Towards a Coalition Focused Neural-Symbolic Generative Policy Model

Dan Cunnington (IBM UK), **Alessandra Russo** (Imperial), **Seraphin Calo** (IBM US), **Elisa Bertino** (Purdue).

Objectives

This paper performs a deep-dive into the emerging field of neural-symbolic machine learning with a view towards enabling future neural-symbolic generative policy models that are capable of analysing both structured and unstructured data present in a coalition operation, whilst providing full transparency and enabling edge of network reasoning capability. Specifically, a technique called DeepProbLog is investigated and applied to a coalition scenario, highlighting key research challenges and questions to be addressed in order to advance the capability of generative policy models for coalition operations.

Technical Challenges

- Enable sensor data in a coalition environment to be analysed for policy generation (e.g. CCTV data as shown in Figure 1).

- How to integrate neural models into generative policy architectures to enable a greater range of contextual understanding, whilst preserving the explainabilty benefits that the symbolic learning components provide.

Fig 1. Sample image indicating the presence of an enemy vehicle, represented as a red London bus

- Ultimately we aim to develop a fully integrated, *end-end* architecture that can perform both forward passing and backward passing over neural and symbolic components.

- Firstly we will develop a *hybrid* approach, that performs forward passing only. This paper investigates a current *hybrid* approach, DeepProbLog and compares it's performance in a pure neural approach.

Military & Coalition Relevance

Consider a Logistical Resupply scenario. The resupply convoy must autonomously make decisions about various actions within the mission, in order to mitigate the risk of adversarial compromise. In our experiments, we assume the resupply convoy has access to a CCTV camera image as shown in Figure 1 that is capable of observing enemy vehicles. Here we utilise the Transport for London (TfL) CCTV dataset[1] and consider red buses to represent enemy vehicles. (continued…)

We also assume the resupply convoy has access to a team of ground troops that may be available or unavailable at a given time. If available, the ground troops are capable of investigating enemy vehicles in a particular location.

Approach

The logistical resupply policy model which forms our learning task is as follows:

- IF *enemy_vehicle_present* AND *ground_troops_available* THEN ***deploy_ground_troops***.
- IF (*enemy_vehicle_present* AND *ground_troops_unavailable*) OR *unsure* THEN ***re_route_convoy***.
- IF *enemy_vehicle_not_present* THEN ***proceed_as_normal***.

We model these rules in probabilistic logic and train a DeepProbLog model end-end based on policy examples annotated with a resulting action. Here – the weights of a CNN used to detect red buses (enemy vehicles) and a CNN used to detect 0 or 1 MNIST images (to represent a decision on ground troop availability) are updated through back-propagation when the DeepProbLog model is trained with logistical resupply policy actions.

Results

Figure 2 details the results of comparing the DeepProbLog approach to a baseline neural approach over 30,000 training iterations on the test dataset. Both experiments were trained with the same 10,000 examples and tested with the same 500 examples. With the addition of the logical rules, you can see that the DeepProbLog approach significantly out-performs the baseline neural approach after 30,000 iterations. This outlines a promising research direction for future, neural-symbolic generative policy models, as unstructured sensor data can be analysed for policy generation within a coalition environment.

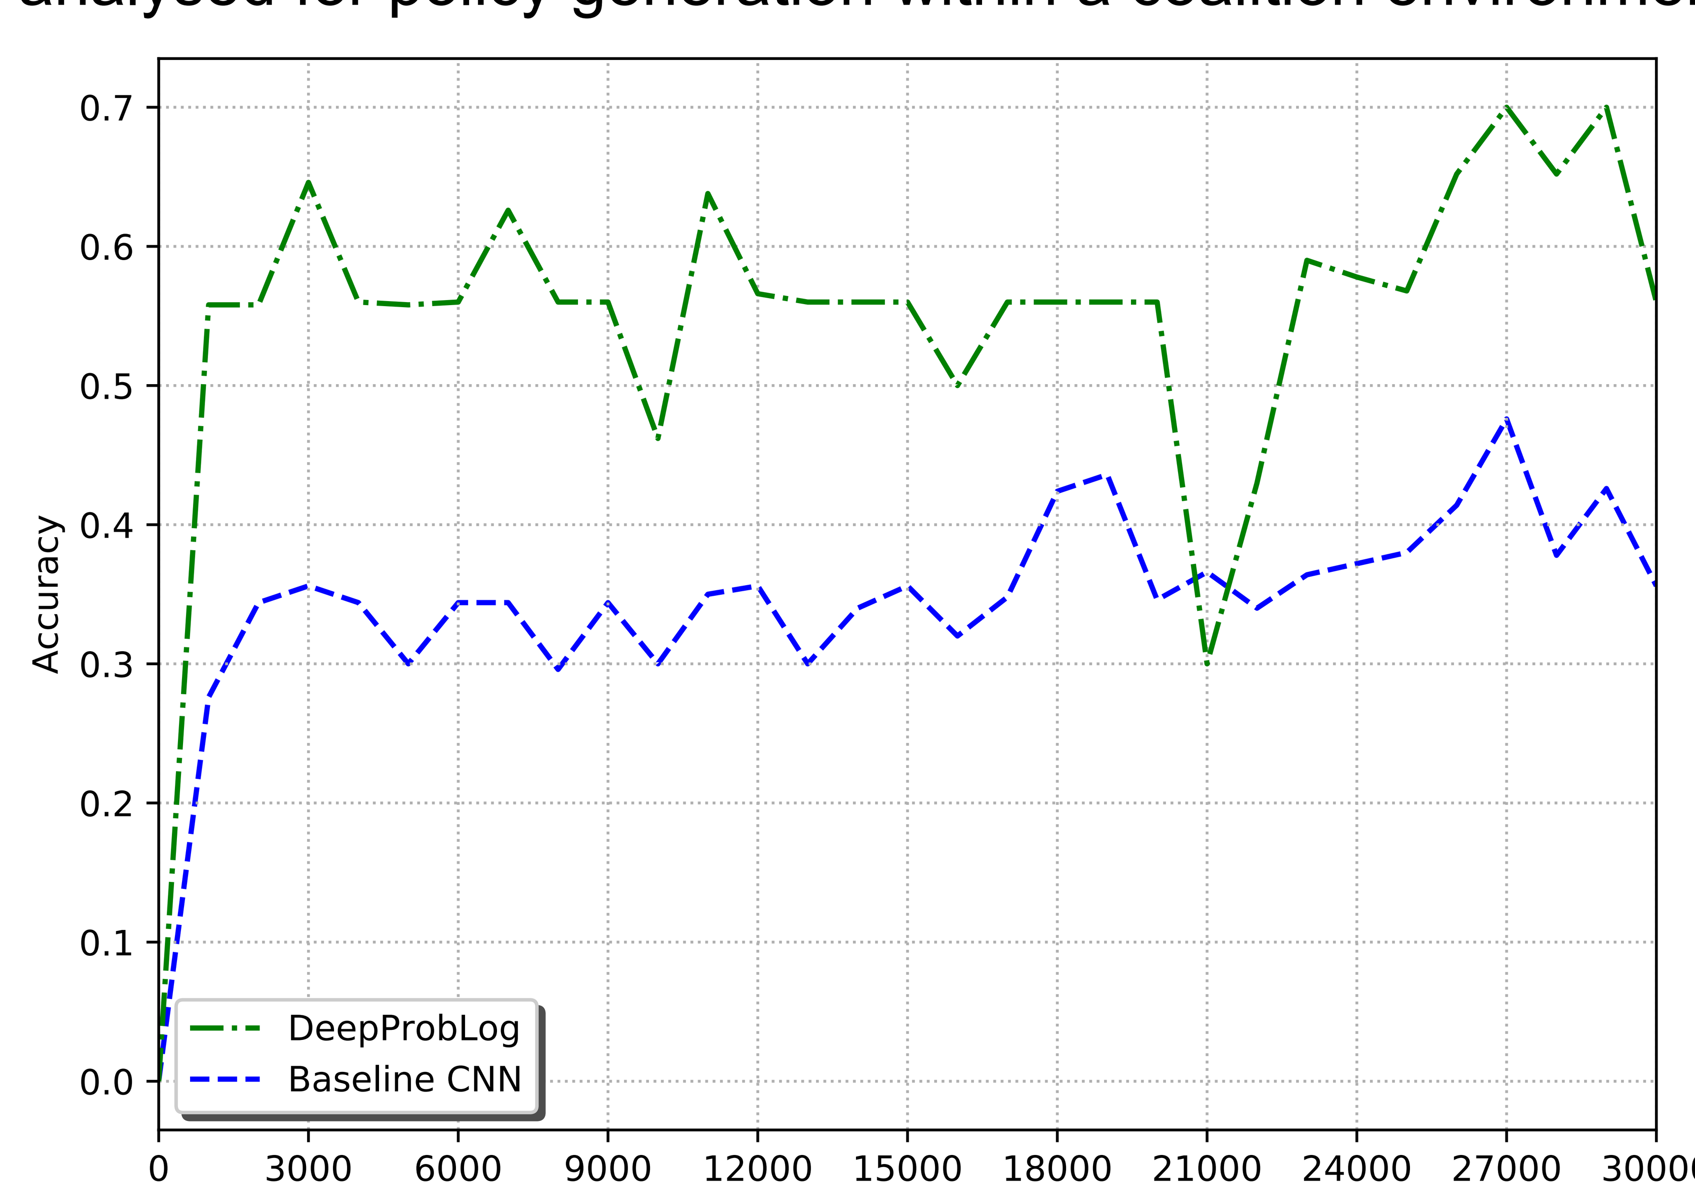

Fig 2. Accuracy of DeepProbLog compared to a baseline CNN for the logistical resupply learning task.

[1] https://data.london.gov.uk/dataset/tfl-live-traffic-cameras

Distributed Analytics and Information Science International Technology Alliance
Annual Fall Meeting, September 2019