Review of "Regular Path Queries and Modal Constraints" by Albert Gutowski

Ian Pratt-Hartmann

December 5, 2025

The primary objective of this thesis is to prove a single theorem, namely, that the problem of determining entailments between unions of regular path queries modulo a theory expressed in the description logic $ALCO_{\mathrm{reg}}$ is decidable. A regular path query is a formula of the form $\exists \bar{x}\, \chi$, where $\bar{x}$ is a tuple of variables and χ a conjunction of atoms whose arguments are variables or individual constants, and whose relations are either unary predicates or regular expressions over an alphabet of binary predicates, interpreted in the expected way. One query entails another when all tuples of objects satisfying the first query in some *finite* structure also satisfy the second. The qualification that entailment is modulo some $ALCO_{\mathrm{reg}}$-theory means that, when determining such entailments, attention is confined to finite structures that are models of the theory in question. The logic $ALCO_{\mathrm{reg}}$ is a restricted, yet non-first-order logic obtained, in essence, by taking the standard translation of multimodal propositional logic and extending it with individual constants and regular expressions over the alphabet of modalities.

The strategy adopted here—quite standard in this area—is to view query satisfaction in terms of graph homomorphisms. A query can be seen as a (labelled, directed) graph whose vertices correspond to the featured variables and whose labelled edges reflect the atomic formulas involved; a structure can likewise be seen as a (labelled, directed) graph in whose set of vertices corresponds to the domain of quantification and whose labelled edges reflect the satisfaction of atomic formulas by individuals. (Where regular expressions are allowed, the precise correspondence becomes complicated.) A Boolean query—one with no free variables—is then true in an interpretation just in case its graph can be homomorphically embedded in the graph corresponding to the interpretation. Thus, the query entailment problem is translated to a homomorphism entailment problem: we want to know, for (appropriately labelled, directed) graphs G_1 and G_2, whether the existence of a homomorphism from G_1 to any graph H in a certain specified class $\mathcal{C}$ entails the existence of a homomorphism from G_2 to H. The specified class $\mathcal{C}$ is defined in a particular way—essentially by sets of $ALCO_{\mathrm{reg}}$-formulas having very small quantifier depth. Note that this $\mathcal{C}$ may be infinite. The key tool employed is the notion of *minimal downsets.* Take any graph H in $\mathcal{C}$, and consider those graphs of a certain bounded size that can be homomorphically embedded in H. The minimal downsets are the smallest such sets of graphs (under set-inclusion) obtained as H ranges over $\mathcal{C}$. The idea is that, if you can compute the minimal downsets (for a given $\mathcal{C}$), you can solve the homomorphism entailment problem. This minimal downsets problem is in

IEPH 8.12.25

turn reduced to a satisfiability problem—roughly, the problem of determining whether the given class C (specified as above) is empty.

To ease the burden on the reader, the version of the problem without regular expressions is considered first. This is the work of Chapters 5 and 6, when things really start to get moving. The reasoning is very detailed, but the key move in the argument appears to be that of Proposition 6.5 (p. 68). This is proved over the course of most of Chapter 6, but the essential idea is simple (and, I believe, familiar): any graph in the class C we are dealing with can be replaced by an equivalent graph with no short cycles. The principal claims then follow fairly routinely from that point. Chapter 7 then extends this to graphs which are annotated with reachability information. (One vertex is reachable from another in a directed graph if there is a path of edges, in the obvious sense, from the first to the second.) This is a much more difficult case, and a natural prelude to the treatment of regular expressions later in the thesis. The material here is actually rather entertaining: the example of Fig 7.4 on p. 83 shows nicely how this case differs essentially from the simpler case of the previous chapter, and thus helps to motivate the ensuing discussion. The key move in the argumentation here is to consider the strongly connected components of the graphs in question. (These are just the maximal sets of mutually inter-reachable vertices.) This makes obvious sense: within a strongly connected component, reachability information goes away; outside them, it is somehow harmless. I think I learnt most from this chapter. Chapters 9 and 10 complete the development to show the decidability of the homomorphism entailment problem for the full class of labelled graph considered. Chapter 11 then relates this back to the problem in logic that originally motivated it, thus proving Theorem 1.1.

The thesis is well-written, in good English, and in an appropriate style. I found no obvious typographical errors or mathematical mistakes—in a work of this length, an incredible feat. There are a few ill-chosen forms of expression here and there. For example, the term "the crux" in the headings to 6.2, 7.5 and 8.3 is unfortunate: there are no such things as cruxes. (True, we say "the crux of the argument", but then we also say "for pity's sake", and there are certainly no such things as sakes.) I also noticed some mid-Atlantic orthography, e.g. "labelled" (British) vs. "gray" (American). But all these matters are perfectly trivial. One very useful aspect of the presentation is the frequent appearance of examples. The Candidate certainly knows his TIKZ.

More seriously, the very abstract presentation, I felt, tended to obscure rather than clarify the essence of the argument. The definitions are so layered one on top of the other, that it is very hard to see where the real work is being done. As an example, consider the definition of the homomorphism entailment problem at the bottom of p. 32. To understand this, we need to know what a ***shallow modal profile*** is. Well, ***shallow*** was defined a few lines up on the same page, so represents no problem. ***Modal profile*** was defined only a few pages pack, on p. 29, l. 20, and refers to an ***n-bisimilarity profile***, which is defined on the same page, l. 14, as a set of ***n-bisimilarity types***. But what's an ***n-bisimilarity type***? We get a clue on l. 1 of the same page, where the n-bisimilarity type ***of a vertex u in a graph G*** is defined—though to understand this deinition, we have to know what a ***bisimilarity type*** (unqualified) and an n-step unravelling are. The ***bisimilarity type*** of a tree (including n-step unravellings) is in turn defined on p. 27 l. 15. (The definition isn't easy.) However, that doesn't exactly resolve the question of what an n-bisimilarity type (with-

IEPH 8.12.25

out reference to u and G) actually is. Presumably, an an n-bisimilarity type is the an n-bisimilarity type of u in G for some vertex u in some graph G. Now this may be what the Candidate is trying to say on p. 27 l. 16, when he writes "n-bisimilarity types are bisimilarity types of all trees of depth at most n". The problem is: this isn't a great English sentence, and, if it means anything (which is debatable), probably asserts that, for all x, x is an n-bisimilarity type just in case it is identical to some y such that for all z, if z is a tree of depth at most n, then y is a bisimilarity type of z." Probably, the Candidate did not mean this. The problem is not that the reader cannot work out what is going on here. I believe I could. Rather, it is that few people will have the patience to unpeel all these definitions, and those that do will not be delighted to find a muddled sentence at the end of the line. More generally—and more constructively—I do wonder if it would not be better to present the ideas in a way which makes it easier to see where the action is really taking place. The numerous examples given in the thesis help, but only to an extent. I cannot in truth say that this thesis was as easy read.

That said, I do believe the principal result of this thesis, and think that the Candidate has made an important contribution in the area of Computational Logic. Few theses display this level of technical sophistication and precision. There is therefore no question in my mind but that the Candidate should be awarded the degree of Ph.D.

Ian Pratt-Hartmann

8.12.25