EPSRC Engineering and Physical Sciences Research Council

ORCHID

Intelligent Tasking for Information Aggregation

Edwin Simpson, Steven Reece, Sarvapali Ramchurn, Stephen Roberts

Machine Learning Research Group, Department of Engineering Science, University of Oxford

DynIBCC

Combine decisions from many agents, people and sensors in a HAC

- Track changing reliability of individuals → Learning, boredom, movement...
- Model: Dynamic Bayesian Independent Classifier Combination

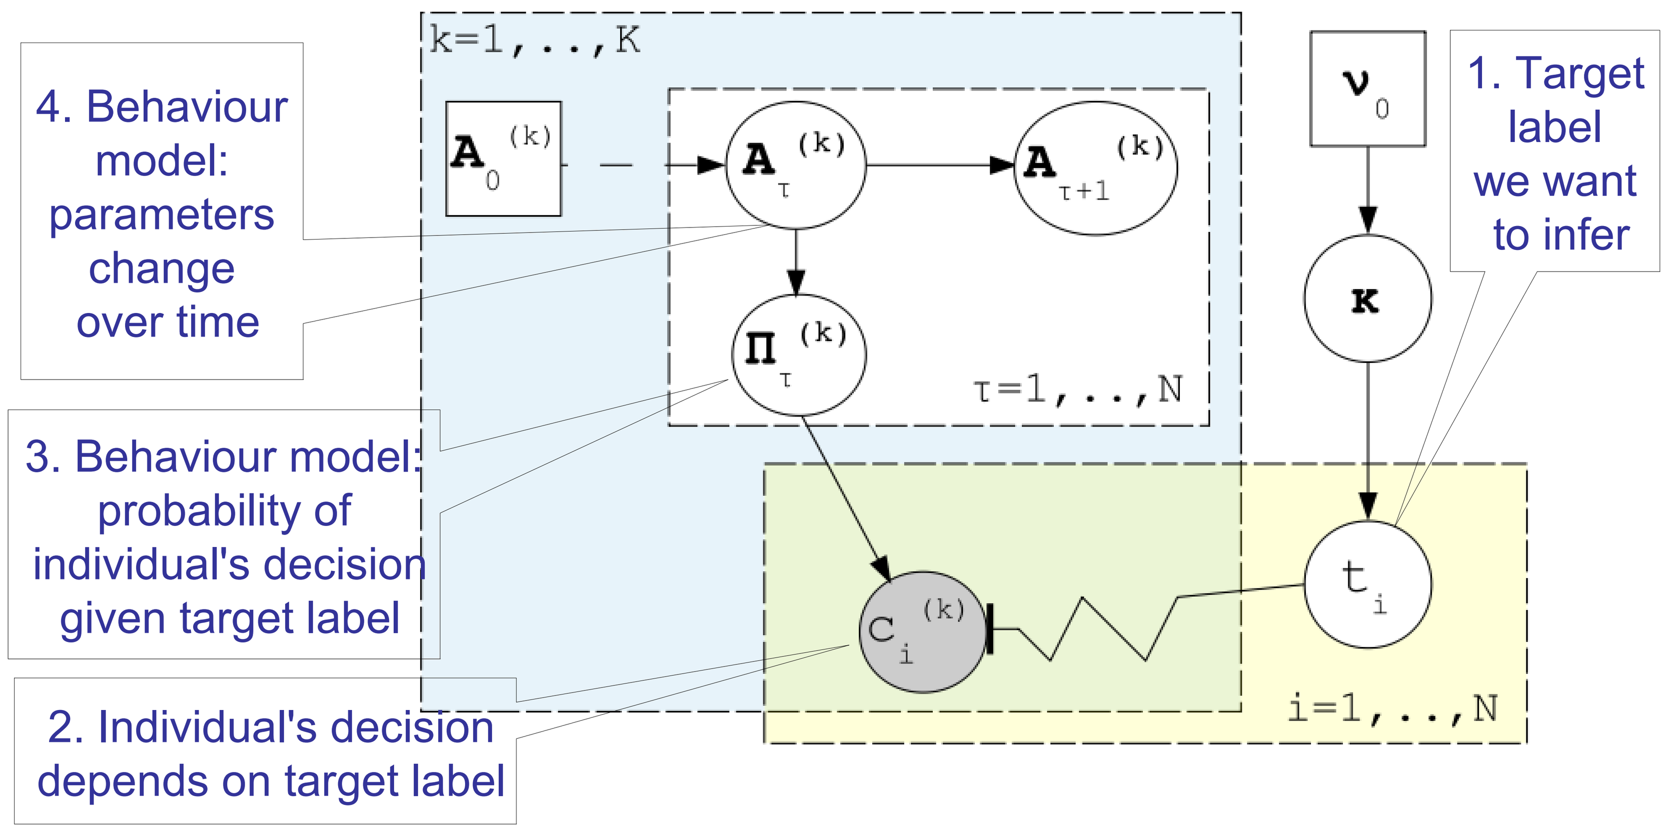

Variational Inference Algorithm

- Semi-supervised → learns distributions over all variables from latent structure in test and training data.

1. Initialise unknown variables
2. Update distribution over true labels given current model parameters
3. Update distribution over the model parameters given current target labels
4. Repeat steps 2 and 3 until converged

DynIBCC Results

Outperforms alternative methods in a range of scenarios:

- Galaxy Zoo Supernovae, GZ Mergers – citizen science
- Financial data – news sources
- HTTP Web Attacks – text features

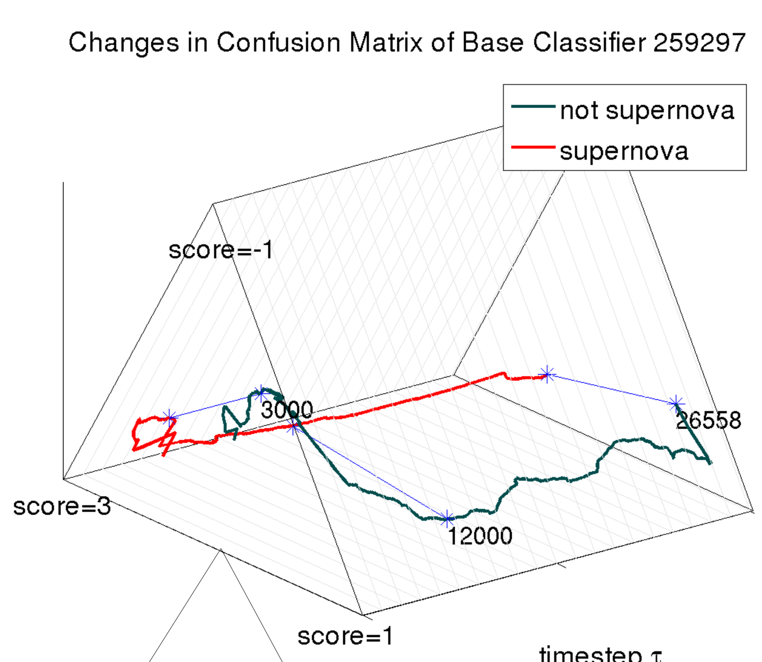

Tracks behaviour of GZ Supernovae volunteers

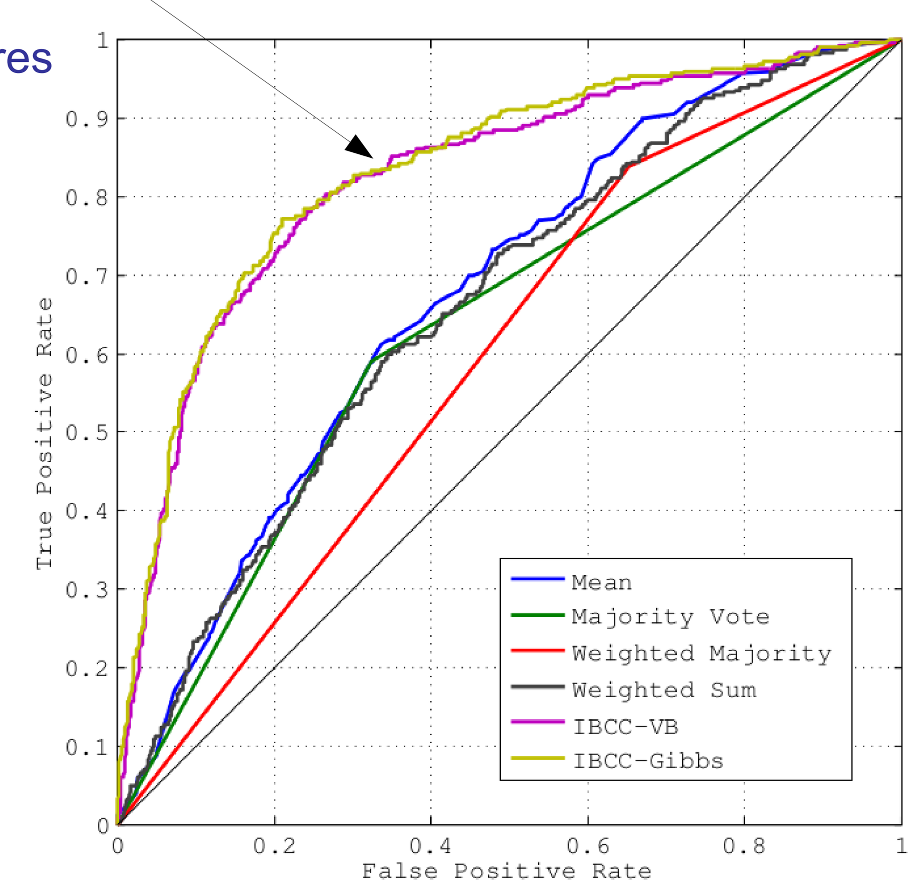

TREC Crowdsourcing Challenge

Live system with real workers using Amazon Mechanical Turk

- Document classification against complex search queries
- Screening mechanism for agents infers trust from 10 gold-labelled tasks
- Reward scheme to incentivise workers to perform more difficult tasks
- LDA text features complement responses from human workers
- Static IBCC outperformed traditional 2-stage classifier when inferring relevance from features + responses

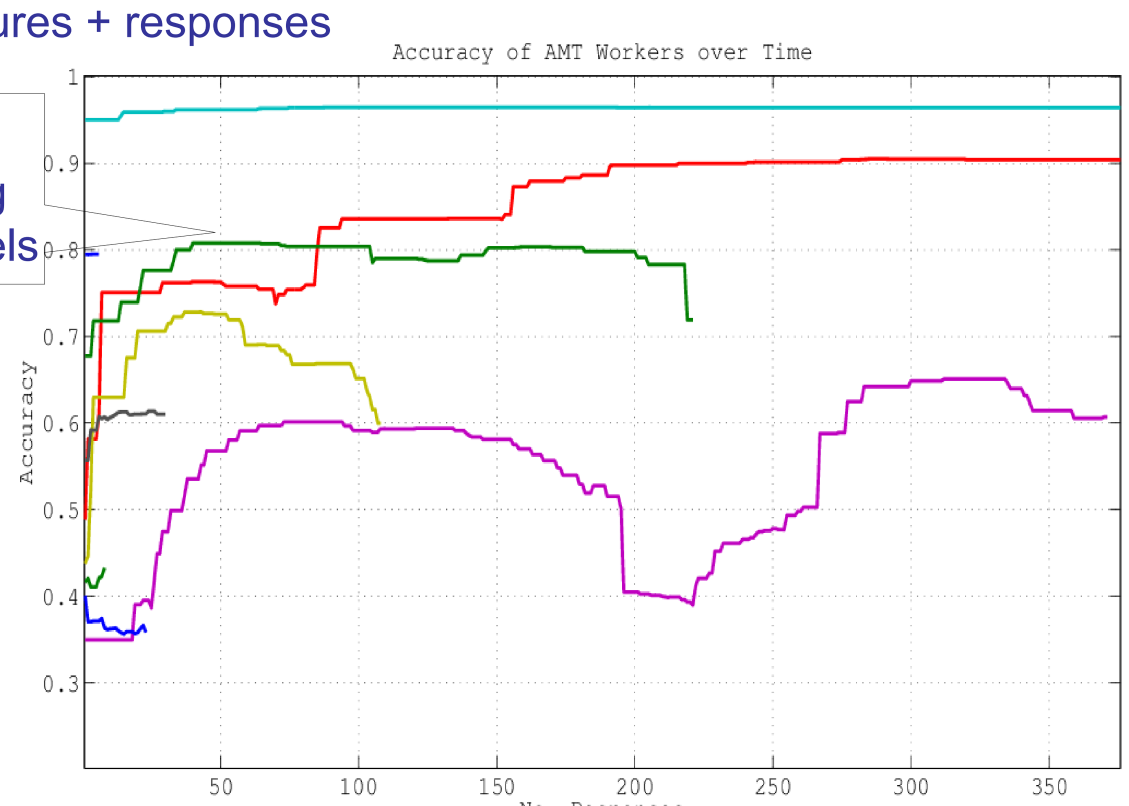

Variations in workers' abilities inferred using DynIBCC with true labels

2nd place in competition using only 2,500 labels, compared to 30,312 for 1st placed entry.

Classifier	AUC
IBCC-VB	0.806

Intelligent Tasking

Adaptively optimise the system as information is received from agents

- Deploy agents to suitable tasks using DynIBCC model
- Balance the need to learn about agents with need to learn target labels
- Train and reward agents automatically depending on benefit to system

Every possible system decision has an expected **utility** defined in terms of **information gain** over target labels. **Utility** of agent *k* completing task *i* given data at time τ:

$$U_\tau(k,i) = \mathbb{E}[I_\tau(\mathbf{t}; c_i^{(k)})] + \mathrm{Cost}(k,i) + \mathbb{E}[I_{\text{future}}(\mathbf{t}; c_i^{(k)})]$$

I_τ is immediate reward; exploits current model and data.

Cost() takes into account...

- Any financial costs (e.g. to pay an expert, rewards)
- Time penalties (e.g. for slow, complex tasks)
- Boredom/motivation cost (e.g. for repetitive tasks)

I_{future} includes future benefits from making this task assignment

- e.g. through training and experience gained by agents
- Exploring agents' behaviour
- **Silver tasking –** learning agent behaviour using unreliable labels
- Estimate I_{future} from changes in behaviour model of similar agents

Error rate on HTTP web attack data. Selecting labels using I_τ and I_{future} three steps into the future.

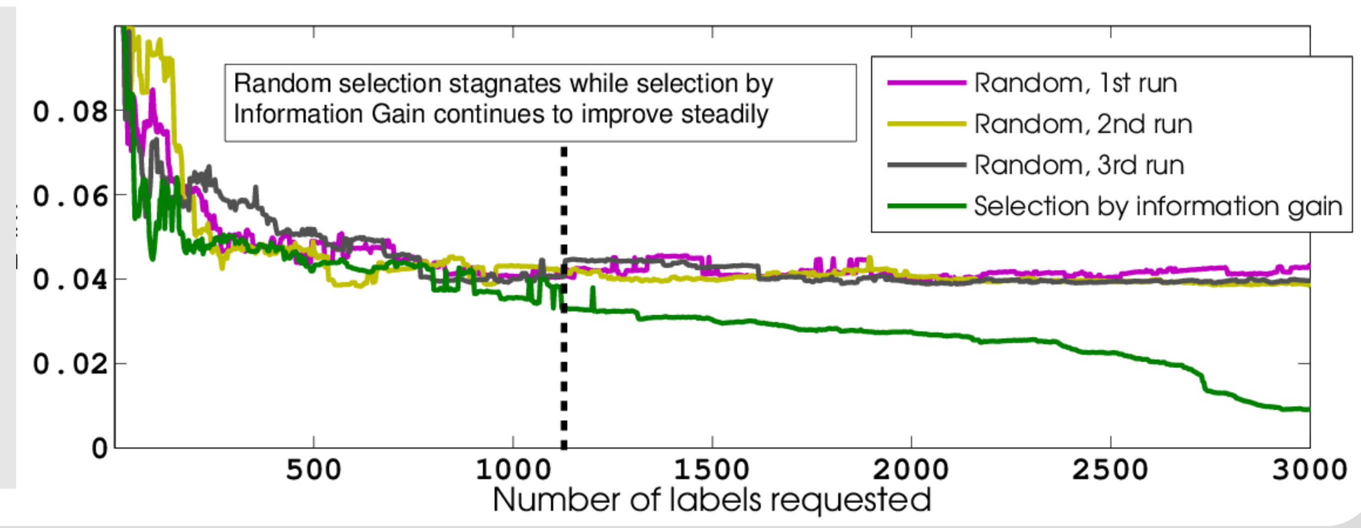

Hiring and Firing

Maintain a good workforce and assign agents to optimal tasks

- Unified, adaptive approach considering only immediate reward I_τ
- Fixed workforce size → when an agent completes a task, either **hire** for optimal task, or **fire** and replace with a new agent

Classifying TREC documents. After random no. classifications, simulated agents are bored and start guessing randomly

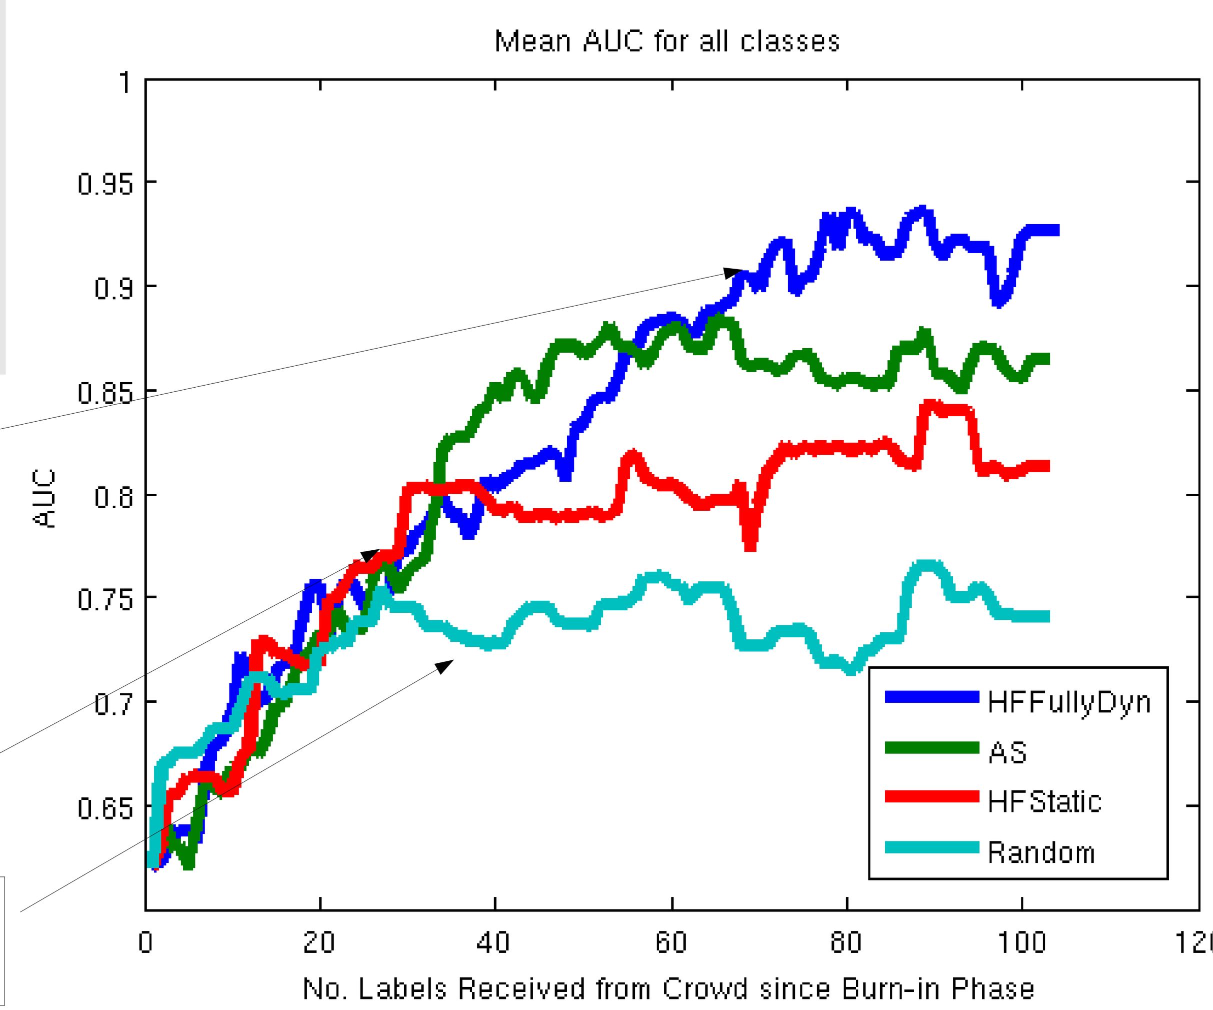

Hiring & Firing continues to improve after workers become bored

More basic variants of H&F

Random selection Method stagnates

Future Work

Agile Teaming: assign sets of tasks to ad-hoc teams; scalable approximations to expedite the search for optimal assignments e.g. using clustering, similarity graphs.

Flexible Autonomy: "weak control" allows agents to remain autonomous, but to influence their behaviour; utility function should consider cases where agents do not respond as the system intends.

Incentive Engineering: adaptively adjust rewards based on task difficulty and information value; develop utility function to model benefits of motivating tasks.

BAE SYSTEMS   ACFR AUSTRALIAN CENTRE FOR FIELD ROBOTICS   SECURE