Reviewer 1 Sorry for the confusion. The argument is roughly that because Theorem 4 establishes that $\mathrm{GND}(\alpha, 1)$ is satisfiable if and only if α is in the countably infinite domain, and Lemma 5 establishes that the introduction of extra names in $\mathrm{GND}(\alpha, z)$ preserves satisfiability, we obtain satisfiability under the larger, common subset of names used in GND^{-} as well. Please see the Appendix of Ref. 1 for a formal proof.

For your second question, note that for our purposes, it is enough to only consider masking functions that mask all but a finite subset of the domain, thus producing finite-size partial models. Thus, we can take both our masking function and the "application function" $\mathrm{app}(M, \Theta)$ to have a countable range, not something with continuum (or larger) cardinality. In particular, then, we get measurability: these discrete output sets can be defined by a countable union over the resulting finite partial examples. The preimage of a single partial model in turn will be a measurable set for M and Θ, given that M is a measurable function: by definition, the preimage for a measurable M of that partial model is a measurable set.

Reviewer 2 The representation lanaguage proper^{+} that we use is emerging as a popular representation language. FOL with universal quantifiers is widely used to express inductive properties in mathematics but also to represent social networks and graphs. For computability results, usually the finite domain assumption is made, but interestingly, proper^{+} seems to allow us to go beyond the closed-world assumption. (And unlike description logics, arity restrictions are also not needed.) We note that beyond the fact that proper^{+} extends incomplete databases (L104–107), for example (Liu and Lakemeyer, 2009) show how to represent a certain family of "local" action models for planning within the fragment of proper^{+} for which Theorems 14/16 give polynomial-time reasoning. There is also a variant for epistemic planning (Muise et al. 2015) where one reasons about the mental states of other agents, and we expect analogous extensions of our work may contribute to that direction too. In particular, our approach applies to infinite domains, or even simply large domains without resorting to directly considering all groundings of the atomic formulas, in contrast to Juba's work. Note that even in moderate size, finite domains, the number of groundings grows exponentially with the arity of the formulas under consideration, and thus quickly grows infeasible to represent as a propositional formula (which is required for Juba's approach).

Implicit learning works because the partial models themselves compactly encode all of the rules that could be learned from those models. So instead of trying to learn a large set of rules from the models and hoping that these rules will permit us to derive the desired conclusions, we use the models directly to answer queries.

Reviewer 3 Sorry for the terse exposition. The *universal closure* is the result of placing a universal quantifier on each free variable appearing in the formula. $\supset$ (L81) denotes implication. Maximality (L107) refers to the database case, in which every true literal is included; by contrast we have a set of clauses that may not specify all of the true literals. $e\theta$ (L111) indeed refers to applying the substitution θ to e. z (L116) is the rank, which yes is an integer. We use $\mathbb{N}$ as the set of names for convenience, but it is not important here that z could be interpreted as a name. In Proposition 8 (L175) we simply take a union bound over the error events which have respective probabilities ϵ_i. Note that $1 - \epsilon'$ validity only requires that the total probability of the error events is *at most* ϵ'. The union bound applies to any set of events and in particular does not require independence. The actual guarantees of the informal discussion on L220 are formalized in Theorem 13. There is not a requirement on a distribution of queries. Rather, what we promise is that I: we will not (significantly) overestimate the validity of a query and II: we guarantee that our estimate of the validity is (approximately) at least the probability that some suitable implicit KB $\mathcal{I}$ is witnessed.

Learning from entailment is pretty different from what we seek here: it asks us to produce a set of formulas H that replicates a desired set of entailment judgments, e.g., that ϕ_1 is entailed but ϕ_2 is not, etc. Our task formulation is much closer to learning from interpretations in ILP, where our partial models are partial intepretations. In that task, one is given a set of background knowledge formulas B and a set of models $x_1, x_2, \ldots$ and seeks an additional set of background knowledge H such that $B \wedge H$ is consistent with the given models. One could subsequently answer entailment queries against $B \wedge H$. The difference is first that ILP only seeks an H that is consistent with the examples, and does not seek to analyze the degree to which the resulting formulas capture an unknown, ground-truth process that produced the example models $x_1, x_2, \ldots$. In particular, there is no sense in which the resulting judgments $B \wedge H \models \alpha$ are "correct" or "incorrect" in ILP, unless we use a "closed-world" assumption or something similar, that leads the set of models to *define* a single "correct" H. Even so, in practice, ILP often requires significant restrictions on the set of clauses permitted in H to ensure that there is a finite Herbrand base of atoms to search through. We will include a discussion on this in the paper.

The time complexity bound is good in the sense that for fixed approximation and confidence parameters γ and δ, the time complexity of querying the implicit KB is equivalent to a constant number of queries for an explicit KB (cf. Theorem 14). So it remains tractable.

References (a) Y. Liu and G. Lakemeyer. On first-order definability and computability of progression for local-effect actions and beyond. In Proc. IJCAI, pages 860–866, 2009. (b) C. J. Muise, V. Belle, P. Felli, S. A. McIlraith, T. Miller, A. R. Pearce, and L. Sonenberg. Planning over multi-agent epistemic states: A classical planning approach. In AAAI, 2015.