

# Machine learning for autonomous AI systems and robotics

Chief Scientist at Dynamic Object Language Labs Inc, Dr Paul Robertson explains how autonomous AI systems are key to achieving synergy between humans and robotics that can work together as a team

**A**rtificial Intelligence (A.I.) and Autonomous Robotics are advancing rapidly and becoming robust and reliable in a number of fields, but there remains much to be done before intelligent robots can work together with us as assistants. Today, most robots are either tele-operated or perform precisely defined missions. In many cases the human overhead required to use a robot exceeds the usefulness of having it.



There is a clear need for human-robot teams in which humans and robots work side-by-side. Homogeneous teams, involve team members doing the same work. In such a case, the value comes from parallelism, such as picking grapes from vines. Heterogeneous teams consists of specialised individuals, or specialised subgroups. Our research concerns this latter kind of team. A.I. systems and autonomous robots have vastly different strengths and weaknesses from human team members. The value of the team comes from mobilising their diverse capabilities. A robot that can fly offers a unique and valuable skill. The benefits that can come from heterogeneous mixed human robot teams are immense.

## Heterogeneous autonomous teams

Heterogeneous human teams are everywhere, from the operating room in a hospital to the football field. It is instructive to look at human/animal

teams since animals are almost as different from humans as robots are.

Sheep herding (1) is essential for sheep farmers. Border collies have a natural herding instinct (2) and sheep have a flocking instinct (3). Trained dogs can understand a number of



voice and whistle commands to perform their innate skills in support of the farmer's herding needs. They can perform as a team of herders. A small team of dogs can control hundreds of sheep without losing any. The dogs know their role in herding, they under-

stand what is expected of them and the owner understands what the dogs are doing. The dogs, for their part, can predict the responses of the sheep. With this understanding, and minimal communication the dogs perform independently to achieve a successful herding mission. While simple, this example represents the core of our research.

## Model-based autonomy

Most of the structure of our world and our activities in it can be modelled. Traditionally, these models have been generated by hand, which is often appropriate, if for example we are modelling a human designed system, such as a photocopier or a road system. More recently, the models can be learned.

The herding dogs soon learn the map of the fields and the whereabouts of the gates. Knowing the map and knowing that the other dogs share the same model permits the dogs to

```
(defclass herding-behaviors []
  ...
  :methods [(defmethod separate-from-flock
              { :doc "Break one or more sheep from the flock"
                :primitive true
                ...}
              [flock set-of-sheep-to-be-separated])
            (defmethod coalesce-two-flocks
              { :doc "Form a new flock from two separate flocks"
                :primitive true
                ...}
              [flock1 flock2])
            (defmethod drive-stray-sheep-into-flock
              { :doc "Drive stray sheep back into the flock"
                :primitive true
                ...}
              [flock set-of-stray-sheep])
            (defmethod drive-flock-in-direction
              { :doc "make the flock moove in a given direction"
                :primitive true
                ...}
              [flock direction])])])
```

depend less on sound cues from the farmer. A model of the sheep allows the dogs to attend to the occasional sheep that stray from the others. Reinforcement learning allows models to be learned in simulation that implement certain operations including: “separate a group”, “coalesce two groups into one”, “move stray sheep back into a group”, and “change the direction of a flock.”

The Pamela modelling language allows these models to be represented concisely and PAT (Pamela Autonomy Toolchain) allows the models to be learned or refined. The symbolic model, such as the one shown above, provides a basis for primitive (innate) behaviours which can be used as building blocks for learning as well as for planning complex behaviours. We can provide primitive complex behaviours, learned in a simulator, which act as a starting point for learning and which can later be abandoned by the superior learned behaviours.

### Learned autonomy

Our starting point is the belief that for an A.I./Robotic system to learn to

work with humans, it will depend upon some level of pre-wired capabilities, including social skills. Our approach is a hybrid of model-based autonomy and learned autonomy. Our research includes several different learning algorithms including reinforcement learning, memory-based learning (10,11), Cluster based learning (8), and constructivist learning (9).

In nature, a mixture of innate skills and learned behaviours account for intelligent teaming behaviour. This will be essential for the next generation of robots that will work side-by-side with humans to perform complex tasks by bringing together the unique strengths of both human and robot capabilities. While much of current research focusses on learning from a blank sheet (4, 5), nature demands learning that converges rapidly and that is achieved by providing an innate starting point. It is interesting what can be learned from a blank sheet after millions of learning episodes, but nature does not have that privilege (6). Our research explores rapid learning from a good starting point.

### References

- “Mesmerising Mass Sheep Herding”  
<https://www.youtube.com/watch?v=IDQw21ntR64>
- Jyh-Ming Lien, O. B. Bayazit, R. T. Sowell, S. Rodriguez and N. M. Amato, “Shepherding behaviors,” IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004, New Orleans, LA, USA, 2004, pp. 4159-4164 Vol.4, doi: 10.1109/ROBOT.2004.1308924.
- Morihiro, K. Isokawa, T. Nishimura, H. Matsui, N. “Emergence of Flocking Behavior Based on Reinforcement Learning” in Knowledge-Based Intelligent Information and Engineering Systems, 2006. Eds: Gabrys, B. Howlett, R. J. Jain, L. C. pp 699-706.
- Baker, B. Kanitscheider, I. Markov, T. Wu, Y. Powell, G. McGrew, B. Mordatch, I. “Emergent Tool Use From Multi-Agent Autocurricula”.  
<https://arxiv.org/abs/1909.07528>
- Locke, J. [1689] 1996. An Essay Concerning Human Understanding II.i, edited by Winkler, K. P. Indianapolis: Hackett Publishing Company. pp. 33–36.
- Pinker, S. “The Blank Slate”. [Pinker.wjh.harvard.edu](http://Pinker.wjh.harvard.edu). Archived from the original on 2011-05-10. Retrieved 2011-01-19.
- Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). “Theory-based Bayesian models of inductive learning and reasoning.” *Trends in Cognitive Sciences*, 10(7), 309–318.
- Robertson, P. Georgeon, O. “Continuous Learning of Action and State Spaces (CLASS)”, PMLR 131:15-31
- Georgeon, O. Robertson, P. Xue, J. “Generating Natural Behaviors using Constructivist Algorithms.” PMLR 131:5-14
- Robertson, P. “Memory-based Simultaneous Learning of Motor, Perceptual, and Navigation Skills” Workshop on Machine Learning IROS 2008, Nice.
- Robertson, P. Laddaga, R. “A Biologically Inspired Spatial Computer that Learns to See and Act.” Spatial Computing Workshop, SASO 2009, San Francisco.



**Dr Paul Robertson**  
**Chief Scientist**

Dynamic Object Language Labs Inc. (DOLL Inc)  
 Tel: +1 617 415 9327  
[paulr@dollabs.com](mailto:paulr@dollabs.com)  
<https://www.dollabs.com>