VACUUM-LOGIC

Jennifer J. Elgot-Drapkin*
Department of Computer Science and Engineering
Box 875406
Arizona State University
Tempe, AZ 85287-5406 USA
drapkin@asu.edu

Sarit Kraus
Department of Mathematics and Computer Science
Bar-Ilan University
Ramat Gan, Israel
sarit@bimacs.cs.biu.ac.il

Michael Miller, Madhura Nirkhe, Donald Perlis
Department of Computer Science
University of Maryland
College Park, MD 20742 USA
{mm,madhura,perlis}@cs.umd.edu

What reasoning mechanism will a room-vacuuming robot need? One might argue that no logic at all is needed, and this would be hard to refute. The matter is too open-ended for hard and fast proofs. But until someone designs a highly successful vacuuming robot that eschews all logic, it is not unreasonable to consider what sort of logic such a robot *might* use. Indeed, even if an alogical robot could be perfected, this is no argument that another design might not make good use of logic. So, this is the issue we are taking up: given that a logic is to provide at least a portion of the "reasoning" in a vacuuming robot, what sort of logic should it be? We argue that a species of logic ("active" or "step" logics) that we have been using for other purposes (commonsense reasoning, planning, language change) may be highly applicable in the vacuuming domain as well (even if only because these other purposes also surface in vacuuming). Below, we identify five such desiderata for a vacuum-logic, which appear to fit nicely into the mold of active logics.

1. Vaccy usually vacuums the living room. Today he finds a note near one end of the room, saying "chemically treated carpet---do not vacuum beyond this point." So, he needs an "embedded" logic, that can relate indexical terms to environmental positions ("beyond this point"). Such a logic will have to provide a means for Vaccy to assign meanings to expressions ("this point") that can change their meanings from one context to another, and can be sensitive to observations.

A related theme here is the *unexpected*: things can go badly wrong, and in so many ways that it is unrealistic to expect Vaccy's designers to have anticipated all problems. For instance, the curtains might get caught in the vacuum; or the

dog's tail might instead; the electricity might go out; the vacuum might start to smoke; etc. It is all too easy to give Vaccy a general-purpose instruction to simply turn off the vacuum and await further instructions if anything does not go according to plan. The trouble is that things almost never go exactly according to plan, and Vaccy would become useless if he simply stops altogether at every *minor* change. So, he needs to distinguish among changes; he needs to observe what transpires and decide, based on *commonsense*, what to do. Now, commonsense is not something that sits on the AI shelf waiting to be plugged into Vaccy. On the other hand, it is something that has been the object of intense study, much of it logic-based.

Our hope is that large doses of nonmonotonic reasoning will be useful here. For instance, there can be axioms instructing Vaccy that certain events are seriously abnormal (so that he should shut down and/or seek help). These need not be sharply specific such as referring to the dog or curtains, but can refer in general terms to "damage", and there can be a list of things that it is okay to damage (small pieces of paper on the floor with no writing on them; mud; lint; etc). In addition, it probably is not enough for Vaccy to switch off the vacuum when illicit damage does occur. In the case of curtain damage it may be okay, but not in the case of smoke, or an injured dog-tail. Vaccy need not know veterinary first-aid, nor fire prevention, but he should recognize an emergency and know how to call for help. Here too logic may be useful: phoning for help is a good idea (unless the phone is dead). Going to the neighbor is also a good idea, but which is better---the phone or the neighbor?---and how long should he deliberate between them? That depends, again, on too many things to put into explicit one-by-one instructions for Vaccy in advance. So, we confront in fairly blunt terms the general issue of formal commonsense reasoning. But even more, it must be reasonably fast as well; we return to this below (points 2 and 5).

Moreover, even a rule such as "get help in case of fire" is too general. Vaccy should not seek help upon observing the gas pilot light on the stove-top, nor a candle burning in a well-protected glass lantern. Nor is the *size* of the fire the proper criterion. A small fire can be extremely dangerous if it is likely to spread. Thus a substantial amount of naive physics seems necessary, if Vaccy is to have enough commonsense about his world to avoid wasteful calls to the Fire Department. Similar comments apply to "get help in case of damage to living things"---for a dog and an ant do not require equal treatment. Whether enough commonsense can be axiomatized---without reverting to an exhaustive separate handling of each special case---is a problematic issue.

2. Another note tells Vaccy: "Vacuum cleaner is in closet; be sure to be done vacuuming before noon." So, he needs a deadline-sensitive logic, that can track the passage of time. An ordinary temporal logic will not do: as Vaccy vacuums through the morning, reasoning about this and that ("move the chair, avoid bumping the vase", etc), he needs to keep his reasoning abreast of the changing time as he goes. He cannot perform a new lightning-fast closed episode of temporal reasoning every minute, starting each one fresh with an updated clock. This is because many subtasks (change the vacuum bag, move the table) may take more than a minute and may require re-thinking as he goes. This is duck-soup for active logics, since they were originally designed for this very thing: keeping track of time-passage [?1;0cduring reasoning, especially for use in deadline-coupled planning.

3. Yet another note, perhaps left over from last week, says: "The vacuum cleaner is in the bedroom," even though in fact it is in the closet as per note 2. So, he needs a contradiction-tolerant logic. Perhaps he should just ignore the mistaken note in this case---but even this calls for a nonstandard logic: how does he

represent the message before he decides it is mistaken? Is he to take *all* notes as suspect until observationally verified?---that would be a risky approach since some notes may be unverifiable, such as the one about the chemically treated carpet. Nor of course can he take *all* notes as true until falsified by observation, since notes 2 and 3 here are contradictory.

4. A fourth note says "When you are done vacuuming, please empty the bag into the bag next to the step-ladder in the kitchen." Vaccy must be able to disambiguate the two uses of "bag", what we call "semantic shift". A related ability Vaccy must have is to be able to add to his lexicon: "step-ladder" might not be an expression he knows, so he must realize this and---if possible---give it a meaning based on context (including what he sees in the kitchen).

5. Moreover, Vaccy is involved in a real-world task and one that interacts with the hectic pace of human society---the room to be vacuumed is in someone's house---the task must be done in reasonably short time and must take up reasonably limited space (at least in the sense of machinery and memory usage). So, the vacuum-logic must avoid the usual logical omniscience that provides infinitely many conclusions. Inference must be fast and narrow. So some sort of guides must contain the logic, perhaps a kind of short-term memory.

Active logics appear to have features well-suited to all these desiderata. Already they have been applied to default reasoning [1, 3, 2], indexicals [4, 5], observations [1, 3], plans and deadlines [6, 7], contradiction detection and recovery [1, 4], and semantic shift [4, 5]. While we have not applied active logics to position-indexicals, we have done so for other indexicals and for meaning change in other settings, and we suspect the same techniques can be applied here. We are currently working on a short-term memory addition to active logics [7] which will ensure the correspondence between real time and internal clock time is maintained as reasoning

progresses.

One might argue that no theoretical work is justified for a real-world agent, yet active logics, unlike most other theoretical formalisms, are specifically designed to go "into" a real-world agent, since they are inherently resource-limited and tractable. In our Workshop presentation, we will attempt to couch active logics in the vacuuming scenario, with detailed axioms and, if time permits, sample executions of implementations. We note that much of the active-logic work is already implemented in other domains.

In the remainder of this position paper, we sketch an English gloss of an episode of active logic inferences as might be employed by Vaccy. We have arbitrarily taken 30 seconds as the step-grain, purely for purposes of illustration here; the step-grain has no built-in value in our formalisms and is quite flexible.

step # (time)	some beliefs at given step #
11:39:30	Note says to finish vacuuming by noon.
11:40:00	It's 11:40, 20 min. left to vacuum.
11:40:30	I can finish the vacuuming in 18 min.
11:41:00	There is enough time to finish before noon.
11:41:30	I'll get the vacuum out of the bedroom.
11:42:00	Vacuum is not in bedroom, note is wrong.
11:42:30	It's now too late to finish vacuuming by noon.

In this scenario, Vaccy discovers the note to vacuum by noon; observes the current time; assesses how long it will take to vacuum; assesses that he does indeed have time to finish by the deadline; proceeds to find the vacuum; upon discovering that it is not in the bedroom, recognizes that the note was wrong; and then realizes he no longer has time to finish the vacuuming job by noon---he spent too much time looking for the vacuum. We see that

Vaccy must not only be able to reason about time, but he must be able to do so *as time is going on.* That is, it is not only necessary to be able to reason *about* time, but Vaccy must also---and this is critical to any real-world agent---reason *in* time. It is this fundamental notion upon which the active logic research is based.

space and computation bounds. In *Common Sense 93*, Austin, Texas, 1993. Also to be presented at the IJCAI-93 Workshop on Spatial and Temporal Reasoning.

References

[1] J. Elgot-Drapkin. *Step-logic: Reasoning Situated in Time.* PhD thesis, Department of Computer Science, University of Maryland, College Park, Maryland, 1988. Available as Technical Report CS-TR-2156, Technical Report UMIACS-TR-88-94, and through UMI order no. 8912283.

[2] J. Elgot-Drapkin. Step-logic and the three-wisemen problem. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 412--417, 1991.

[3] J. Elgot-Drapkin and D. Perlis. Reasoning situated in time I: Basic concepts. *Journal of Experimental and Theoretical Artificial Intelligence*, 2(1):75--98, 1990.

[4] M. Miller. *A view of one's past and other aspects of reasoned change in belief.* PhD thesis, Department of Computer Science, University of Maryland, College Park, Maryland, 1993. (Directed by D. Perlis.).

[5] M. Miller and D. Perlis. Presentations and this and that: logic in action. In *15th Annual Conference of the Cognitive Science Society*, Boulder, Colorado, 1993.

[6] M. Nirkhe, S. Kraus, and D. Perlis. Fully deadline-coupled planning: One step at a time. In *Proceedings of the sixth international symposium on methodologies for intelligent systems*, Charlotte, NC, 1991.

[7] M. Nirkhe, S. Kraus, and D. Perlis. Situated reasoning within tight deadlines and realistic