Are You Doing What I Think You Are Doing? Criticising Uncertain Agent Models

Stefano V. Albrecht
s.v.albrecht@sms.ed.ac.uk

Subramanian Ramamoorthy
s.ramamoorthy@ed.ac.uk

Introduction

- Key for effective interaction in many multiagent systems is to *hypothesise* behaviour of other agents
- **Question:** given history H and hypothesis π_j^* for behaviour of agent j, does j really behave according to π_j^*?

  $\Rightarrow$ No universal theory to contemplate question
- If answer is no, can hypothesise alternative behaviour or resort to default strategy (e.g. maximin)

t	(a_1^t, a_2^t)	π_2^*
1	$(1,2)$	$\langle .3, .1, .6 \rangle$
2	$(3,1)$	$\langle .2, .3, .5 \rangle$
3	$(2,3)$	$\langle .7, .1, .2 \rangle$
4	$(2,3)$	$\langle .0, .4, .6 \rangle$
5	$(1,2)$	$\langle .4, .2, .4 \rangle$

a_i^t is action taken by agent i at time t

π_2^* is hypothesised action probabilities for agent 2

$\Rightarrow$ **Does agent 2 really behave according to π_2^*?**

Behavioural Hypothesis Testing

- Observe $\mathbf{a}_j^t = (a_j^0, ..., a_j^{t-1})$

  Sample $\hat{\mathbf{a}}_j^t = (\hat{a}_j^0, ..., \hat{a}_j^{t-1})$ using π_j^*

  Question: $\mathbf{a}_j^t$ and $\hat{\mathbf{a}}_j^t$ generated from same behaviour (π_j^*)?
- Decide question as frequentist hypothesis test

$$p = P\left(|T(\tilde{\mathbf{a}}_j^t, \hat{\mathbf{a}}_j^t)| \geq |T(\mathbf{a}_j^t, \hat{\mathbf{a}}_j^t)|\right)$$

  where $\tilde{\mathbf{a}}_j^t \sim \delta^t(\pi_j^*) = (\pi_j^*(H_i^0), ..., \pi_j^*(H_i^{t-1}))$
- Test statistic T based on *score functions* z_k

$$T(\tilde{\mathbf{a}}_j^t, \hat{\mathbf{a}}_j^t) = \frac{1}{t}\sum_{\tau=1}^{t} T_\tau(\tilde{\mathbf{a}}_j^\tau, \hat{\mathbf{a}}_j^\tau)$$

$$T_\tau(\tilde{\mathbf{a}}_j^\tau, \hat{\mathbf{a}}_j^\tau) = \sum_{k=1}^{K} w_k \left(z_k(\tilde{\mathbf{a}}_j^\tau, \pi_j^*) - z_k(\hat{\mathbf{a}}_j^\tau, \pi_j^*)\right)$$

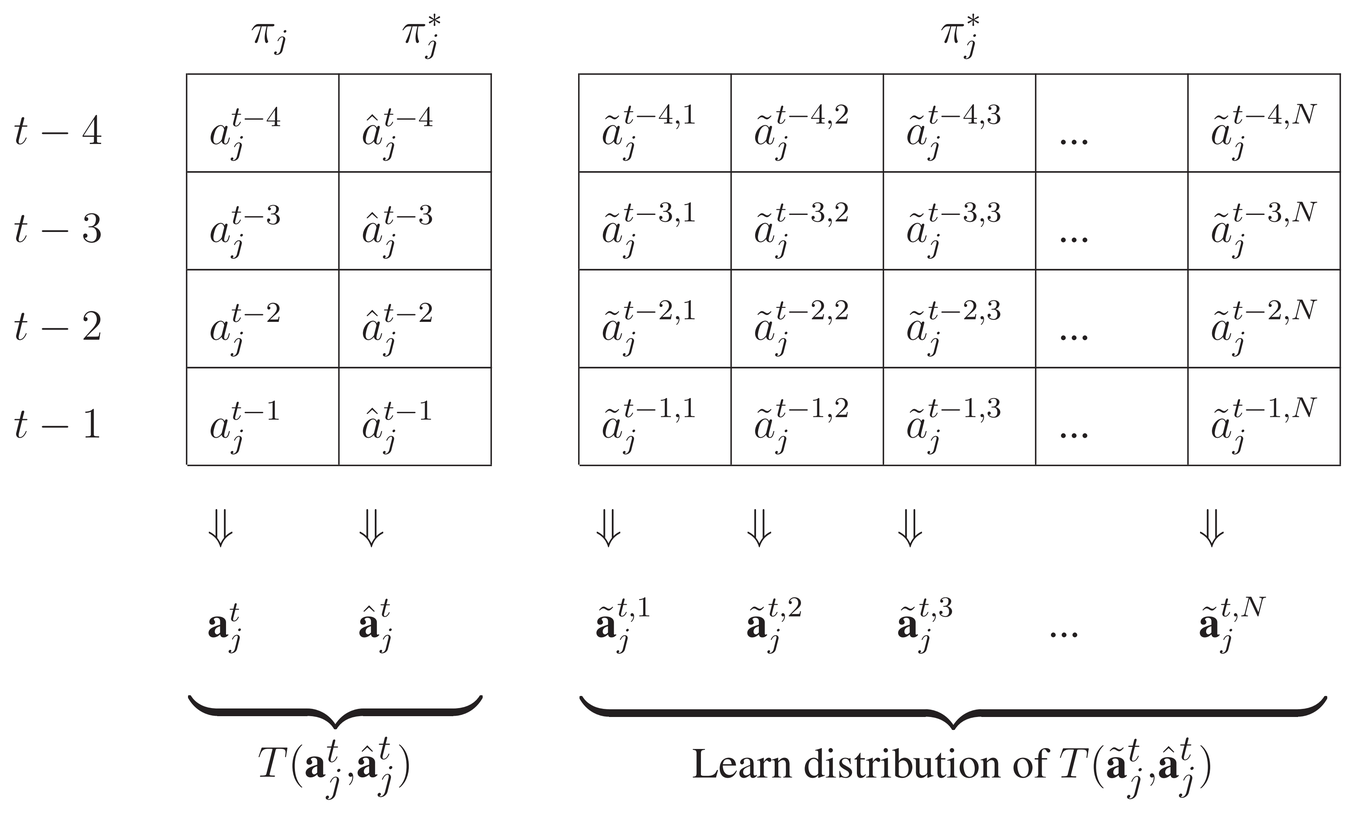

Experiments

- Three score functions:

$$z_1(\mathbf{a}_j^t, \pi_j^*) = \frac{1}{t}\sum_{\tau=0}^{t-1} \frac{\pi_j^*(H_i^\tau)[a_j^\tau]}{\max_{a_j \in A_j} \pi_j^*(H_i^\tau)[a_j]}$$

$$z_2(\mathbf{a}_j^t, \pi_j^*) = \frac{1}{t}\sum_{\tau=0}^{t-1} 1 - \mathbb{E}^{a_j \sim}_{\pi_j^*(H_i^\tau)} \left|\pi_j^*(H_i^\tau)[a_j^\tau] - \pi_j^*(H_i^\tau)[a_j]\right|$$

$$z_3(\mathbf{a}_j^t, \pi_j^*) = \sum_{a_j \in A_j} \min\left[\frac{1}{t}\sum_{\tau=0}^{t-1}[a_j^\tau = a_j]_1, \frac{1}{t}\sum_{\tau=0}^{t-1}\pi_j^*(H_i^\tau)[a_j]\right]$$

- Four classes of behaviours:

  - Random behaviours

  - LFT, CDT, CNN (Albrecht et al., 2015)

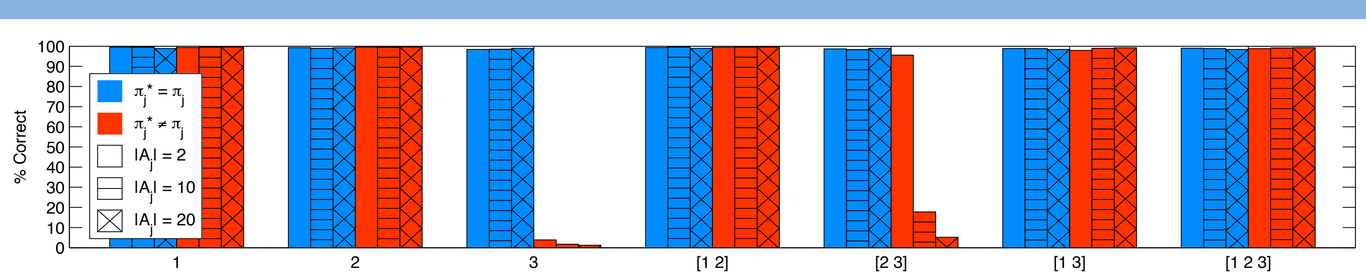

Average accuracy with random behaviours after 10000 time steps, for $|A_j| = 2, 10, 20$ and $N = 50$. X-axis shows score functions z_k used in test statistic.

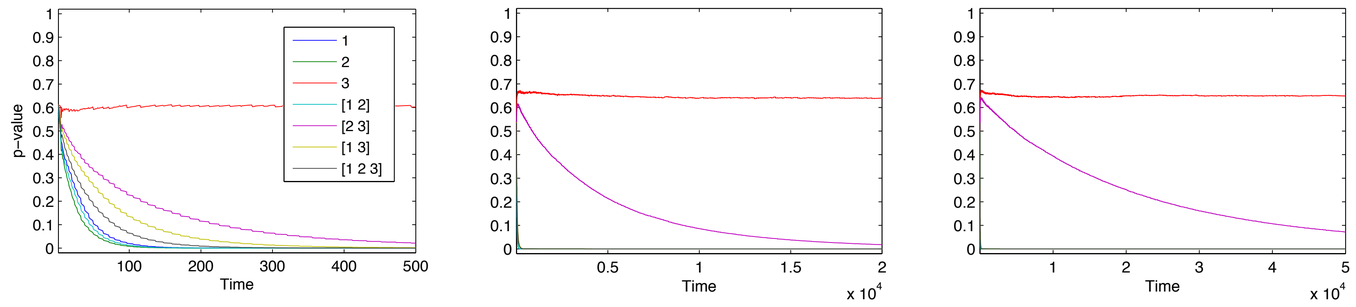

(a) $|A_j| = 2$ (b) $|A_j| = 10$ (c) $|A_j| = 20$

Average p-values with random behaviours, for $N = 50$ and $\pi_j^* \neq \pi_j$ (i.e. hypothesis wrong). Legend shows score functions z_k used in test statistic.

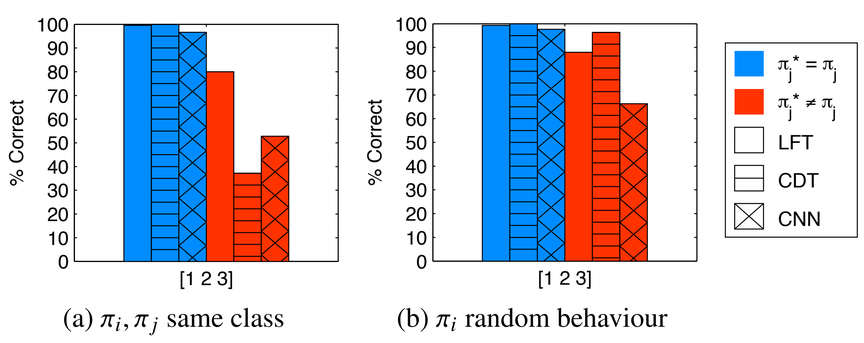

(a) π_i, π_j same class (b) π_i random behaviour

Average accuracy with behaviour classes LFT, CDT, CNN for $N = 50$. Results shown for $[z_1, z_2, z_3]$ test statistic.

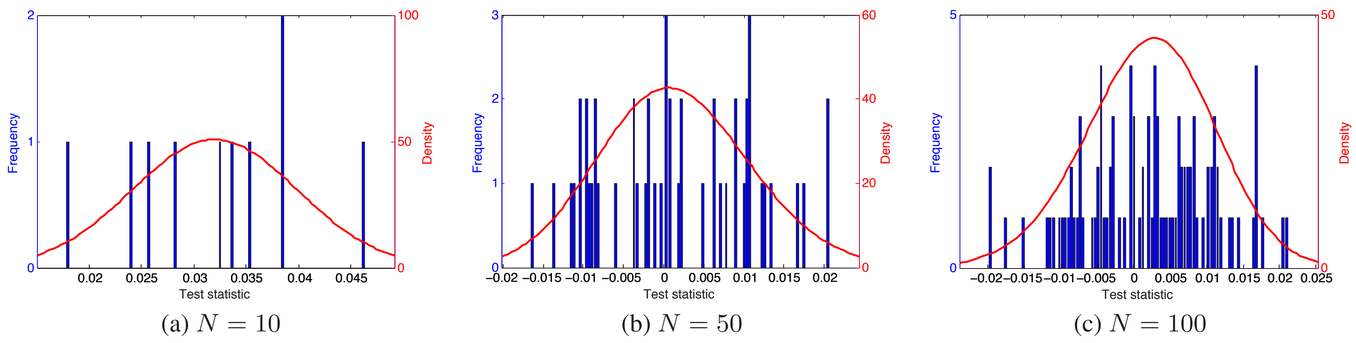

(a) $N = 10$ (b) $N = 50$ (c) $N = 100$

Example histograms and fitted skew-normal distributions (red curve) after 1000 time steps, for random behaviours with $|A_j| = 10$, using score function z_1.

References:

Albrecht, S., Crandall, J., Ramamoorthy, S. (2015). An empirical study on the practical impact of prior beliefs over policy types. In Proceedings of the 29th AAAI Conference on Artificial Intelligence, pp. 1988–1994.

See paper for more details:
http://svalbrecht.de/docs/uai15.pdf