A Demonstration of Generative Policy Models in Coalition Environments

Daniel Cunnington[1], Graham White[1], Mark Law[2], and Geeth de Mel[1]

[1] IBM Research, UK
{dancunnington,gwhite,geeth.demel}@uk.ibm.com
[2] Imperial College London, UK
mark.law09@imperial.ac.uk

Abstract. Autonomous systems are expected to have a major impact in future coalition operations to assist humans in achieving complex tasks. Policies are typically used by systems to define their behavior and constraints and often these policies are manually configured and managed by humans. This paper presents a demonstration of a recent Generative Policy-based Model (GPM) approach applied to generating coalition policies for asset serviceability. This demonstrates the flexibility of the approach for generating policies in a distributed coalition environment to facilitate effective collaboration between coalition partners.

Keywords: Generative Policy Models · Generated Policies · Symbolic Learning · Coalition Environments

1 Introduction

In the future battlespace, coalition systems and devices will be required to operate in challenging environments that impose certain constraints such as lack of or low-bandwidth connectivity to back-end services, rapidly changing environmental conditions and the requirement to abide by legal regulations and mission directives [5]. Therefore, coalition systems and devices require the capability to adapt and evolve such that they can behave autonomously 'at the edge' in previously unseen contexts. Crucially, systems need to understand the bounds in which they can operate based on their own (and that of other systems) capability, constraints of the environment and safety requirements. According to Bertino, et al. [1], policies can be seen as directives given by a managing party to one or more managed parties in order to guide their behavior in coalition missions and collaborative activities. Policies are usually expressed as technology independent rules aiming to enhance the hard-coded functionality of the managed parties by the introduction of an interpreted logic that can be dynamically changed without modifying the underlying implementation. Policy-based management therefore significantly increases the self-managing aspects of coalition operations. However, conventional approaches to policy-based management that are typically based on a top-down approach by which policies are specified through multiple refinement steps by a centralised policy administrator do not transfer effectively to the military environment. This is especially true of distributed, coalition scenarios where centralized policy management may not be possible due to environmental constraints. This paper presents a demonstration of recent work in GPM architectures showcasing their applicability to generating policies in a distributed, coalition environment. The paper is structured as follows. In Section 2 we outline the main purpose of our work, in Section 3 we present the associated demonstration and finally we conclude in Section 4.

2 Main Purpose

The main purpose of our work is to develop a flexible, secure and scalable generative policy architecture that is capable of meeting the complex requirements of distributed coalition environments. Within the International Technology Alliance for Distributed Analytics and Information Sciences (DAIS-ITA), we have proposed the notion of Generative Policy-based Models (GPMs) [7], where managed parties are provided with an initial policy specification. Each managed party can dynamically generate its own policies from the initial specification and possibly evolve its policies over time based on a change in operating context. As part of our work we have explored using a symbolic approach for learning policy models [2, 3], utilising state-of-the-art Answer Set Grammars (ASGs) [6].

Within a coalition environment, coalition partners are required to work together which may involve sharing assets with one another. An asset could be a physical device such as a surveillance camera, an autonomous system such as an Unmanned Aerial Vehicle (UAV) or an asset could be a virtual service such as a database or a specific type of machine learning model such as a face recognizer or weaponry detector. Given the constantly evolving context and potentially dynamic asset pool caused by coalition partners frequently bringing assets online/offline during a mission, local and global serviceability policies that define the type of requests that an asset or coalition partner can service will be constantly in flux and it wont be possible for a human operator to generate all of these policies for all assets in all possible contexts in a timely manner. Also, given the distributed nature of a partners assets, the coalition partner as a whole requires means to keep track of the serviceability of its asset pool through a generated policy and also is responsible for communicating this updated policy in a distributed manner to its assets.

3 Demonstration

For the purposes of demonstration, we have utilised an agile knowledge representation framework alongside synthetic ground truth generation [4] in order to create a dataset of annotated asset requests, which is released online[3]. This approach enables customised annotations of ground truth data to facilitate complex

[3] https://github.com/dais-ita/coalition-data/

evaluations of GPMs. In our demonstration, the United Kingdom (UK) is working alongside the United States (US) and a fictitious third party 'Kish' as part of a coalition undertaking a particular mission. Other assets currently present in the environment have learned to accept or reject serviceability requests from other coalition parties and each asset stores a log of annotated request examples, with respect to a given context and its own capabilities and constraints. Initially these annotations could be captured by monitoring human behaviour or configured by a managing party. The UK then introduces a new autonomous UAV into the environment and in order to enable the UAV to function as part of the coalition, the UAV is required to learn a GPM such that it can generate policies that govern its behaviour in varying contexts, i.e. autonomously decide whether or not to service a request from another coalition party. Firstly, the UAV communicates with nearby assets to obtain annotated serviceability request logs. The UAV then constructs and solves an ASG that represents a GPM which can be used to generate a set of valid behaviours, i.e. requests that can be serviced within a given context. The demonstration shows the interaction with and querying of the generated policy in natural language whilst also explaining how the GPM can be used to generate a new set of valid behaviours given a different contextual environment. Figure 1 shows the UAV rejecting a serviceability request due to a no-fly zone, as indicated by its generated policy in this context. Having learned a GPM, the coalition can operate in a distributed manner, where assets can behave autonomously to service requests for various capabilities. Also, this can be extended to make use of a Software Defined Network to generate a coalition-wide policy that defines how different requests can be serviced by different coalition devices, without using a centralised asset management system.

4 Conclusions

This demonstration paper has presented a symbolic, ASG-based GPM applied to distributed coalition environments. The demonstration shows how real-time constraints that occur in the operating context such as adverse weather conditions, restricted operating capability and evolving trust relationships between coalition partners can be learned in order to generate policies for coalition devices to collaborate effectively in distributed environments.

Acknowledgement

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.