How much can experimental cost be reduced in active learning of agent strategies?

Imperial College London

Céline Hocquette & Stephen H. Muggleton

Learning agent strategies from observations

- Experimentation requires energy, time and resources
- Automated experimentation with active learning

Learning agent strategies from observations

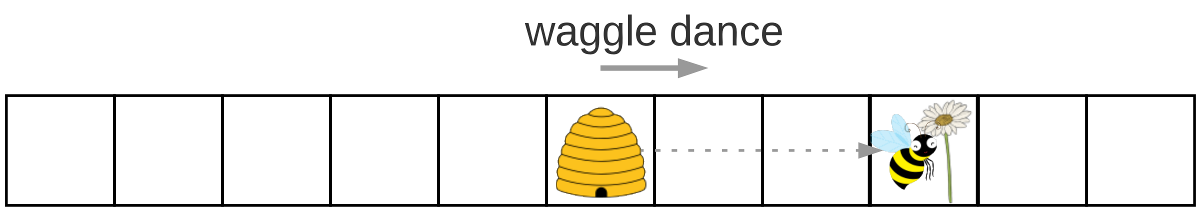

Hypothesis 1	Hypothesis 2

`f(A,B):-f1(A,C),grab(C,B).`
`f1(A,B):-until(A,B,at_flower,move_right).`
	`f(A,B):-f2(A,C),grab(C,B).`
	`f2(A,B):-until(A,B,at_flower,f1).`
	`f1(A,B):-ifthenelse(A,B,waggle_east,move_right,move_left).`

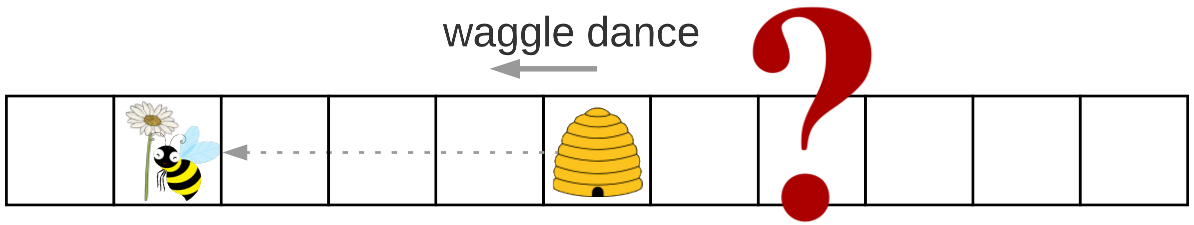

Related work

	Size of the hypothesis space considered	Active Learning	Target hypotheses learned
Robot Scientist (King et al, 2004)	Finite (15)	yes	Abductive bindings
MetaBayes (Muggleton et al, 2014)	infinite	no	logic programs
Efficiently Learning Efficient Programs (Cropper, 2017)	Reduced with Abstractions	no	strategies
Bayesian Active MIL (2018)	infinite	yes	strategies

Related work

- Active Learning
  - Widely studied for identifying classifiers
  - Other applications, among them Object Detection in Computer Vision (Roy et al., 2016), Natural Language Processing (Thompson et al., 1999)

- Relational Reinforcement Learning

Framework

- Meta-Interpretive Learning

- Bayesian prior probability distribution over the hypothesis space

- Active Learning

$$ent(e) = p\ log(p) + (1-p)log(1-p)$$

Framework

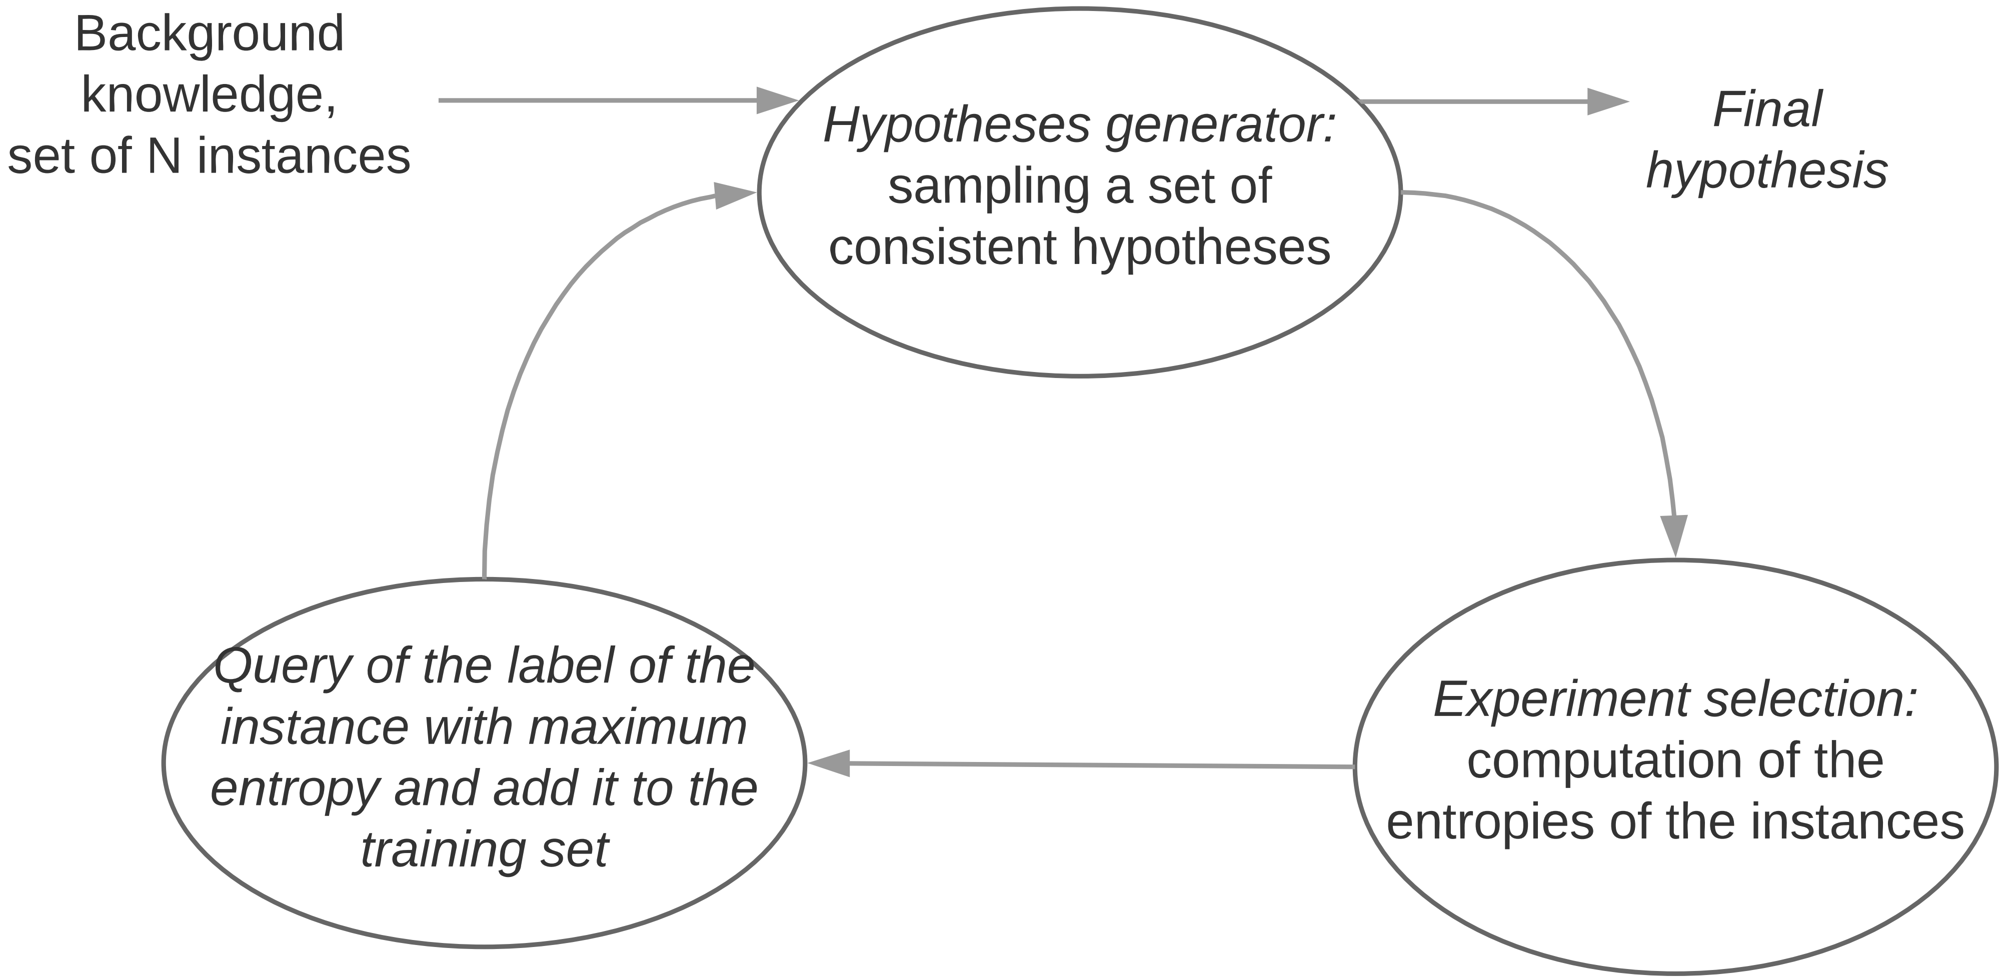
Background knowledge,
set of N instances
Hypotheses generator: sampling a set of consistent hypotheses
Final hypothesis
Experiment selection: computation of the entropies of the instances
Query of the label of the instance with maximum entropy and add it to the training set

Implementation

- Regular Sampling (MetaBayes, 2014)
- Entropy of the instances measured from the sampled set of hypotheses

Theoretical Analysis

What is the probability of selecting an instance ε-close to the entropy maximum?

- ***Active learner:*** selects the instance with maximum entropy among a set of N sampled instances

P_{active} (pi < p_{ε}) = $(1-\varepsilon)^N$ P_{active} ($p_{\varepsilon} \leq p_i$) = N ε - o(ε)

Probability distribution

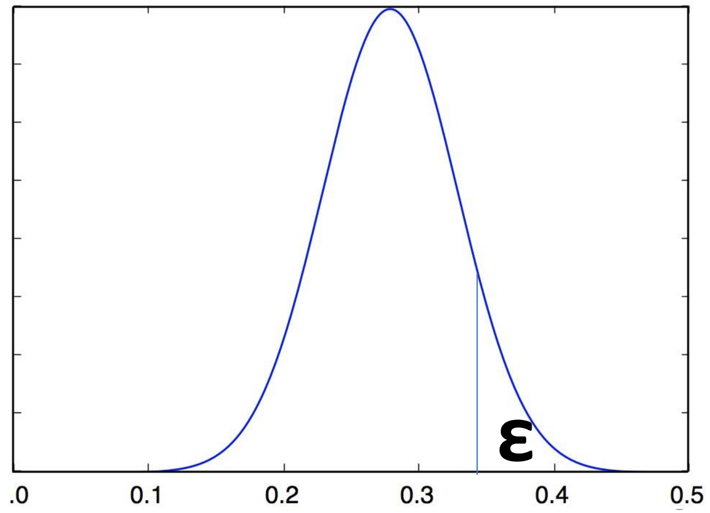

- ***Passive learner:*** random selection

$P_{passive}$ ($p_{\varepsilon} \leq p_i$) = ε

Results: Learning a Regular Grammar

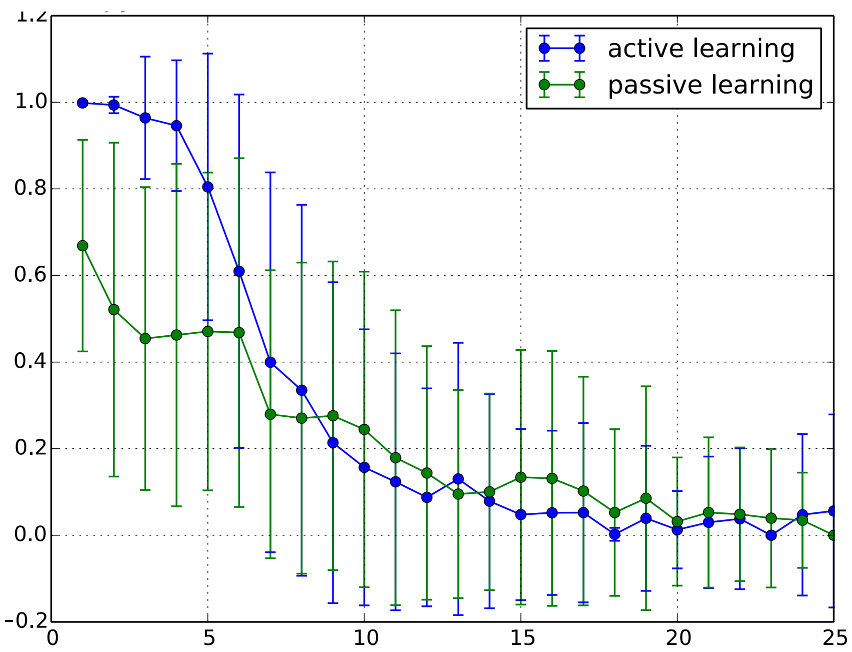

Entropy versus the number of iterations

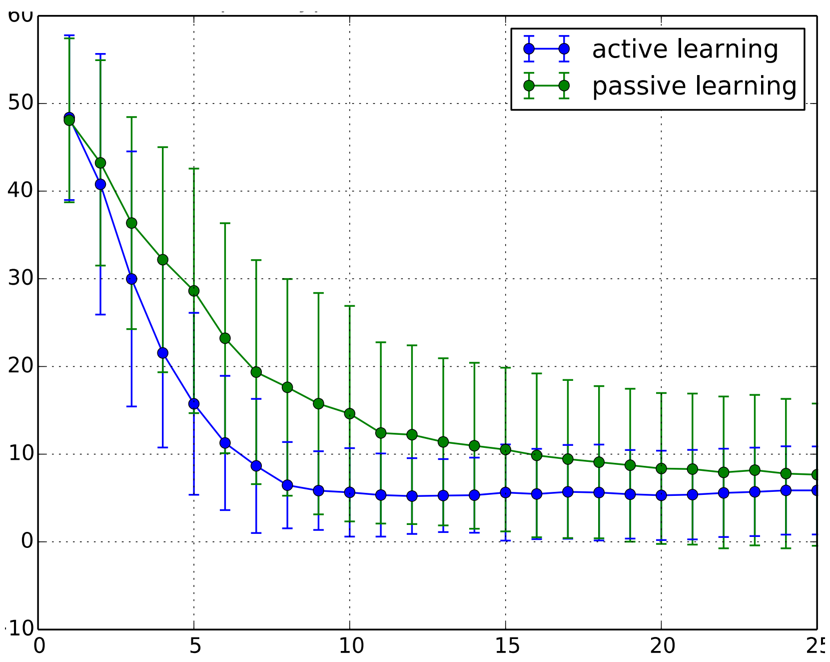

Number of hypotheses versus the number of iterations

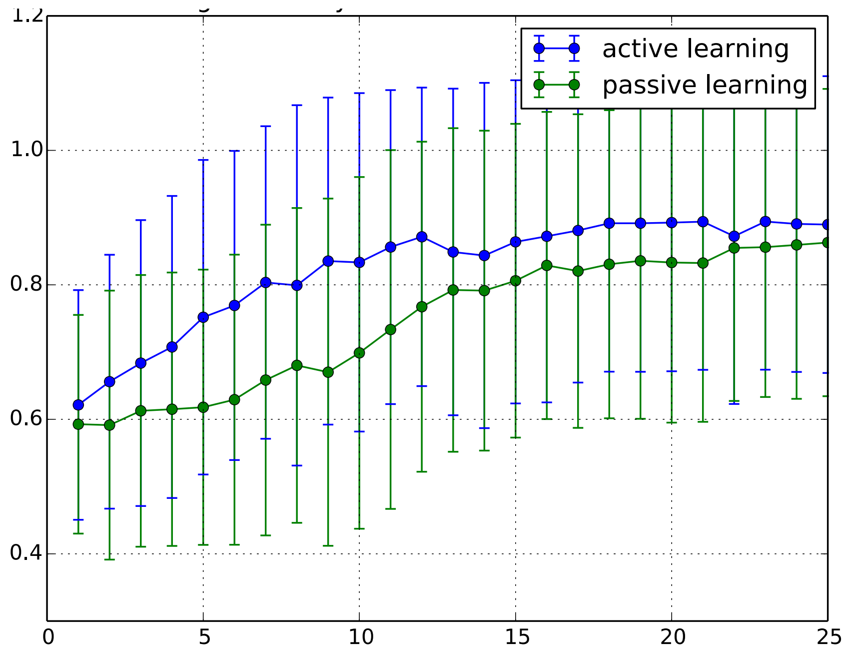

Accuracy versus the number of iterations

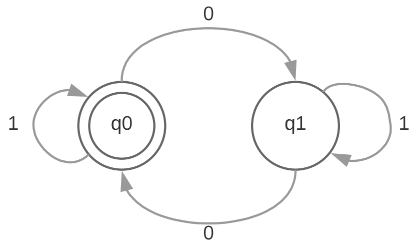


```
q0([0|A],B) :- q1(A,B).
q0([1|A],B) :- q0(A,B).
q0([0|A],B) :- q0(A,B).
q1([1|A],B) :- q1(A,B).
q0([],[]).
```

Results: Learning a Bee Strategy

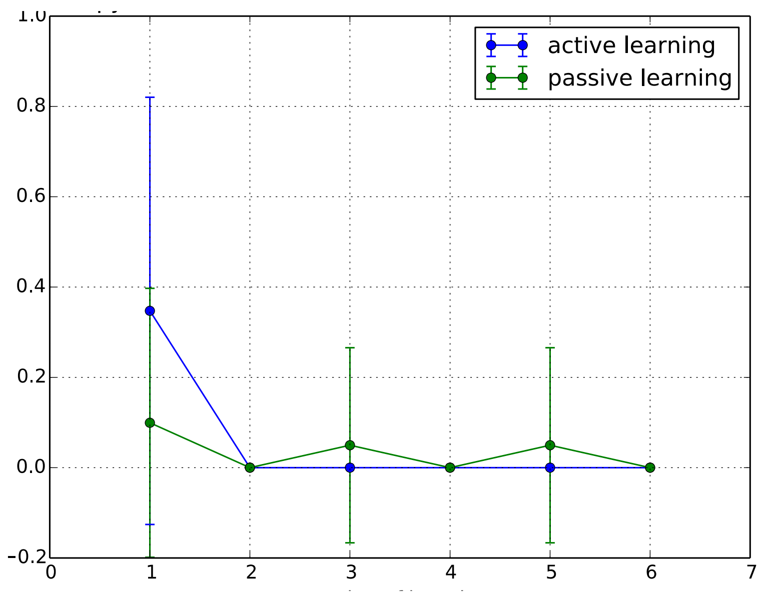

Entropy versus the number of iterations

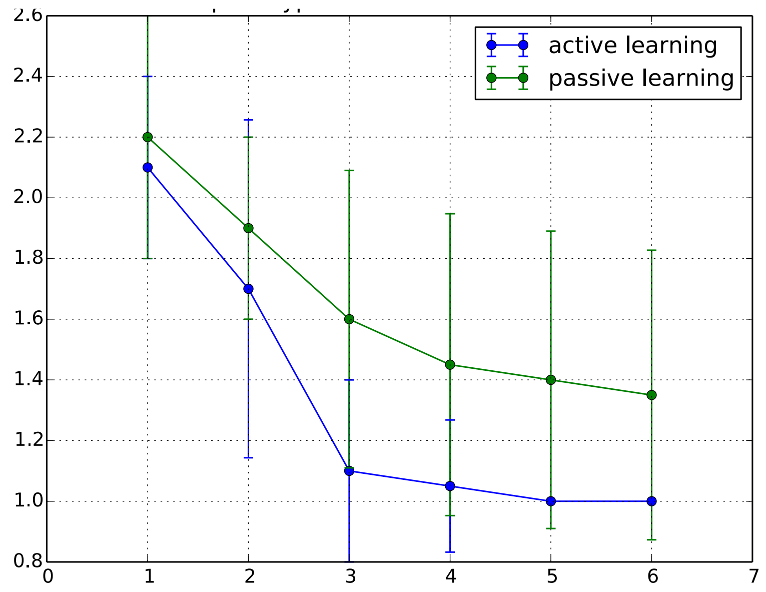

Number of hypotheses versus the number of iterations

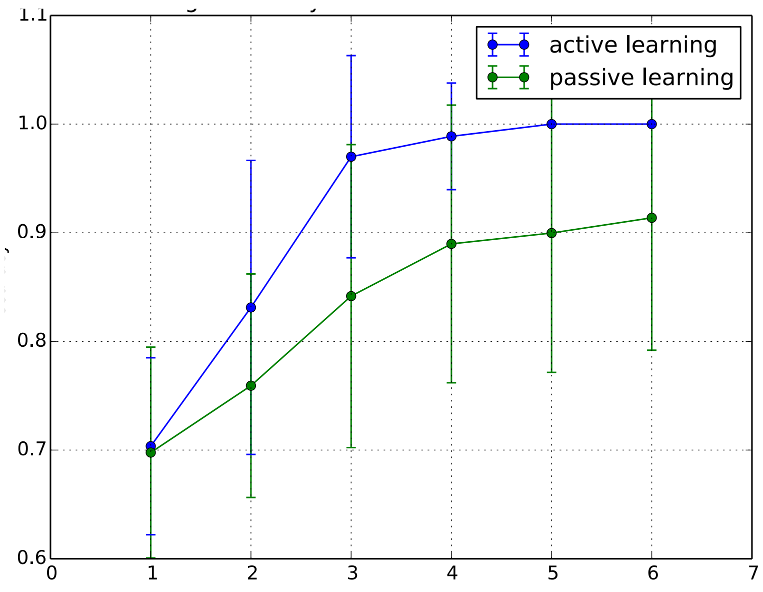

Accuracy versus the number of iterations

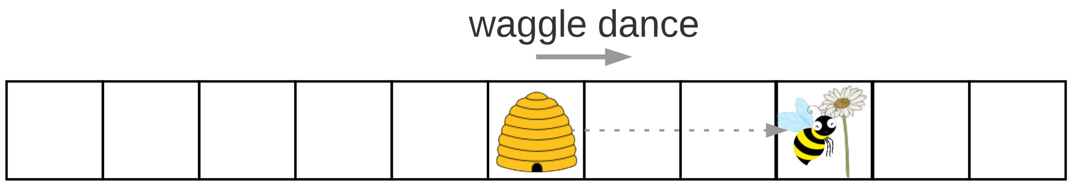


```
f(A,B):- f2(A,C),grab(C,B).
f2(A,B):- until(A,B,at_flower,f1).
f1(A,B):- ifthenelse(A,B,waggle_east,move_right,move_left).
```

Conclusion

- Automated experimentation with active learning for learning efficient strategies while making efficient use of experimental materials

- Wide range of applications such as modelling butterfly behaviors

Future work: learning probabilistic models

- Generation of SLP by Super-Imposition

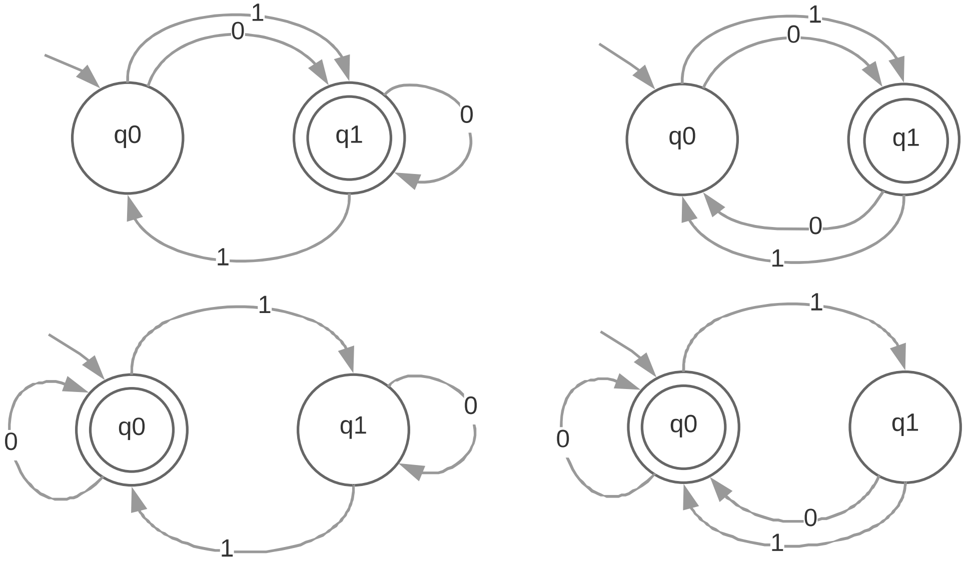

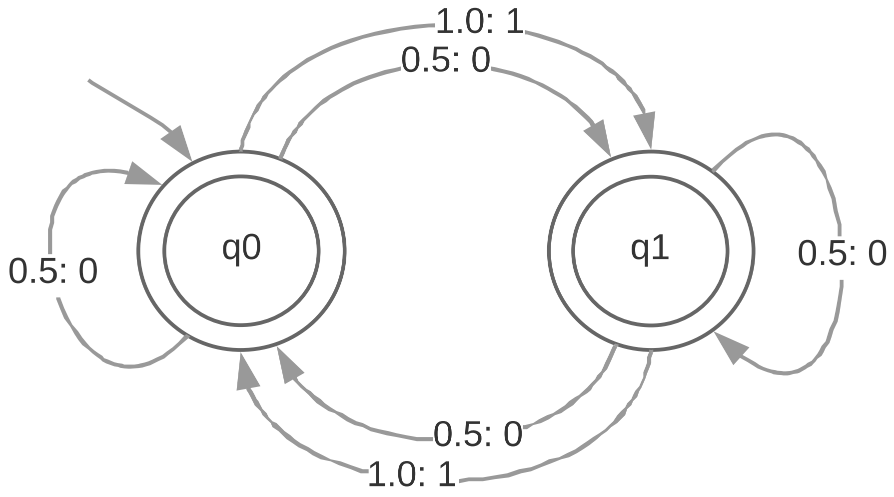

- Model scoring: sum of log posterior probabilities

$$Score(M) = \sum_{e\ in\ Test\ Set} \log(P(M|e)) = \sum_{e\ in\ Test\ Set} \log\big(P(e|M)\big) + \log\big(P(M)\big) - \log\big(P(e)\big)$$

Future work: multi-agents

- Learning a strategy for describing the behavior of an agent adapting in an evolving environment
- Applications: 2 player games

Thank you

celine.hocquette16@imperial.ac.uk
s.muggleton@imperial.ac.uk

References

- A. Cropper. Efficiently learning efficient programs. PhD thesis, Imperial College London, 2017.

- R.D. King, K.E. Whelan, F.M. Jones, P.K.G. Reiser, C.H. Bryant, S.H. Muggleton, D.B. Kell, and S.G. Oliver. Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature*, 427:247-252, 2004.

- S.H. Muggleton, D. Lin, J. Chen, and A. Tamaddoni-Nezhad. Metabayes: Bayesian meta-interpretative learning using higher-order stochastic refinement. In Gerson Zaverucha, Vitor Santos Costa, and Aline Marins Paes, editors, *Proceedings of the 23rd International Conference on Inductive Logic Programming (ILP 2013)*, pages 1-17, Berlin, 2014. Springer-Verlag. LNAI 8812.

- Roy, S., Namboodiri, V.P.n Biswas, A., Active learning with version spaces for object detection, ArXiv e-prints, 2016

- Thompson, C. A., Califf, M. E., Mooney, R. J., Active Learning for Natural Language Parsing and Information Extraction, in Proceedings of the 16th International Conference on Machine Learning, ICML 1999, Morgan Kaufmann Publishers Inc.