

Provenance narratives and visualisation to support understanding and trust

A. Cunningham ^a, M. Nowina-Krowicki ^b, J. Walsh ^a, J. Chung ^b,
S. Wark ^b and B. H. Thomas ^a

^a*Australian Centre for Immersive Virtual Environments, University of South Australia, Australia,*

^b*Defence Science and Technology Group, South Australia, Australia*

Email: Marcin.Nowina-Krowicki@dst.defence.gov.au

Abstract: We present a system for automatically generating and presenting provenance narratives from World Wide Web Consortium (W3C) compliant provenance records. The interactive narrative visualisation produced allows an analyst or decision maker to explore complex provenance information in a simple, engaging and intuitive manner. Our approach provides a summarised view of the provenance narrative, that can additionally be expanded, collapsed and explored to provide varying levels of detail to the user. Expansion information provides further context to help aid understanding and gain insight. The analyst or decision maker can use the narrative visualisation to explore how: a particular result was achieved; the influence of any particular agent on the result; and the flow-on effects of a particular change in input, constraint or processing that has taken place.

Understanding provenance information in a complex Command and Control system employing AI and automation provides the use case and initial motivation for this work. However, this approach is generalisable, and could be transferred to any W3C Provenance Data Model compliant system. The solution is domain agnostic and capable of handling any type of information, be it tactical, operational or strategic in nature. While still in its early stages, this approach shows promise and may provide a general approach for delivering a level of Explainable AI and transparency, on demand, to a user anywhere where provenance tracking can be instrumented to capture the processing flow of information through a system.

To generate the provenance narratives Rhetorical Structure Theory (RST) is used to provide a framework for achieving and maintaining coherence (Andre 2000). Each element in narrative presentation has a rhetorical relationship to another element that describes its narrative role.

This work presents three main contributions: 1) we examine the novel approach of presenting provenance models as narrative visualisations to increase understanding and trust, 2) we present models to map RST structures to W3C Provenance Data for narrative visualisations, and 3) we present a novel interactive visualisation system based on narrative visualisation that demonstrates the applicability of the proposed approach.

Keywords: *Provenance, narrative visualisation, storytelling, explainable AI, autonomy*

1. INTRODUCTION

The challenge of making sense of data in command & control (C2), data analysis, and information fusion environments is increasingly being handled by introducing automation and Artificial Intelligence (AI) to curate, integrate, and reason with the data to extract information products that will support the decision maker or analyst. To build understanding and trust in the results being presented, and indeed correct any errors that may arise, the decision maker or analyst will often need the system to explain how these results were obtained, which will involve weaving together the explanations of multiple different systems using different AI approaches.

The World Wide Web Consortium's (W3C) PROV provenance model allows a system to track the production and flow of information through the system, including: how a particular result or product was obtained; how changes to a particular input will affect the results obtained; and how a particular agent or actor in the system has been influenced by, or has influenced, other parts of the system. In a similar way, an analyst is often required to make sense of complex simulations involving many different entities and relationships. Advances in AI technologies could allow the automated identification of decisive points and significant relationships within the simulation and their impact on the results. Capturing the provenance of the analysis allows the analyst to understand and have confidence in these results and correctly apply them to their analytical question.

In a complex system, however, this provenance information can rapidly become complex in its own right. Autonomous Warrior 2018 demonstrated how the complexity of this provenance information in the TTCP AIM C2 system can be understood by transforming it into a structured narrative that allows explanations to be generated at different levels of abstraction to suit the information needs of a decision maker. Even with this structured narrative data, the question remains *how to best present this data to support the decision-making process?*

In recent years, the information visualisation community has increasingly recognised the value for data visualisation as a storytelling and communication tool. *Narrative visualisation* is concerned with how to design visualisation systems that leverage narrative techniques to communicate an underlying "data story". In this paper we present a narrative visualisation approach for displaying a structured provenance narrative that provides an intuitive visualisation of the information flows and interactions through a system, and allows a user to interact with the data to understand the story. This provides a general approach to automatically generating explanations using provenance that is independent of the domain and the individual AI components within the system.

2. BACKGROUND

Autonomous Warrior 2018 (AW18), held in Jervis Bay, Australia during November 2018 explored the military utility of emerging technologies and innovations in the Autonomy space. Amongst the technologies demonstrated was the Allied Intelligent Multi-UxV Planner with Adaptive Collaborative Control Technologies (AIM). The AIM C2 system (see **Figure 1**) was developed by an international team as part of The Technical Cooperation Program's (TTCP) Autonomy Strategic Challenge (ASC), and allows a single human operator to supervise a team of over a dozen heterogeneous autonomous vehicles, a task traditionally requiring teams of operators (Draper *et. al.* 2018). During AW18, AIM was trialled and evaluated using a combination of both real and simulated vehicles executing various use cases that explored complex real-world military scenarios.

In a complex C2 system such as AIM, a change in a single piece of information can have unexpected flow-on effects on: the fused picture displayed; the operating constraints applied to the vehicles; what tasking options are available to the operator; and what the expected outcomes of these options may be. From these changes in their tasking, it is difficult for the operator to keep track of the effects, and understand what may have caused an unexpected change or response in the system. Furthermore, these changes often occur while the operator is occupied on other tasking or systems that require different mental models than that to which the change has taken place.

A provenance tracking system has been developed in AIM to track the dependencies between the information flowing through the system as it is transformed and fused by the various AI components. In order to understand the significance of any particular change, the operator may need to switch context and associated mental models as they deal with multiple, varying, vehicles and changes. To achieve this, multimedia

narrative is used to ‘set the scene’ and explain these changes, so that the operator can understand the context of an unfolding situation and make the appropriate decisions to achieve the mission goals (Wark *et al.* 2017) (see **Figure 1**).

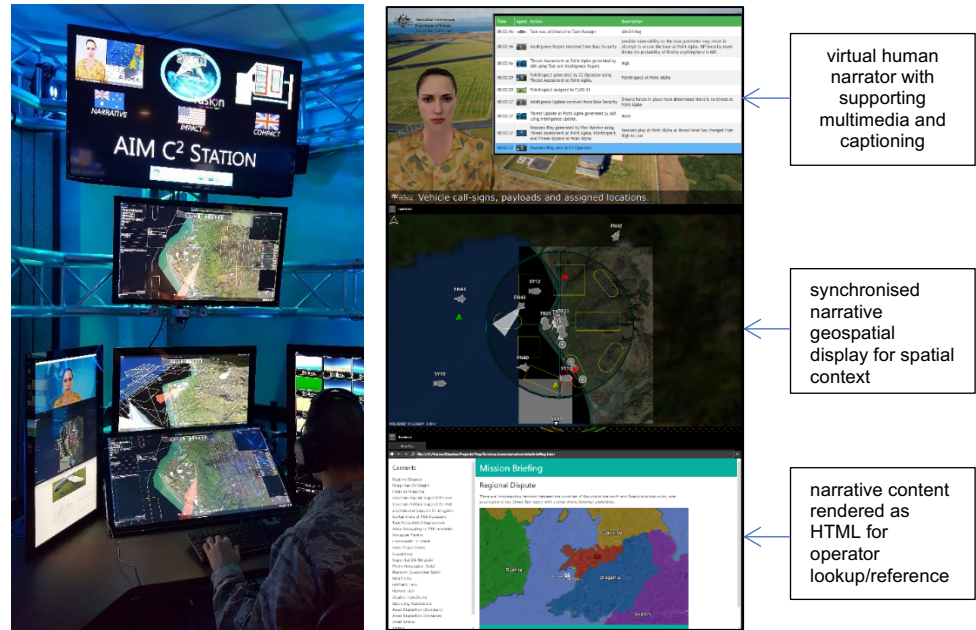


Figure 1. The AIM System allows a single operator to control a team of heterogeneous, real, and/or simulated, autonomous vehicles. (Left) Operator sitting at C2 Station. (Right) The Multimedia Narrative system displayed on the explanation display on AIM’s C2 Station.

2.1. Provenance

Provenance provides a record that describes the people, institutions, entities, and activities involved in producing, influencing, or delivering a piece of data or a thing. The W3C Provenance Specification, offers a standardised data model to allow the capture, translation and interchange of provenance information between systems (Moreau *et al.* 2013). In a complex C2 environment, many different systems interact to produce, ingest, transform, and aggregate the data used for producing some action or output.

In the AIM C2 System, the provenance of information is tracked and recorded by a component that is connected to the central messaging hub. The provenance system listens to the many different components connected to the hub as they interact and enact change. The provenance information is recorded using a PROV Java library, which implements the W3C Provenance Specification. The PROV data model enables recording the relationships between *entities* (reports/assessments/tasks), *activities* (actions/events) and the *agents* (software components/vehicles/people) in the system (see **Figure 2**), describing how entities were created, changed or delivered (Lucero *et al.* 2018). This PROV structure not only provides a framework for explaining how a particular report, assessment or task was obtained, but in turn allows us to track the corresponding flow-on effects of any particular change.




Figure 2. PROV data model is used to capture the relationships and interactions between Agents, Activities and Entities in a system

Provenance information in a complex system, such as AIM, can become very complex in its own right. To deal with this complexity the details must be managed and presented appropriately to allow the decision maker or analyst to gain the key insights and understanding required for the current task.

2.2. Narrative Generation

Narrative is a common method employed by humans to effectively convey understanding of complex situations to each other. Narrative immerses the audience within a situation and allows them to understand the main events and actors, the important relationships between them, and what the consequences may be (Gershon & Page 2001). Multimedia narrative coordinates spoken or written narration, with other display modalities are used to rapidly convey key features and relationships. Different display modalities can be chosen to convey the appropriate level of trust in the information provided.

As with human interactions, a textual narrative may not always be the most appropriate method of conveying information. The level of detail needed in the narrative depends on the information requirements of the target audience, based on their prior experience and interaction with the system. Ideally the narrative should reduce to traditional visualisations as the audience becomes contextually situated in the ‘story’. This is achieved in AIM by maintaining a model of the operator’s requirements, prior knowledge, and interactions with the narrative system.

The narrative system we implemented in AIM uses Rhetorical Structure Theory (RST) to provide a framework for achieving and maintaining coherence within a multimedia presentation (Andre 2000). Each element in the multimedia presentation has a rhetorical relationship to another element that describes its narrative role. A multimedia presentation structured in this way forms a graph, or more precisely a hierarchical tree, that links down to the central conceptual element of the presentation. Various pruning strategies can be applied to the graph to produce coherent narratives with different levels of detail that are tailored to meet the current information requirements of the operator (Wark *et al.* 2016). This becomes particularly important when generating a provenance narrative, as the size of a provenance graph can be prohibitively large and complex.

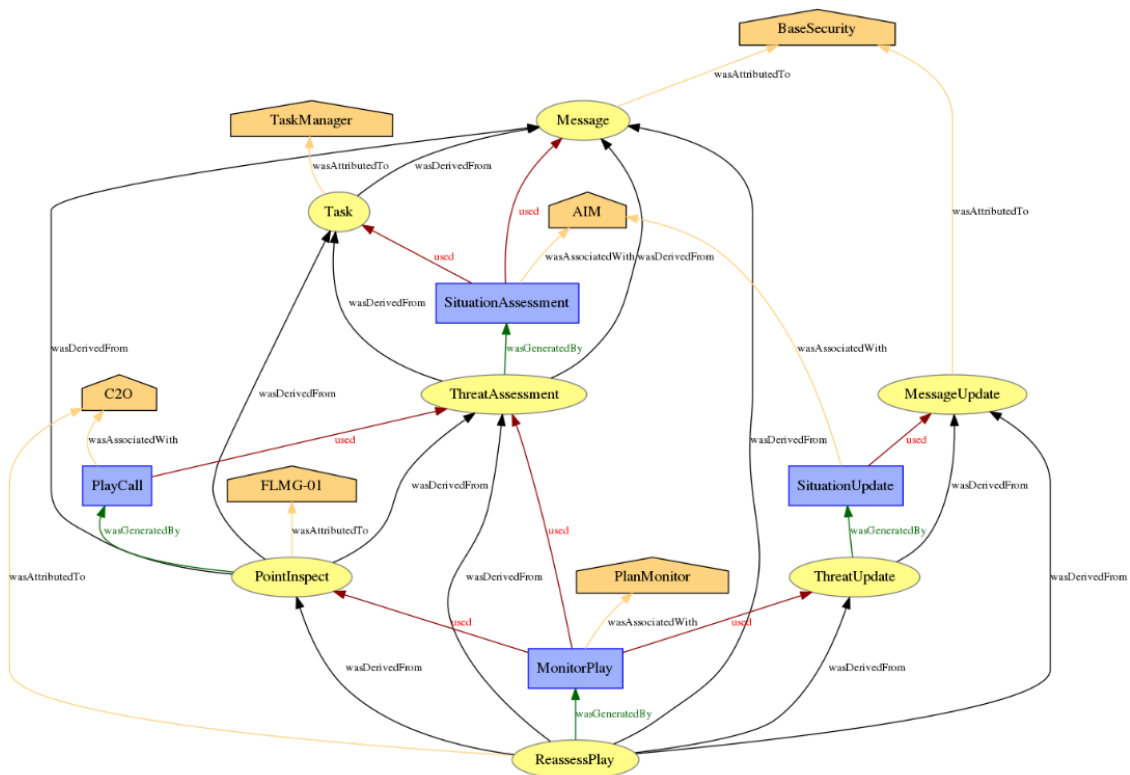


Figure 3. PROV graph for a reassessment task in AIM due to changes in intelligence message.

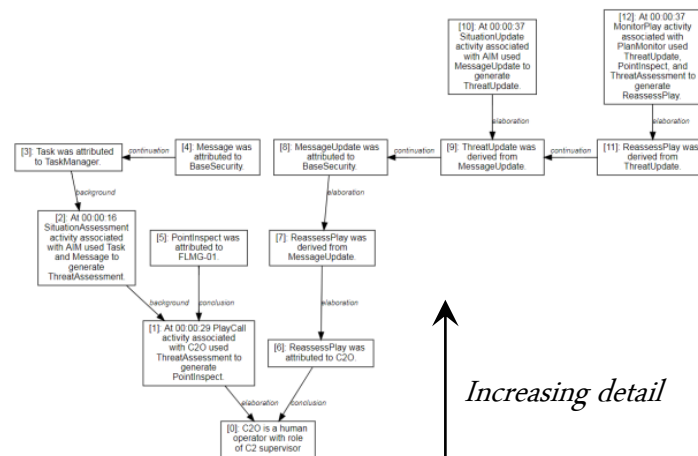


Figure 4. Corresponding (skeleton) narrative RST structure shown from a C2 Operator (C2O) agents perspective for the tasking reassessment provenance graph shown in Figure 3. The actual narrative content is not shown here, skeleton structure only for brevity and illustration purposes

To generate a provenance narrative the records captured in the PROV model in AIM are used to automatically construct a narrative explanation of the changes. A narrative generation module maps the PROV representation to a corresponding RST structure (see **Figure 3** and skeleton in **Figure 4**). The RST narrative structure represents different levels of detail in the narrative where the narrative can be generated from different perspectives, such as from a particular PROV ‘agents’ or ‘entities’ perspective.

3. NARRATIVE VISUALISATION OF PROVENANCE

Narrative visualisation commonly refers to data visualisation systems that employ elements of narrative discourse in order to convey information. Effective communication in narratives involves relaying the *motivation* of entities within the narrative, and communicating the *causality* of events that have taken place. In this work, we explore the use of narrative visualisation approaches to display the structured narrative of a PROV model. Our aim is to provide an intuitive visualisation of the information flows and interactions through the system, and allow the user to interact with the data to understand the story behind it.

The visualisation is designed around a horizontal timeline view, where events are placed on the timeline in a chronological layout that can be controlled by the user. Agents and objects within the visualisation are given their own row in the visualisation. Icons for the agents and entities appear at the point where they begin interacting within the scenario. Interactions are represented on the timeline as boxes with a narrative description of the interaction. This allows an analyst or decision maker to visually scan along a particular agent or object’s timeline to understand how they interact in the scenario.

The visualisation system consists of two main parts: a narrative engine to model the scenario, and the visualisation itself (see **Figure 5**). The model consists of a simplified representation of the W3C PROV structure, modelling the scenario on the “active” agents to take action, the activity performed, and the agents who are the passive recipients of any activity. This is referred to as a subject-verb-object tuple, where some subject has taken action with or against some other object. Each entry on the timeline can contain a mix of these tuples, defining numerous different activities that have occurred relating to each other.

The narrative model for the visualisation is generated using a two-stage approach. The first stage generates the RST narrative structure from the W3C Prov Data Model provenance records. This is done by recursively following the PROV relationships associated with the agent or entity of focus and mapping them to an RST structure. The mapping chosen was designed to produce a hierarchy that summarised the derivations of entities at different levels in the generated hierarchy. Each level of the hierarchy provides progressively more information about the agents and entities that have influenced the targeted agent or entity. The second stage uses both the RST narrative structure and the original PROV records to construct the narrative model required for the visualisation. Two passes are used to generate all the subject-verb-object tuples for the model, as well as the necessary expansion information to be displayed in the visualisation. In the first pass the RST narrative structure is traversed to provide the path through the data, here we can calculate a score for

each respective ‘agent and ‘entity’ based on their depth and position in the RST hierarchy. The score generated can later be used, for example, as a means of filtering in the visualisation. The first pass also marks the initial set of ‘agents’ and ‘entities’ of interest. The ‘of interest’ marking can, for example, be used to determine which elements should appear expanded first in the visualisation. The second pass decorates the model with all the additional information contained in the provenance records in the PROV model. This allows expansion information to be displayed in the visualisation, for example, on hover over we can display a baseball card view of an ‘agent’ or ‘entities’ properties.

Analysts are initially presented with a collapsed view of the PROV graph based on the RST narrative structure generated. With high-level agents listed down the left side of the screen, and the timelines for each agent extending horizontally, the user can scroll the timeline to view the activities they were involved in. In this view, the analyst is able to see the agents that initiated the scenario as well as the outcome. From here,

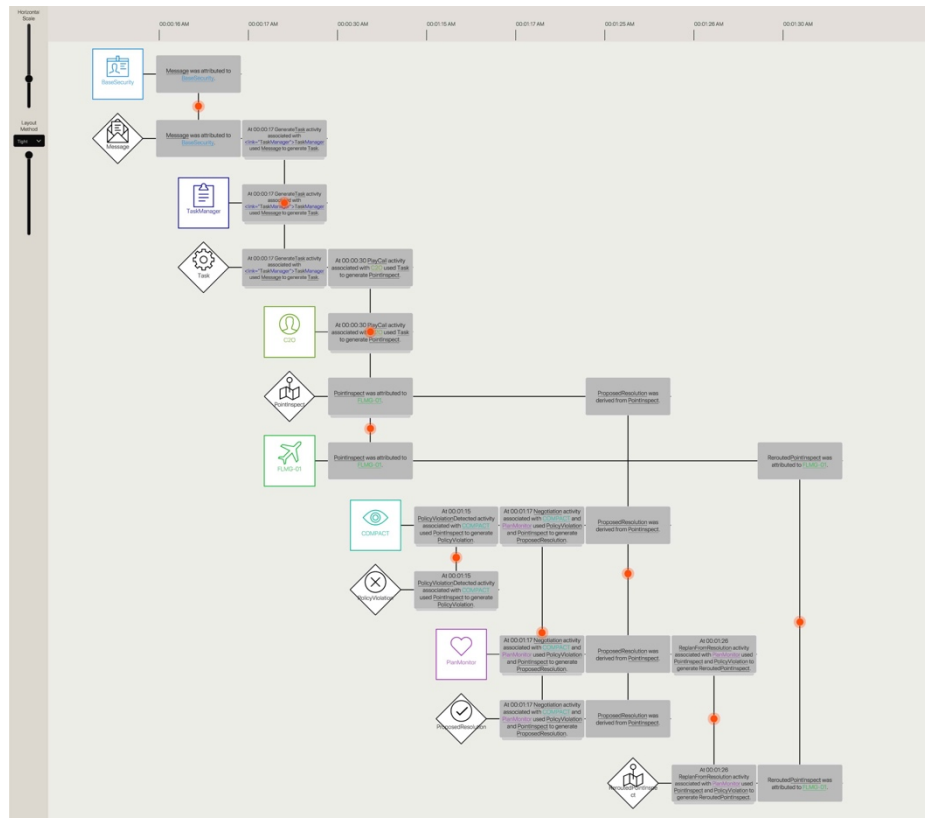


Figure 5. The interactive Narrative Visualisation system showing the same provenance information represented in Figure 3 and Figure 4 visualised.

the analyst can click on activities and the graph is expanded to show the next level of agents and timelines involved with that activity, thereby interactively revealing more detail into the constituent components of the graph. Detailed metadata information is available for each agent and entity via roll over with a summary card view.

The system enables a hierarchical view of agents, for example, we can prioritise people above systems, and systems above messages. This means that upon clicking an activity on a person’s timeline, the activity expands to show the systems involved. Click the same activity again on either the original person’s timeline, or any of the newly revealed system timelines, shows the message agents and their timelines. This allows the activities, and the agents involved, to be abstracted to only the highest level desired by the user.

This provides a general approach to automatically generating explanations using provenance that is independent of the domain and the individual AI components within the system.

4. DISCUSSION

AW18 demonstrated how provenance narratives could be generated from PROV records and presented to the operator using multimedia narrative. The use case chosen for our initial demonstration of this capability was

deliberately simple. Intricate enough to demonstrate the value of such narratives when there are a number of different interacting agents, entities and activities occurring in the system, but manageable enough in size and complexity to avoid any scalability issues.

This demonstration was based around a use-case of tracking the flow-on effects of updates to an intelligence report. Reporting to the operator, the consequent need to reassess tasking for an air vehicle, and explaining why this tasking should be reassessed. A number of different options are available for presenting these narratives to the AIM operator. Our initial implementation rendered the timeline of activities in a table while a virtual human provided a narration describing the sequence of activities and highlighting any associated locations or physical actors on a geospatial display. Given the AIM operator may need to attend to multiple events concurrently, the system provided them with a notification when the provenance system has detected a change in the system state and a narrative is available for presentation.

A limitation of presenting the provenance narrative in this way is that for a large provenance structure a substantial amount of culling may be required to present a manageable, time constrained, narrative to the user. The RST narrative structure allows culling of content while maintaining coherence, but, while the narrative may still convey the key points of the provenance, a culled version may also miss details that could prove useful in better understanding and instilling the appropriate level of trust in the result or output generated. These limitations have led us to explore the use of an interactive narrative visualisation approach to provide an alternative, yet complimentary, method of presenting the provenance narrative to a user that should scale to far larger and more complex provenance narratives. In future a combination of these two approaches, may be used.

5. CONCLUSION

This paper discusses preliminary work exploring how provenance narratives can be dynamically generated and presented in an interactive narrative visualisation system to help build understanding and trust in the results presented. While the example presented uses a complex C2 system as the source of the provenance information, the solution is generic and flexible, designed to allow the generation and presentation of provenance narratives for any system that adheres to the W3C Provenance Data Model. Future work will explore how well this solution scales with arbitrarily large and complex prov structures and how the narrative visualisation can be improved to provide an analyst or decision maker key insights and understanding from the information presented.

ACKNOWLEDGMENTS

The authors wish to acknowledge the Data To Decisions Cooperative Research Centre for supporting the development of the Narrative Visualisation system as part of the Integrated Law Enforcement program, and The Technical Cooperation Program's Autonomy Strategic Challenge team that developed the AIM C2 system. Particular thanks go to our provenance collaborators at NIWC PACIFIC, Crisrael Lucero and Douglas Lange for their contributions.

REFERENCES

- Andre, E. (2000). The Generation of Multimedia Presentations. Handbook of Natural Language Processing, Marcel Dekker Inc., New York NY, 314-338.
- Draper, M., et al. (2018). Realizing Autonomy via Intelligent Adaptive Hybrid Control: Adaptable Autonomy for Achieving UxV RSTA Team Decision Superiority. Technical Report
- Gershon, N., Page, W. (2001). What Storytelling Can Do for Information Visualization. Communications of the ACM, 44(8): 31-37.
- Lucero, C., Coronado, B., Hui, O., Lange, D.S., 2018. Exploring Explainable Artificial Intelligence and Autonomy through Provenance, 2nd Workshop on Explainable AI, IJCAI-ECAI-18, 85-89.
- Moreau, L. et al. (2013). PROV-DM: The PROV Data Model. In W3C Technical Reports.
- Wark, S., Nowina-Krowicki, M. (2016). Generating Multimedia Narrative for Virtual Humans, Proceedings of the Australasian Computer Science Week (ACSW) '16, 2-5 February, Canberra, Australia.
- Wark, S., Nowina-Krowicki, M., Dall, I., Chung, J., Argent, P., Bowering, G. (2017). Multimedia Narrative For Autonomous Systems, Workshop on Impedance Matching In Cognitive Partnerships (CogPartner), IJCAI-17, (19-25 August), Melbourne, Australia.