern, and “not much slower” than the polynomial time method, even if the size of the shortest proof that the pattern holds grows super-exponentially.

Definition 4 (Quickly computable). A sequence of sentences $\{\phi_n\}$ is called quickly computable if there exists a Turing machine that outputs all the sentences in the sequence in order, and outputs ϕ_n by a time polynomial in n .

Definition 5 (Uniform coherence). An approximation scheme M is called uniformly coherent with respect to T if it satisfies the following three conditions:

1. $\lim_{n \rightarrow \infty} M_n(\perp) = 0$
2. $\lim_{n \rightarrow \infty} M_n(\phi_n)$ converges whenever $\{\phi_n\}$ is quickly computable and $\phi_n \rightarrow \phi_{n+1}$ is provable in T for all n .
3. $\lim_{n \rightarrow \infty} M_n(\phi_n) + M_n(\psi_n) + M_n(\chi_n) = 1$ whenever $\{\phi_n\}$, $\{\psi_n\}$, and $\{\chi_n\}$ are quickly computable, and for all n , it is provable in T that $\{\phi_n, \psi_n, \chi_n\}$ is a partition of truth (meaning exactly one of them is true).

To gain an intuition for how Definition 5 guarantees that a uniformly coherent M recognizes patterns in quickly computable sequences “not much slower” than the Turing machine q that quickly computes them, imagine that $q(n)$ outputs pairs (ϕ_n, ψ_n) in time polynomial in n such that each ϕ_n is provably equivalent to ψ_n , but the shortest proof of equivalence grows super-exponentially fast in n . We want M to eventually, on some iteration n and thereafter, “recognize the pattern” and start assigning roughly the same probabilities to ϕ_n and ψ_n . Note that this is a claim about the iteration n by which M must start recognizing the pattern identified by q , not a claim about the runtime of $M_n(\phi)$ on an individual ϕ , which may be exponential or worse.

Theorem 1. If M is uniformly coherent, and $\{\phi_n\}$ and $\{\psi_n\}$ are quickly computable, and ϕ_n is provably equivalent to ψ_n for all n , then

$$\lim_{n \rightarrow \infty} M_n(\phi_n) - M_n(\psi_n) = 0. \quad (2)$$

Proof. If $\{\phi_n\}$ is quickly computable then so is $\{\neg\phi_n\}$. The constant sequence $\{\perp\}$ is quickly computable. Apply property 3 to the sequences $\{\phi_n\}$, $\{\neg\phi_n\}$, and $\{\perp\}$, and then to $\{\psi_n\}$, $\{\neg\phi_n\}$, $\{\perp\}$. Subtracting the results, we have

$$\lim_{n \rightarrow \infty} M_n(\phi_n) + M_n(\neg\phi_n) - M_n(\psi_n) - M_n(\neg\phi_n) = 0, \quad (3)$$

because $\lim_{n \rightarrow \infty} M_n(\perp) = 0$. \square

Provable equivalence is only one type of pattern that a uniformly coherent M exploits before it can find the associated proofs. As a second example, if there is any quickly computable method for identifying sentences that are provable (even if the proofs are very long), then M must recognize those patterns as well.

Theorem 2. If M is uniformly coherent, for any quickly computable sequence $\{\phi_n\}$ of theorems of T we have

$$\lim_{n \rightarrow \infty} M_n(\phi_n) = 1. \quad (4)$$

Proof. Apply property 3 to the sequences $\{\phi_n\}$, $\{\perp\}$, $\{\perp\}$. Since each ϕ_n is provable, we have partitions of truth as desired, and $\lim_{n \rightarrow \infty} M_n(\perp) = 0$. \square

This implies that the criterion of uniform coherence captures some of what it means to reason well about computations that are too expensive to run. Recall that statements making true claims about the output of a computation are theorems, because the statement can be proven by providing an execution trace of the computation. Theorem 2 implies that if there is a polynomial-time method for generating true statements about computations, then after some number of iterations n , M_n incorporates that pattern into its probability assignments.

Property 1 of Definition 5 is fairly trivial. Property 2 implies that $M_n(\phi_n)$ converges if $\{\phi_n\}$ is quickly computable and all the ϕ_n are provably equivalent; the more general form of Property 2 is used in Theorem 4. Property 3 is the meat of uniform coherence; it ensures M recognizes exclusivity relationships between provable sentences. While Definition 5 only mentions partitions of truth of size 3, it is strong enough to guarantee M recognizes arbitrarily large partitions of truth.

Theorem 3. *If M is uniformly coherent, for any list of k quickly computable sequences of sentences, $\{\phi_n^1\}, \dots, \{\phi_n^k\}$, such that for each n , it's provable that ϕ_n^i is true for exactly one i , we have that*

$$\lim_{n \rightarrow \infty} \sum_{i=1}^k M_n(\phi_n^i) = 1. \quad (5)$$

Proof. The proof works by induction on k . For $k = 1$, this is Theorem 2. For $k = 2$, apply property 3 to the sequences $\{\phi_n^1\}$, $\{\phi_n^2\}$, $\{\perp\}$. For $k = 3$, this is exactly property 3.

For $k > 3$, consider the list $\{\phi_n^1\}, \dots, \{\phi_n^{k-2}\}, \{\phi_n^{k-1} \vee \phi_n^k\}$; the list $\{\phi_n^{k-1}\}, \{\phi_n^k\}, \{\neg(\phi_n^{k-1} \vee \phi_n^k)\}$; and the list $\{\phi_n^{k-1} \vee \phi_n^k\}, \{\neg(\phi_n^{k-1} \vee \phi_n^k)\}$.

These lists are of length $k - 1$, 3, and 2 respectively and all satisfy the conditions of this theorem. Apply this theorem to all three lists, add the first two results and subtract the third. This gives $\lim_{n \rightarrow \infty} \sum_{i=1}^k M_n(\phi_n^i) = 1 + 1 - 1 = 1$. \square

Intuitively, a uniformly coherent M is good at identifying any pattern in what is provable that can be expressed using the properties of Definition 5. As an example, observe that if q can quickly compute an infinite sequence of provably mutually exclusive sentences, M must eventually start noticing that those sentences are mutually exclusive, not much later than q :

Theorem 4. *If M is uniformly coherent, then for a quickly computable sequence of mutually exclusive sentences, $\{\phi_n\}$, we have $\lim_{n \rightarrow \infty} M_n(\phi_n) = 0$.*

Proof. Define ψ_n to be the disjunction of all $\phi_{i \leq n}$. Applying property 2 to $\{\psi_n\}$, we have that $\lim_{n \rightarrow \infty} M_n(\psi_n)$ converges to some p . Applying property 3 to the sequences $\{\psi_n\}$, $\{\neg\psi_n\}$, and $\{\perp\}$, we have that $\lim_{n \rightarrow \infty} M_n(\neg\psi_n)$ converges to $1 - p$. Therefore,

applying property 3 to $\{\psi_{n-1}\}$, $\{\phi_n\}$, and $\{\neg\psi_n\}$ shows that $\lim_{n \rightarrow \infty} M_n(\phi_n) = 0$ as desired. (Note that $\{\psi_n\}$, $\{\neg\psi_n\}$, and $\{\psi_{n-1}\}$ are all quickly computable if $\{\phi_n\}$ is.) \square

Does a uniformly coherent M quickly identify *all* quickly identifiable patterns in claims about computations? Probably not; limitations are discussed in Section 4. However, we can show that M will *eventually* identify all patterns in which sentences are provable, by showing that a uniformly coherent M is coherent in the limit. This may be surprising at first glance, given that Definition 5 only mentions convergence for sequences that can be computed in polynomial time. The trick is that any constant sequence $\phi_n := \phi$ is “quickly computable,” by the Turing machine that ignores n and always outputs ϕ .

Theorem 5. *If M is uniformly coherent, then*

$$P(\phi) := \lim_{n \rightarrow \infty} M_n(\phi) \quad (6)$$

is well-defined, approximable, and coherent.

Proof. That $P(\phi)$ is well-defined follows from property 2 and the fact that the constant sequence $\{\phi\}$ is quickly computable for any ϕ . Approximability then follows trivially from the definition of $P(\phi)$. $P(\phi)$ is in $[0, 1]$ because $M_n(\phi)$ is in $[0, 1]$, and the limit of any sequence in $[0, 1]$ is in $[0, 1]$. It remains to show that P is coherent.

The first property of coherence follows from properties 1 and 3 with the partition $\{\phi, \perp, \perp\}$. The second follows with the partition $\{\phi, \perp, \top\}$. The third follows with the partition $\{\phi \wedge \psi, \phi \wedge \neg\psi, \neg\phi\}$, using the fact $P(\phi) = 1 - P(\neg\phi)$ which we get from property 2 and the partition $\{\phi, \neg\phi, \perp\}$. \square

Thus we see that we are justified in saying that a uniformly coherent M must both quickly identify some patterns in what is provable, and eventually identify all patterns. This implies a uniformly coherent M would assign reasonable probabilities to sentences describing the behavior of computations, even before the computation can be run, because statements about computations are decidable. If there is *any* polynomial-time method for accurately noticing relationships between computations, M will eventually recognize it and distribute its probability mass accordingly.

One way to look at uniform coherence is this: Coherence in the limit requires that each *individual sentence* is eventually assigned a reasonable probability, possibly only after that sentence is decided. Uniform coherence requires that there is some uniform bound past which *all* theorems that can be quickly identified as theorems start to be assigned high probability, as quickly as they can be identified.

3 A Uniformly Coherent Approximation Scheme

We now turn our attention to providing a uniformly coherent approximation scheme M^* which approximates a coherent distribution P^* . We do this by defining a variant of the distribution and approximation scheme defined by [8], and showing that our variant is uniformly coherent.

We begin by defining an uncomputable process (which we later show to be approximable) that builds a random complete, consistent extension T^* of T . We define $P^*(\phi)$ to be the probability that $\phi \in T^*$ if T^* is sampled according to this process. To build T^* , we sample random Turing machines according to a simplicity prior, and interpret their outputs as claims about which sentences are true. If the machine makes claims that are consistent with T^* so far, we add those claims to T^* and repeat. Otherwise we discard that machine and repeat. Continuing indefinitely, T^* will (with probability 1) be a complete consistent extension of T in the limit.

To formalize this idea, fix a universal Turing machine U with an advance-only output tape, with its input tape initialized to a random infinite bitstring b . We can interpret an infinite bitstring as a self-delimiting encoding of a Turing machine followed by a specification of the initial state of that machine's tape. Fix an enumeration of sentences in the language \mathcal{L} , and interpret the output of U on input tape b as a sequence of sentences. Write $U(b)$ for the set of sentences output by U on the input b , and write $U_t(b)$ for the finite set of sentences output during the first t steps of operation. For a finite bitstring x , write $U(x) = S$ if $U(b) = S$ for all b which have x as a prefix, and similarly for $U_t(x)$. Note that if the length of x is at least t , then $U_t(x)$ is always well-defined, because U cannot read more than t input bits in t steps. The (uncomputable) process converging on T^* is then defined by Algorithm 1.

Algorithm 1: A method for constructing a complete, consistent extension of T .

Input: b_1, b_2, \dots , an infinite list of infinite bitstrings.
Data: $T^* \leftarrow T$, the extension of T under construction.
for b **in** b_1, b_2, \dots :
 if $T^* \cup U(b)$ **is consistent**:
 $T^* \leftarrow T^* \cup U(b)$

We define $P^*(\phi)$ to be the probability that $\phi \in T^*$ when the b_i are chosen uniformly at random (e.g., by fair coin tosses; recall that a single stream of coin tosses can encode an infinite sequence of infinite bitstrings).

P^* has the desirable property that, for every noncontradictory sentence ϕ , $P^*(\phi)$ is lower-bounded by the complexity of the Turing machine that outputs only ϕ . To see this, let w_ϕ be the bitstring encoding that machine with respect to U ; the chance that b_1 starts with

w_ϕ is at least $2^{-|w_\phi|}$.

To see that $P^*(\phi)$ is coherent, note that with probability 1 T^* is a complete consistent extension of T , so P^* is isomorphic to a distribution μ over complete consistent extensions of T , which means it is coherent [7]. P^* is uncomputable, but can be approximated by Algorithm 2:

Algorithm 2: Computable approximation scheme for P^*

```

def  $\text{Claims}_n(b_1, \dots, b_{2^n})$ :
     $\Phi \leftarrow$  the first  $n$  axioms of  $T$ 
    for  $i$  in  $0 \dots 2^n$ :
         $S \leftarrow U_{2^n}(b_i)$  interpreted as a list of sentences
        if  $\text{Con}_n(\Phi \cup S)$ :
             $\Phi \leftarrow \Phi \cup S$ 
    for  $\phi$  in  $\Phi$ :
        output( $\phi$ )

def  $\text{Con}_n(\Phi)$ :
    if a proof of length  $\leq 2^n$  proves  $\Phi$  inconsistent:
        return false
    for  $S \subset \Phi$ :
        for  $\phi$  of length  $\leq 2^n$  such that  $S \cap \{\phi, \neg\phi\}$  is empty:
            if  $\neg\text{Con}_n(S \cup \{\phi\})$  and  $\neg\text{Con}_n(S \cup \{\neg\phi\})$ :
                return false
    return true

def  $M_n^*(\phi)$ :
    return the probability  $\text{Claims}_n(b_1, \dots, b_{2^n})$  outputs  $\phi$  when the  $b_i$  are uniform random bitstrings of length  $2^n$ .

```

Of note is the function Con_n , which checks whether a set of sentences Φ is “consistent enough” for time n . It checks not only whether Φ can be proven inconsistent with a proof of length 2^n or less, but also whether there is a subset $S \subset \Phi$ and sentence ϕ such that both $\Phi \cup \{\phi\}$ and $\Phi \cup \{\neg\phi\}$ can be proven inconsistent in length 2^n . (This implies that Φ is inconsistent, but the proof may be longer than length 2^n unless one of ϕ or $\neg\phi$ is added.) This gives Con_n a convenient closure property.

Theorem 6. M^* is a uniformly coherent approximation scheme which approximates P^* .

We show that M^* recognizes quickly computable theorems, which is suggestive. The rest of the proof is in Appendix A.

Lemma 1. If $\{\phi_n\}$ is a quickly computable sequence of theorems, then

$$\lim_{n \rightarrow \infty} M_n^*(\phi_n) = 1. \quad (7)$$

Proof. Write Φ_n for a random run of $\text{Claims}_n(b_1 \dots b_{2^n})$ when the b_i are chosen uniformly at random. We want to show that, for n large enough, $\mathbb{P}(\phi_n \in \Phi_n) > 1 - \varepsilon$. Since $\{\phi_n\}$ is quickly computable, there exists a finite bitstring w_ϕ such that $U(w_\phi)$ outputs the sentences $\{\phi_n\}$ in order, and for all sufficiently large n , $\phi_n \in U_{2^n}(w_\phi)$.

There exists a N_0 such that with probability at least $1 - \varepsilon/2$, at least one of $b_1 \dots b_{2^{N_0}}$ starts with w_ϕ , and $\phi_n \in U_{2^n}(w_\phi)$ for all $n \geq N_0$. There also exists a N_1 such that with probability $1 - \varepsilon/2$, for every subset S of $\{1, \dots, 2^{N_1}\}$, either $T \cup \bigcup_{i \in S} U(b_i)$ is consistent or $\neg \text{Con}_{N_1}(T \cup \bigcup_{i \in S} U_{2^{N_1}}(b_i))$; simply choose N_1 large enough that any inconsistencies can be uncovered with sentences output by time 2^{N_1} and proofs of length less than 2^{N_1} . (This is possible because only finitely many proofs of inconsistency are needed.)

Choose $N_1 \geq N_0$. For all $n \geq N_1$, with probability at least $1 - \varepsilon/2$, one of the sampled machines (namely w_ϕ) outputs ϕ_i for all $i \leq n$. Then, with probability at least $1 - \varepsilon/2$, this implies that this machine will end up contributing to Φ_n because any machine before w_ϕ inconsistent with w_ϕ (which outputs only theorems) has been discarded. Therefore, with probability at least $1 - \varepsilon$, $\phi_n \in \Phi_n$. \square

4 Conclusions

We have proposed *uniform coherence* as a strengthening of coherence in the limit. Uniform coherence requires that computable distributions assign probabilities to claims about computations that are reasonable before they're able to run these computations. Specifically, if there is any polynomial-time method for identifying patterns in what is provable, a uniformly coherent M must eventually recognize and exploit that pattern, eventually assigning probabilities that are coherent with respect to that pattern. This implies that uniform coherence captures some of what we mean when we ask for a probability distribution that assigns reasonable probabilities to claims about computations.

However, a uniformly coherent M doesn't necessarily recognize all patterns in the behavior of computations. For example, consider: Is the 10^{100} 'th decimal digit of π a 7? It seems that in lieu of additional knowledge and the ability to compute the digit, a reasonable estimator should assign this event 10% probability. Reasonable predictors of computations should be able to recognize similar patterns, such as "this computation returns an error one time in ten," and assign probabilities accordingly.

More formally, imagine we have some sequence of deterministic computations that output a one on $1/10$ of their inputs. Imagine further that there is no polynomial-time algorithm that has better average squared error, when predicting this sequence, than the algorithm "output $(1/10)$." It seems reasonable to ask that a predictor of computations start assigning probability $1/10$ to the next element in the sequence eventually, until it has enough resources to compute the actual answer. However, we have no reason to expect that a uniformly coherent M would have this property. [30] study computable distributions that can do this; it is not yet clear how to reconcile our framework with theirs.

This demonstrates that further constraints on ap-

proximation schemes are likely necessary before we can define computable distributions that are able to recognize all the patterns in the behavior of computations that humans can easily recognize. Uniform coherence gives us approximate distributions that have some desirable properties in their predictions about computations, but more work is needed before we can say we understand how to assign reasonable uncertainty to claims about computations in general.

A Proof that M^* is Uniformly Coherent

Let b_1, b_2, \dots be an infinite list of infinite bitstrings, generated uniformly at random. Let T^* be the complete extension of T that Algorithm 1 converges to on the input b_1, b_2, \dots . Let $\Phi_n = \text{Claims}_n(b_1, b_2, \dots)$. Note that Claims_n only reads the first 2^n bits of the first 2^n bitstrings.

Lemma 2. *If $\{\phi_n\}$ is a quickly computable sequence of sentences, then*

$$\lim_{n \rightarrow \infty} M_n^*(\phi_n) + M_n^*(\neg\phi_n) = 1. \quad (8)$$

Proof. Since $\{\phi_n\}$ is quickly computable, there is a prefix w_ϕ such that $U(w_\phi x)$ outputs ϕ_n in polynomial time if x encodes n . We can choose w_ϕ such that x encodes n in $2 \log_2(n)$ bits. Therefore, if x is an infinite uniform random bitstring, $U(w_\phi x)$ will output ϕ_n with probability at least $2^{-2 \log_2(n)} = n^{-2}$.

Similarly, there exists a finite bitstring w'_ϕ such that if x is a uniform random bitstring, $U(w'_\phi x)$ outputs the single sentence $\neg\phi_n$ in polynomial time with probability at least n^{-2} . If w_ϕ and w'_ϕ are each at most k bits long, the probability that $w_\phi x$ and $w'_\phi x$ are prefixes for some $b_i, b_j \in \{b_1, b_2, \dots, b_{2^n}\}$ is at least $1 - (1 - 2^{1-k} n^{-2})^{2^n}$. This converges to 1 as n goes to ∞ . Therefore, with probability converging to one, there exist b_i and b_j with $i, j \leq 2^n$ such that $U_{2^n}(b_i)$ is the singleton containing ϕ_n and $U_{2^n}(b_j)$ is the singleton containing $\neg\phi_n$. Without loss of generality, assume $i < j$. We want to show that if both of these sentences are sampled, exactly one of ϕ_n or $\neg\phi_n$ is in Φ_n .

Clearly, the set $\{\phi_n, \neg\phi_n\}$ and all its supersets will be rejected by Con_n , so ϕ_n and $\neg\phi_n$ are not both in Φ_n . Let Φ_n^k be the value of Φ in $\text{Claims}_n(b_1, b_2, \dots)$ after k iterations of the for loop. Assume that neither ϕ_n nor $\neg\phi_n$ are in Φ_n . Thus, Con_n rejects $\Phi_n^i \cup \{\phi_n\}$ and $\Phi_n^j \cup \{\neg\phi_n\}$. Adding more sentences to the input of Con_n cannot cause it to accept, so Con_n rejects $\Phi_n^j \cup \{\phi_n\}$ as well. Thus, Con_n accepts Φ_n^i but rejects $\Phi_n^j \cup \{\phi_n\}$ and $\Phi_n^j \cup \{\neg\phi_n\}$. This is a contradiction, so as n goes to ∞ , with probability approaching 1, exactly one of ϕ_n or $\neg\phi_n$ is in Φ_n . \square

Lemma 3. $\lim_{n \rightarrow \infty} M_n^*(\perp) = 0$.

Proof. From Lemma 1, we have $\lim_{n \rightarrow \infty} M_n^*(\neg \perp) = 1$. From Lemma 2, we have $\lim_{n \rightarrow \infty} M_n^*(\perp) + M_n^*(\neg \perp) = 1$. Together, this gives $\lim_{n \rightarrow \infty} M_n^*(\perp) = 0$. \square

Lemma 4. $\lim_{n \rightarrow \infty} M_n^*(\phi_n)$ converges whenever $\{\phi_n\}$ is quickly computable and $\phi_n \rightarrow \phi_{n+1}$ for all n .

Proof. Because each ϕ_n implies ϕ_{n+1} , in every complete consistent extension of T , either there is some greatest index k such that ϕ_k is false, or ϕ_n is always false. For $k \in \{0, 1, \dots, \infty\}$, let $\{\phi_n^k\}$ be the sequence given by $\phi_n^k = \phi_n$ if $n \geq k$ and $\phi_n^k = \neg \phi_n$ otherwise. Note that for each k , $\{\phi_n^k\}$ is quickly computable, and that with probability 1, there exists exactly one k such that $\phi_i^k \in T^*$ for all i . Note that

$$M_n^*(\phi_n) = \sum_k \mathbb{P}(\phi_n \in \Phi_n \mid \forall i \phi_i^k \in T^*) \mathbb{P}(\forall i \phi_i^k \in T^*), \quad (9)$$

so it suffices to show that $\mathbb{P}(\phi_n \in \Phi_n \mid \forall i \phi_i^k \in T^*)$ converges for each k , because a weighted average of bounded sequences each of which converge also converges.

Fix a $k \in \{0, 1, \dots, \infty\}$. Let $p = \mathbb{P}(\forall i \phi_i^k \in T^*)$ and let $\varepsilon > 0$. Using the same approach as in the proof of Lemma 1, we can choose N_1 such that, for $n > N_1$, with probability at least $1 - p\varepsilon$, we have $(\forall i \phi_i^k \in T^*) \rightarrow \phi_n^k \in \Phi_n$. The probability of the conjunction $P(\phi_n^k \in \Phi_n \wedge \forall i \phi_i^k \in T^*)$ is then at least $p - p\varepsilon$. Therefore, $\mathbb{P}(\phi_n^k \in \Phi_n \mid \forall i \phi_i^k \in T^*) \geq 1 - \varepsilon$ for all $n \geq N_1$, so

$$\lim_{n \rightarrow \infty} \mathbb{P}(\phi_n^k \in \Phi_n \mid \forall i \phi_i^k \in T^*) \quad (10)$$

converges. For $k < \infty$, $\lim_{n \rightarrow \infty} \mathbb{P}(\phi_n \in \Phi_n \mid \forall i \phi_i^k \in T^*)$ converges, because $\{\phi_n^k\}$ is eventually just $\{\phi_n\}$. For $k = \infty$, it also converges, because $\phi_n^k = \neg \phi_n$, and as seen in the proof of Lemma 2, with probability converging to 1, exactly one of ϕ_n and $\neg \phi_n$ is in Φ_n . \square

Lemma 5. $\lim_{n \rightarrow \infty} M_n^*(\phi_n) + M_n^*(\psi_n) + M_n^*(\chi_n) = 1$ whenever $\{\phi_n\}$, $\{\psi_n\}$, and $\{\chi_n\}$ are quickly computable, and for all n , $\{\phi_n, \psi_n, \chi_n\}$ is a partition of truth.

Proof. First, observe that

$$\{(\phi_n \wedge \neg \psi_n \wedge \neg \chi_n) \vee (\neg \phi_n \wedge \psi_n \wedge \neg \chi_n) \vee (\neg \phi_n \wedge \neg \psi_n \wedge \chi_n)\} \quad (11)$$

is a quickly computable sequence of theorems. By Lemma 1, each of these sentences must be in Φ_n once n is large.

As seen in the proof of Lemma 2, when n is large we also have that with probability 1, exactly one of ϕ_n and $\neg \phi_n$ is in Φ_n , exactly one of ψ_n and $\neg \psi_n$ is in Φ_n , and exactly one of χ_n and $\neg \chi_n$ is in Φ_n .

Since Φ_n contains no set of sentences from which one can prove a contradiction in fewer than 2^n steps, this means that Φ_n must eventually contain exactly one of ϕ_n , ψ_n , and χ_n . Therefore,

$$\lim_{n \rightarrow \infty} \mathbb{P}(\phi_n \in \Phi_n) + \mathbb{P}(\psi_n \in \Phi_n) + \mathbb{P}(\chi_n \in \Phi_n) = 1. \quad (12)$$

\square

Lemma 6. $\lim_{n \rightarrow \infty} M_n^*(\phi) = P^*(\phi)$.

Proof. It suffices to show that for all ϕ ,

$$\lim_{n \rightarrow \infty} \mathbb{P}(\phi \in \Phi_n) = \mathbb{P}(\phi \in T^*). \quad (13)$$

Note that in the proof of Lemma 4, if we take $\{\phi_n\}$ to be the sequence which is constantly ϕ and considering $k = 0$ and $k = \infty$, we showed that for sufficiently large n , $\mathbb{P}(\phi \in \Phi_n \mid \phi \in T^*) \geq 1 - \varepsilon$ and $\mathbb{P}(\neg \phi \in \Phi_n \mid \neg \phi \in T^*) \geq 1 - \varepsilon$. Since ϕ and $\neg \phi$ cannot both be in Φ_n for sufficiently large n , this means that $|\mathbb{P}(\phi \in \Phi_n) - \mathbb{P}(\phi \in T^*)| < \varepsilon$, so $\lim_{n \rightarrow \infty} \mathbb{P}(\phi \in \Phi_n) = \mathbb{P}(\phi \in T^*)$. \square

References

- [1] Philipp Hennig, Michael A. Osborne, and Mark Girolami. “Probabilistic numerics and uncertainty in computations”. In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 471.2179 (2015).
- [2] Nicholas Hay and Stuart J Russell. *Metareasoning for Monte Carlo tree search*. Tech. rep. Electrical Engineering and Computer Sciences, University of California at Berkeley, 2011.
- [3] Tom Rainforth et al. “Bayesian Optimization for Probabilistic Programs”. In: *Workshop on “Black Box Learning and Inference” at NIPS 2015*. 2015. URL: <http://www.robots.ox.ac.uk/~twgr/assets/pdf/rainforth2015B0PP.pdf> (visited on 01/15/2016).
- [4] Sam Owre, John M Rushby, and Natarajan Shankar. “PVS: A prototype verification system”. In: *Automated Deduction—CADE-11*. Springer, 1992, pp. 748–752.
- [5] Gerwin Klein et al. “seL4: Formal verification of an OS kernel”. In: *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*. ACM, 2009, pp. 207–220.
- [6] Stuart Russell. “Unifying logic and probability”. In: *Communications of the ACM* 58.7 (2015), pp. 88–97.
- [7] Haim Gaifman. “Concerning Measures in First Order Calculi”. In: *Israel Journal of Mathematics* 2.1 (1964), pp. 1–18. DOI: 10.1007/BF02759729.
- [8] Abram Demski. “Logical Prior Probability”. In: *Artificial General Intelligence. 5th International Conference, AGI 2012, Oxford, UK, December 8–11, 2012. Proceedings* 7716 (2012). Ed. by Joscha Bach, Ben Goertzel, and Matthew Iklé, pp. 50–59. DOI: 10.1007/978-3-642-35506-6_6.
- [9] Marcus Hutter et al. “Unifying Probability and Logic for Learning”. In: *IJCAI-13 Workshop on Weighted Logics for Artificial Intelligence (WL4AI-2013)*. 2013, pp. 65–72.
- [10] Mikael Cozic. “Impossible States at Work: Logical Omniscience and Rational Choice”. In: *Contributions to Economic Analysis* 280 (2006), pp. 47–68.

- [11] Joseph Y Halpern and Riccardo Pucella. “Dealing with logical omniscience: Expressiveness and pragmatics”. In: *Artificial Intelligence* 175.1 (2011), pp. 220–235.
- [12] Jerzy Łoś. “On the axiomatic treatment of probability”. In: *Colloquium Mathematicae*. Vol. 2. 3. 1955, pp. 125–137.
- [13] Theodore Hailperin. “Probability logic.” In: *Notre Dame Journal of Formal Logic* 25.3 (1984), pp. 198–212.
- [14] Dana Scott and Peter Krauss. “Assigning probabilities to logical formulas”. In: *Aspects of inductive logic* 43 (2000), pp. 219–264.
- [15] Dan Suciu et al. “Probabilistic databases”. In: *Synthesis Lectures on Data Management* 3.2 (2011), pp. 1–180.
- [16] David Picado Muiño. “Measuring and repairing inconsistency in probabilistic knowledge bases”. In: *International Journal of Approximate Reasoning* 52.6 (2011), pp. 828–840.
- [17] Nico Potyka and Matthias Thimm. “Probabilistic Reasoning with Inconsistent Beliefs Using Inconsistency Measures”. In: *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI-15), Buenos Aires, Argentina*. 2015, pp. 3156–3163.
- [18] Stephen Muggleton and Luc De Raedt. “Inductive logic programming: Theory and methods”. In: *The Journal of Logic Programming* 19 (1994), pp. 629–679.
- [19] Shan-Hwei Nienhuys-Cheng and Ronald De Wolf. *Foundations of Inductive Logic Programming*. Vol. 1228. LNAI. Springer, 1997.
- [20] Luc De Raedt et al. *Probabilistic Inductive Logic Programming. Theory and Application*. Springer, 2008.
- [21] Matthew Richardson and Pedro Domingos. “Markov logic networks”. In: *Machine Learning* 62.1-2 (2006), pp. 107–136.
- [22] Raymond Ng and Venkatramanan Siva Subrahmanian. “Probabilistic logic programming”. In: *Information and Computation* 101.2 (1992), pp. 150–201.
- [23] Stephen Muggleton. “Stochastic logic programs”. In: *Advances in Inductive Logic Programming* 32 (1996), pp. 254–264.
- [24] Taisuke Sato and Yoshitaka Kameya. “PRISM: a language for symbolic-statistical modeling”. In: *Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI-97), Nagoya, Japan*. 1997, pp. 1330–1339.
- [25] David Poole. “The independent choice logic for modelling multiple agents under uncertainty”. In: *Artificial Intelligence* 94.1 (1997), pp. 7–56.
- [26] Liem Ngo and Peter Haddawy. “Answering queries from context-sensitive probabilistic knowledge bases”. In: *Theoretical Computer Science* 171.1 (1997), pp. 147–177.
- [27] Daphne Koller and Avi Pfeffer. “Learning probabilities for noisy first-order rules”. In: *Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI-97), Nagoya, Japan*. 1997, pp. 1316–1323.
- [28] Thomas Lukasiewicz. “Probabilistic Logic Programming”. In: *13th European Conference on Artificial Intelligence (ECAI-98)*. Wiley, 1998, pp. 388–392.
- [29] Kristian Kersting, Luc De Raedt, and Stefan Kramer. “Interpreting Bayesian logic programs”. In: *AAAI-2000 Workshop on Learning Statistical Models from Relational Data*. 2000, pp. 29–35.
- [30] Scott Garrabrant et al. “Asymptotic Logical Uncertainty and The Benford Test”. In: *arXiv:1510.03370 [cs.LG]* (2015).