Dear Paul Rosenbloom,

We regret to inform you that your paper, "Towards a New Cognitive Hourglass: Uniform Implementation of Cognitive Architecture via Factor Graph" (#607), has not been accepted for presentation at the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09). The competition was especially strong this year, with only 331 papers accepted out of 1,290 submissions. Attached to this message are a description of the review process and comments from the reviewers of your paper, which should explain their recommendation and the decision made by the Program Committee.

We understand that every paper results from careful thought and hard work by its authors. We hope that the reviewers' comment help you to improve your paper and that you find an appropriate venue for publication of your work (including future IJCAIs).

Thank you for your interest in the IJCAI-09 conference.

Sincerely,

Craig Boutilier
IJCAI-09 Program Chair

DESCRIPTION OF THE REVIEW PROCESS

Every paper received at least three reviews. There were approximately 600 program committee members, 91 senior program committee members, and 22 area chairs who oversaw the final decision-making process. The reviews, online discussion, author feedback, and often the papers themselves were read by at least one of the senior program committee members and area chairs. Detailed discussion was initiated before and after the author feedback phase, and, when necessary, additional reviews were collected. For most papers, and in all cases where the decision was controversial or unclear, the senior program committee member and/or the area chair wrote a meta-review, summarizing the positive and negative aspects of the paper. Final decisions were then made by the area chairs and the program chair based on all information: reviews, discussion and recommendations. While no process of this scope is flawless, the tremendous efforts put in by so many people ensured that process was as careful, fair and transparent as possible.

==================

------------- Review from Reviewer 1 -------------

Relevance : 8
Significance : 2
Technical Soundness : 2
Novelty : 5
Quality of Evaluation : 1
Clarity : 5
OVERALL SCORE : 3
CONFIDENCE SCORE : 5

-- Comments to the author(s):

Relevance

Cognitive architectures - the subject at the core of this paper - are a central problem of AI. The presented material is therefore highly relevant to this conference.

Significance

The authors propose factor graphs as an interface layer (or as a uniform implementation layer) capable of producing a narrow waistline (the hourglass) in a cognitive architecture. Since the paper does not present any qualitative or quantitative results to compare cognitive architectures based on this idea with those not using factor graphs as a uniform implementation layer, I can only evaluate the significance of the presented idea: the use of factor graphs in cognitive architectures.

In the introduction, the authors refer to Domingo's call for an interface layer in AI. And the authors propose that factor graphs are an answer to that call. I wholeheartedly concur with the importance of an appropriate interface layer (or layers) for AI and for cognitive architectures. However, I remain unconvinced that factor graphs help us to define this layer.

Looking at Domingo's table 1 in the cited paper, it seems striking that every interface layer hides some detail and provides an appropriate abstraction for a higher level. And, of course, to achieve cognition, it seems imperative to identify relevant, lower-dimensional subspaces of the overall "AI problem". This then is the important function of an interface: based on a functional and abstraction-based subdivision, it provides an interface between two levels or two capabilities that ultimately lets each of the interfaced parts solve their respective problem, while simplifying the synergistic collaboration between them. To enable this, the interface needs to help reduce the dimensionality of the overall problem by providing an appropriate factorization. The interface thus must be restrictive, it must impose structure!

In contrast, as pointed out by the authors in the introduction, factor graphs subsume a large number of algorithms and representations used in AI. They do not impose structure, they act as a universal interface. Therefore, at least in the opinion of this reviewer, the use of factor graphs as an interface or implementation layer in cognitive architecture in itself, does not represent an advance toward cognition or towards better cognitive architectures. In fact, it does not even point towards a concrete architecture. Instead, it enables an effective, unifying representation of a wide range of architectures.

In general, it would be very insightful to to support claims about progress towards cognitive architectures with experimental results that demonstrate cognitive behavior. Without such a validation, claims about progress cannot be validated.

Technical soundness

The presentation of factor graphs is sound. The paper does not contain any information that would allow the reviewer to infer conclusively that factor graphs are indeed a technically sound interface layer in cognitive architectures.

Novelty

The idea presented in this paper is novel.

Quality of Evaluation

The paper does not contain any experimental validation of the central claim made by the authors.

Clarity

The presentation and style of writing are adequate.

---------- End of Review from Reviewer 1 ----------

Response to Reviewer 1:

This paper was intended to provide a radically new vision for cognitive architecture, an argument for its plausibility, and a concrete step towards its realization. It is indeed still a long way from a full cognitive architecture based on factor graphs. A vision should be evaluated in the short term by novelty and potential, and in the long term by results. Asking for evaluation of a full architecture at this point is a bit like asking that any work in AI be run against the Turing test. The concrete step –

production match in factor graphs – is the first of its kind. It needs to be evaluated primarily by sufficiency, as was done.

The reviewer focuses on the concept of an interface layer, which is not quite the same as that of the implementation level used in the submission. An implementation level is not about relating higher to lower layers. The focus is on the uniform integration of capabilities required in cognition rather than on a better or more restrictive interface between lower and higher layers.

Factor graphs do subsume a large number of algorithms and representations, but this is a positive as far as the kind of breadth required in a cognitive architecture. Factor graphs are not arbitrarily malleable though in terms of the number of very disparate ways they can support specific capabilities. Thus they provide both breadth in capability and constraint within capability, exactly what is desirable in an implementation level for cognitive architecture.

------------- Review from Reviewer 2 -------------

Relevance : 5
Significance : 3
Technical Soundness : 2
Novelty : 2
Quality of Evaluation : 2
Clarity : 3
OVERALL SCORE : 3
CONFIDENCE SCORE : 7

-- Comments to the author(s):

As far as I could understand this paper, it aims at discussing the need for cognitive architectures that establish a balance between the diversity of cognitive behaviors they support and the uniformity of mechanisms they use to explain and implement those behaviors. The authors claim that factor graphs is the right mechanism for the implementation of cognitive architectures and they briefly report how soar's reactive layer can be implemented by means of factor graphs. The application of factor graphs to implement the reactive layer of the soar architecture seems to be the main message of this paper and its only relevance to AI.

The discussion on cognitive architectures is unclear and the idea of applying factor graphs for implementing soar's reactive layer is not significant. Many unjustified statements are made in the paper and the relation to the literature is vague and imprecise.

The proposed application of factor graphs and the implementation of soar's reactive layer are not evaluated. There are not much technical details, formal proofs, or experimental evidence in the paper to illustrate the plausibility of the proposed approach. These make it hard to assess the technical quality of the paper.

Neither the soar architecture nor the implementation of its reactive layer is new. The application of factor graphs does not support the claim of the paper in a convincing manner; it only shows a possible way to implement the matching and application of production rules of the soar's reactive layer.

The most problematic aspect of this paper is related to its readability. I found the paper very hard to follow and understand. The authors do not formulate the motivation, the technical problem, and the result of the paper in a concrete and explicit manner. They also do not conclude the paper with an evaluation of the proposed approach. There are many references to existing works on cognitive architectures, but the relation between the presented approach and the referred literature remains unclear and vague.

---------- End of Review from Reviewer 2 ----------

Response to Reviewer 2:

This paper was intended to provide a radically new vision for cognitive architecture, an argument for its plausibility, and a concrete step towards its realization. The one concrete claim is that factor graphs, a novel approach to implementing production match, are sufficient for this. This claim is supported in the submission. The overall message for AI is broader though. A vision should be evaluated in the short term by novelty and potential, and in the long term by results it engenders. To start off by engendering a novel match algorithm yielding improved cost bounds is a good sign on the results side.

The reviewer states that the concrete result is neither novel nor significant. If the algorithm is not novel, I would appreciate a pointer to prior work that has done this, as I am unaware of any such. The statement about lack of significance, and much of the rest of the review – concerning readability and vagueness in relationship to the cognitive architecture literature – lead to a concern that the reviewer is unfamiliar with both cognitive architectures in general and production system match algorithms in particular, and perhaps comes from the graphical community instead. This kind of issue is a potential risk whenever an attempt is made to bridge two areas that presently lack familiarity with each other's goals and results. But novel interdisciplinary work must go on, and should be publishable at IJCAI.

------------- Review from Reviewer 3 -------------

Relevance : 6

Significance : 3
Technical Soundness : 4
Novelty : 3
Quality of Evaluation : 2
Clarity : 2
OVERALL SCORE : 5
CONFIDENCE SCORE : 6

-- Comments to the author(s):

Towards a New Cognitive Hourglass

The authors discuss the possibility that a cognitive architecture capable of human-level performance will be uniform or diverse in its structure. While an architecture must be able to complete many tasks (diverse at the top), and by the strong AI hypothesis, can be implemented on any computational substrate (diverse at the bottom), it's not clear whether there is diversity/complexity in the structure throughout or whether some simple and uniform system (or principle) exists as a middle layer.

The authors propose that factor graphs---i.e., algebraic decompositions of functions---are a good candidate for a uniform implementation level, and support this claim by noting the plethora of problems that factor graphs (including bayesian networks and markov random fields) have been successfully applied to. As the authors are familair with Soar, they purpose to re-implement the system on top of factor graphs. In that vein, the authors describe an encoding of the "match" algorithm in the production system Soar as a factor graph, and describe how the application of other work on speeding up inference to this inference task results in a 500 factor speed up over a naive implementation of factor graphs.

I am not convinced by the paper. First, expressing production system algorithms as factor-graph inference seems like the wrong direction. Probability theory seems to be the more appropriate calculus for uncertain reasoning. How does using factor graphs (which can be used to integrate all data and make coordinated inferences) improve the production system?

Second, a fixed, finite factor graph will never be able to handle our complex world where the number and existence of objects will require that, at the very least, the topology of the graph be random as well. Probabilistic languages like BLOG, Church (Goodman et al UAI 2008), IBAL (Pfeffer 2001), or even Markov Logic (Domingos and Richardson) are a much better candidate for a universal glue to hold together a cognitive architecture. At any moment in one of the languages, their inference engines may have a representation in memory isomorphic to a factor-graph. But critically, this structure changes over time in a way which is controlled by the inferences the model makes.

Other comments:

Looking into Blaise, it's characterization in this paper as "generalized factor graphs with limited relational concepts" seems very inaccurate. First, I don't see what Blaise has to do with relational concepts at all. Second, it seems that Blaise's notion of Kernels is fairly important and makes its use of a factor-graph like data structure (States/Densities) a superficial quality. Kernels change the structure of the underlying representation (which is a graph); this is not a factor-graph... it's more like a stochastic automata.

Also, BLOG is not about factor graphs... its about representing knowledge by probabilistic programs. In fact, in that vein there are other examples like IBAL, Church, lambda_circle (Park, Pfenning, Thrun POPL 2008) which go far beyond factor graphs.

---------- End of Review from Reviewer 3 ----------

Response to Reviewer 3 (and their answers):

I'm puzzled by the statement that production system algorithms as FGs seem like the wrong direction. Is this a general paradigm argument about the use of production systems in AI? The submission focuses on a particular architecture that uses them, but the ultimate goal includes probabilistic reasoning (where FGs also work). The main functional advantage of the FG match implementation is the change in the cost bound, but the bigger message is that match is implementable in factor graphs.

I'm warming up to the idea of a "softer" computational substrate.

I'm also puzzled by the notion that no fixed FG can handle our complex world, and that this complexity implies that the topology must be random. The graph will need to evolve over time through learning, and probabilistic reasoning will also be needed, but how does this make BLOG or Markov Logic more appropriate?

"Learning" a graph topology is a point estimate version of treating it as random and doing inference over the random graph. ML is no improvement in this vein, actually. However, BLOG is because you can describe the distribution over the random topology in the language itself and do inference. You don't have to leave the system to express uncertainty in the topology. Factor graphs cannot do this reflection.

There is no claim that FGs are more appropriate, but they are not obviously less so. The main proposal is that factor graphs in particular, and graphical models more generally, which have previously been largely disjoint from cognitive architecture, are a promising approach to implementing them.

Sure.

The submission never states that BLOG is based on factor graphs; it mentions “other variants of graphical models”. With respect to Blaise, Bonawitz’s thesis states that its SDK is a “graphical modeling language that generalizes factor graphs”, and section 6.3 is about relational models. Thus the high-level summary of Blaise does not appear patently incorrect.

It's not incorrect, but its not a good description either. Also, Blaise can characterize uncertainty of the topology of the graph it is doing inference in. I'm skeptical of basing the future of AI on a system that cannot express uncertainty over its structure.