

Cognitive Information Processing Challenges for Homeland Security

Paul Kogut, Kenneth Kisiel

Lockheed Martin Integrated Systems & Solutions, King of Prussia, PA

paul.a.kogut@lmco.com, kenneth.w.kisiel@lmco.com

Abstract

This paper describes the goals and emerging results of an ongoing research project in order to solicit input and feedback from the AI for Homeland Security community. The objective is to gain a greater understanding of the performance issues and computational requirements associated with real world AI applications for intelligence analysis and homeland security. These computational requirements will be used to design next generation hardware/software architectures for cognitive information processing.

Problem

Many AI techniques are notoriously computationally intensive. What are the performance requirements for the application of various AI techniques to real world homeland security and intelligence analysis? Will existing AI techniques scale-up to handle increasing volumes of raw data? What kinds of AI applications become feasible and practical if we can break through the computational limitations of current hardware/software architectures? This paper focuses on the computational requirements of next generation hardware/software architectures for cognitive information processing – not on the design of these architectures.

Approach

There is a renewed interest in Cognitive Information Processing (CIP) and cognitive architectures [Cassimatis, and Winston, 2004]. Our long term research goal is to integrate new software and hardware technologies into an architecture that reconfigures itself to optimally perform various AI tasks. Current general purpose computer systems were not designed for CIP. CIP includes theorem proving and forward/backward chaining systems that exhibit exponential behavior on the size of the input. CIP also includes reasoning with distributed ontologies and knowledge-bases, probabilistic/Bayesian reasoning, machine learning and AI planning. Our team is studying CIP requirements in two mission challenge areas: intelligence analysis and unmanned aerial vehicles (UAVs). This paper will focus on computational

requirements for intelligence analysis relevant to homeland security.

The approach is to identify activities that an analyst does which are currently automated with AI technology or will potentially be automated in the next ten years. Our application study team is currently in the process of identifying a balanced set of AI applications that:

- Support diverse needs of a broad range of Intelligence and Homeland Security organizations
- Cover a wide spectrum of analysis activities
 - Collection, exploitation, analysis, dissemination
- Cover a variety of cognitive services
 - Deduction, induction, probabilistic reasoning
- Cover a variety of sensor inputs and reference knowledge
 - Text, imagery, video, signals, speech, geospatial

We have identified two main CIP application areas to focus on: raw data screening and evidence marshaling. These two areas are described in the following sections. Other intelligence community CIP application areas not discussed in this paper include sensor collection planning, personal assistant agents, question answering and agent-based modeling and simulation.

Raw Data Screening

Analysts are swamped by an ever increasing volume of various types of collected data, such as intelligence reports, annotated imagery, telephone intercepts and open source data, such as web sites and broadcast news [Waltz, 2003]. The analyst needs automated screening tools to help her figure out what to look at first (i.e., tools that filter and prioritize incoming data). A broad range of CIP supports data screening including speech recognition, machine translation, natural language processing and video interpretation. Video interpretation includes real-time monitoring of news broadcasts and surveillance camera networks. These data screening techniques will increasingly rely on automated deductive reasoning in the future. The induction of rules to support deduction will also become a significant CIP requirement.

Evidence Marshaling

The analyst also needs automated support for evidence marshaling (i.e., putting knowledge together to formulate and support multiple hypotheses). During a typical investigation like counter-terrorism, intelligence analysis, or criminal investigation, the primary challenge for the analyst is determining which evidence is important and how the evidence is combined into the most plausible explanation from among all possible explanations. Humans are fraught with cognitive biases [Heuer, 1999] that prevents them from naturally considering alternative explanations. Evidence marshaling will rely heavily on automated deductive reasoning (e.g., temporal, spatial, common-sense) and probabilistic reasoning. We do not claim that the analyst will be replaced by a computer anytime in the foreseeable future. We do claim that automated reasoning will eventually be able to make useful suggestions to the analyst.

Conclusions

Based on the application areas described above we believe that deductive reasoning is the biggest requirements driver of CIP for intelligence analysts. Induction and probabilistic reasoning are also very important. Current practical deductive reasoning techniques include description logics and automated theorem proving [Sutcliffe]. Researchers are only beginning to apply description logic reasoning to intelligence analysis [Kogut et. al., 2005]. We plan to analyze existing deductive reasoning performance data (e.g., [Guo et. al., 2004] and [Sutcliffe & Suttner]) to help formulate CIP requirements.

This paper describes the goals and emerging results of an ongoing research project in order to solicit input and feedback from the AI for Homeland Security community. Please contact the authors if you think important AI techniques have not been considered or you have insights or real world data that might contribute to this effort.

References

- Cassimatis, N., and Winston, P., editors, *AAAI Fall Symposium: Achieving Human-Level Intelligence through Integrated Systems and Research*, Washington, DC, Oct 2004. AAAI Press.
- Guo, Y., Pan, Z., Heflin, J. An Evaluation of Knowledge Base Systems for Large OWL Datasets. *Third International Semantic Web Conference*, Hiroshima, Japan, LNCS 3298, Springer, 2004, pp. 274-288.
- Heuer Jr., R.J., *Psychology of Intelligence Analysis*, Washington D.C., CIA Center for the Study of Intelligence, 1999
- Kogut, P., Leung, Y., Ryan, K., Gohari, L., Kokar, M., Letkowski, J., Applying Semantic Web Technologies to Counter-Terrorism To appear in *21st Century Information Technologies*

and Enabling Policies for Counter-Terrorism Editors R. Popp J. Yen, IEEE Press

Sutcliffe, G. Overview of Automated Theorem Provers
<http://www.cs.miami.edu/~tptp/OverviewOfATP.html>

Sutcliffe, G. and Suttner, C. The TPTP Problem Library for Automated Theorem Proving <http://www.cs.miami.edu/~tptp/>

Waltz, E. 2003 *Knowledge Management in the Intelligence Enterprise*, Artech House