

Representation and Contribution-Integration Challenges in Collaborative Situation Assessment

Daniel D. Corkill

University of Massachusetts Amherst
Amherst, Massachusetts 01003
U.S.A.
corkill@cs.umass.edu

Abstract – *Blackboard systems are an ideal architecture for situation assessment involving large data volumes and heterogeneous data and knowledge sources. However, the ad hoc confidence and belief values used in traditional blackboard applications has led to criticism of the blackboard approach and spawned efforts to combine collaborative blackboard-system techniques with more “principled” graphical-network representations. We discuss two important collaborative-assessment challenge areas: 1) principled blackboard representations and 2) principled integration of contributions made by independent knowledge-source entities. The complexity of these challenges is highlighted using a very simple assessment scenario.*

Keywords: Blackboard systems, multi-agent systems, representation, integration, information fusion, situation analysis.

1 Introduction

Traditional data-fusion approaches have operated within a single fusion level [1], however accuracy and speed at each level can be improved by utilizing the evolving inferences from other levels. The DARPA/IXO Information Fusion Workshop held at Captiva Island, Florida, on December 10–13, 2001¹ was a visible milestone in multi-level data fusion activity. At the workshop, the technical foundations for an advanced high-level data fusion approach were debated by 34 multidisciplinary participants, and an architectural concept was recommended that coupled the powerful and flexible control capabilities of AI blackboard systems [2, 3, 4, 5] with a more “principled” blackboard representation.

Blackboard systems The effectiveness of blackboard systems is the product of a number of architectural capabilities working in concert. The first important capability is the control flexibility provided by indirect, anonymous, and temporally-disjoint interaction among software entities. The blackboard-system control shell can delay execution of a knowledge-source (KS) execution without having to modify an explicit process or worry about managing the data needed by the delayed KS—they remain on the

blackboard. Similarly, KS activations can be executed earlier than normal—whenever there appears to be sufficient information for them to perform useful work.

Blackboard-system applications usually involve multiple representation levels with different abstractions or aggregations. This provides additional control flexibility and the opportunity to work bottom-up (*data-directed*) or top-down (*model-directed*) as is most appropriate given the current state of processing activities and information on the blackboard. Such opportunistic processing is especially important in large-data-volume applications such as situation assessment.

Anonymous and temporally disjoint sharing of contributions go hand in hand with opportunistic control. The collaborative power of blackboard systems [6] results from individual KS executions sharing what is important, unusual, or likely to be useful to another KS—without knowing exactly what specific KSs may benefit from the shared information or even when or if specific shared information will be used. Sharing contributions is crucial for effective collaboration, and blackboard systems provide relatively inexpensive sharing mechanisms. KS executions cannot share everything, as that leads to unwanted distraction [7], but they must share something. The cost of sharing in blackboard systems enables an aggressive sharing strategy where significant amounts of possibly useful information are posted on the blackboard. KS activations are typically triggered by only part of the information they need to operate, so they must find what additional information has been shared previously that meets their needs. Because what information might be relevant is not known when the information is shared, locating anonymously contributed information must be both fast, complete, and highly relevant.

Principled blackboard representations The extensive use of graphical models, such as Bayesian networks, in modern AI applications [8, 9, 10] has led to criticism of the ad hoc confidence and belief values used in traditional blackboard applications. These ad hoc belief values were involved in everything from making control decisions to determining solutions and the system’s confidence in them. This criticism has generated considerable interest in developing what have been termed *Bayesian blackboard* systems. The idea is to replace ad hoc representations of the relation-

This work is supported by the “Fusion Based Knowledge for the Future Force” ATO Program at the U.S. Army RDECOM CERDEC Intelligence and Information Warfare Directorate, Fort Monmouth, NJ, under contract W15P7T-05-C-P621. The views contained in this paper are the author’s.

¹The Captiva Island workshop was originally scheduled the week of September 17th 2001, but was postponed until December due to the events of 9/11.

ships among blackboard objects with incrementally generated graphical models. Recent techniques in constructing belief networks using network fragments [11, 12] and in hierarchical object-oriented Bayesian networks [13, 14] have been suggested as candidate technologies that can be extended to create more principled blackboard reasoning.

Preliminary efforts in applying graphical belief networks to blackboard systems have focused on a principled representation of the developing *solution* on the blackboard [15]. Current beliefs are represented on the blackboard as a possibly disconnected graphical network. A first-order extension to belief networks can be used to collapse similar entities into a single node-type and set of arguments that in combination uniquely specifies a node on the blackboard. Time complicates graphical-network representations, and can involve significantly different temporal scales. To address this, multiple temporal representations have been used: a discrete approach where each node is indexed by the time it occurs² and a duration-interval approach where nodes have a start and an end time that are themselves represented as nodes in the network. Complex spatial representation and reasoning are also problematic for graphical-network approaches, and *procedural KSs* are often used to perform geometric reasoning. Such reasoning is both difficult and highly inefficient to represent explicitly using a Bayesian-network fragment. These concessions to complexity weaken the foundation of principled reasoning sought by a Bayesian approach and real-world situation assessment representation and inferencing requirements present considerable challenges to the use of Bayesian techniques.

Principled integration of contributions The emphasis on developing a principled blackboard representation of the developing solution is misplaced. The emphasis should be on making the *integration* of the contributions made by diverse entities well founded. This can only be achieved by modeling how these contributions are generated and how they relate to one another.³ Blackboard applications work incrementally, with independently developed KSs adding their contributions to those already present on the blackboard. For example, if two KSs use the same data and produce similar results using different computational approaches, how independent are the results? Are they redundant (with no added certainty in the results) or complementary (in the sense that each has the potential to make mistakes on certain data values, but these mistakes are fully independent of one another)? In the latter case, the integration model needs to reflect the additional certainty produced by the corroborating contribution.⁴ To date, “principled” approaches have simply assumed independence of contributions, used fuzzy averaging, and other ad hoc approaches to integration. Note that recording the pedigree of

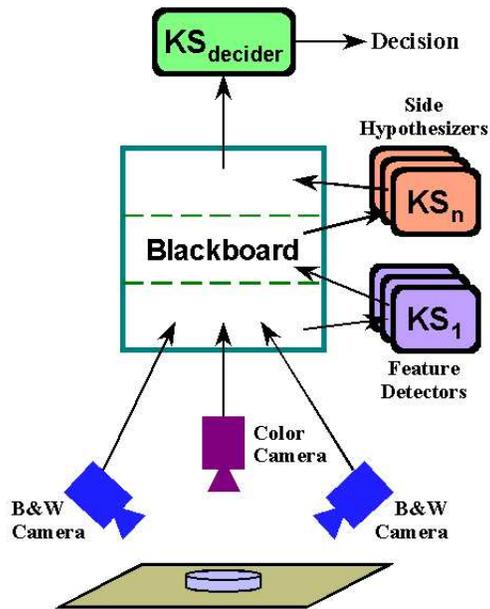


Fig. 1: Fair-Coin Detector

the data used in making contributions is insufficient in understanding how they relate. We also need to understand the context-specific behavior of KS activities and the collective relationships among them.

The Fair-Coin Problem To illustrate these challenges, consider a simple collaborative-assessment problem of deciding if a U.S. quarter is a fair coin (has a head and a tail) by observing a series of coin flips. A priori we are told that there is a 50% chance that the quarter is either two-headed or two-tailed. We have a table top that can be viewed by three cameras: two black-and-white cameras and a color camera (Fig. 1). Images feed into our assessment architecture that includes a number of KSs. There are low-level KSs that attempt to identify coin features, higher-level KSs that aggregate features to hypothesize coin sides, and a decider KS that makes the fair or non-fair-coin designation. The goal is to make a principled determination with a specific confidence with as few flip observations as possible.

Adding to the complexity is the U.S. 50 State Quarter[®] program, where a new quarter with a state-specific reverse side is issued every 10 weeks in the order that the states were admitted into the Union (Fig. 2). Started in 1999, the last state quarter will not be minted until 2008, and designs for the later quarters are not finalized. We do not know what specific quarter might be used (including the non-state eagle reverse design or the 1976 bicentennial commemorative reverse where a colonial drummer replaced the eagle). So, every 10 weeks we must add to or refine our KSs to handle the latest reverse design.

A number of control and integration issues face our blackboard application. Assume we have obtained performance characteristics (error rates and conditions) on each KS. How do we use those characteristics in concert? How should we relate the results of applying feature-identification KS #1 to the images of multiple cameras? Of applying feature-identification KS #2 to the same image? Does our confidence in a feature detected from B&W cam-

²Making the graphical network a Dynamic Bayesian Network [16].

³Principled integration applies to contributions shared via a blackboard or exchanged among agents in a multi-agent system (MAS)—an issue that has received surprisingly little attention from MAS researchers.

⁴A simple example of such a model was used in RESUN [17, 18], where hypotheses included symbolic reasons provided by their contributors that characterized how they were uncertain.



Fig. 2: Example Reverse and Obverse Sides

era #1 increase with detection of that feature from B&W camera #2? From the color camera? How certain are we that it is the same feature? From the same flip? Is the additional processing beneficial, redundant, or somewhere in between? Are the error characteristics correlated (such as the rotational angle of the coin relative to the cameras and lighting)? How are these confidences passed along to higher-level KSs and to the decider KS in a principled way?

Principled integration requires understanding how the collaborating KSs (cameras and software entities) operate in conjunction with one another. As each new quarter is issued and new or modified KSs are added, these interrelationships and the confidence and error characteristics of computations will also change. In open, collaborative environments, such as our fair-coin blackboard application, understanding the interrelationships among collaborating software entities is antithetical with independent and anonymous KS interaction. If we seek to maintain the flexibility to add, remove, and change KSs and to use them opportunistically during processing, we must develop techniques that do not require constructing a rigid, global model of application processing with every KS modification.

One approach to this challenge is to develop and bundle entity-specific behavioral specifications of contribution-generation characteristics with each KS. So, a control component or integrator can use these specifications to characterize the conditions under which an entity is likely to make mistakes, what information and factors it uses in making its contributions, etc. These specifications can be provided by the developers of each software entity, or it may be possible to develop software entities that use learning techniques to expand and refine their control and integration knowledge themselves. Note that the degree that results are shared among entities has a direct relation to the complexity required of the contribution-integration models. The integration model need only address contributions to the extent that they are made available to others.

References

[1] Data Fusion Subpanel of the Joint Directors of Laboratories Technical Panel for C3. Data fusion lexicon. Technical report, F. E. White, Code 4202, NOSC, San Diego, California, 1991.

- [2] Lee D. Erman, Frederick Hayes-Roth, Victor R. Lesser, and D. Raj Reddy. The Hearsay-II speech-understanding system: Integrating knowledge to resolve uncertainty. *Computing Surveys*, 12(2):213–253, June 1980.
- [3] Daniel D. Corkill. Blackboard systems. *AI Expert*, 6(9):40–47, September 1991.
- [4] Robert S. Engelmore and Anthony Morgan, editors. *Blackboard Systems*. Addison-Wesley, 1988.
- [5] V. Jagannathan, Rajendra Dodhiawala, and Lawrence S. Baum, editors. *Blackboard Architectures and Applications*. Academic Press, 1989.
- [6] Daniel D. Corkill. Collaborating software: Blackboard and multi-agent systems & the future. In *Proceedings of the International Lisp Conference*, New York, New York, October 2003. (Invited presentation.).
- [7] Victor R. Lesser and Lee D. Erman. Distributed interpretation: A model and experiment. *IEEE Transactions on Computers*, C-29(12):1144–1163, December 1980. (Also published in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, editors, pages 120–139, Morgan Kaufmann, 1988.).
- [8] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [9] Judea Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2000.
- [10] Spirtes, Glymour, and Scheines. *Causation, Prediction and Search*. MIT Press, 2nd edition, 2001.
- [11] Robert P. Goldman and Eugene Charniak. A language for construction of belief networks. *IEEE Transactions of Pattern Analysis and Machine Intelligence*, 15(3):196–208, 1993.
- [12] Kathryn Blackmond Laskey and Suzanne M. Mahoney. Network fragments: Representing knowledge for constructing probabilistic models. In *Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-97)*, pages 302–313. Morgan Kaufman, 1997.
- [13] Daphne Koller and Avi Pfeffer. Object-oriented Bayesian networks. In *Uncertainty in Artificial Intelligence: Proceedings of the Thirteenth Conference (UAI-97)*, pages 302–313, San Francisco, California, 1997.
- [14] Avi Pfeffer, Daphne Koller, Brian Milch, and Ken T Takusagawa. SPOOK: A system for probabilistic object-oriented knowledge representation. In *Proceedings of the Fourteenth Annual Conference on Uncertainty in AI (UAI-99)*, pages 541–550. Morgan Kaufman, 1999.
- [15] Charles Sutton, Clayton Morrison, Paul R. Cohen, Joshua Moody, and Jafar Abidi. A Bayesian blackboard for information fusion. In *Proceedings of the Seventh International Conference on Information Fusion (Fusion 2004)*, pages 1111–1116, Stockholm, Sweden, June 2004.
- [16] Kevin P. Murphy. *Dynamic Bayesian Networks: Representation, Inference, and Learning*. PhD thesis, University of California at Berkeley, Berkeley, California, July 2002.
- [17] Norman Carver and Victor Lesser. A new framework for sensor interpretation: Planning to resolve sources of uncertainty. In *Proceedings of the National Conference on Artificial Intelligence*, pages 724–731, Anaheim, California, August 1991.
- [18] Norman Carver and Victor R. Lesser. A planner for the control of problem solving systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(6):1519–1536, 1993.