Australian Government

Department of Defence

Defence Science and Technology Group

Machine Cognition
Open Questions

Dr Ian Dall

Tyche Theme Leader – Machine Cognition

Gap

- Acknowledge the gap between the Autonomy we desire and the automation we have.
- What, from a machine cognition point of view is missing. Why can't we build them now?
- Do we need
  - New logic?
  - New architectures?
  - Better Knowledge representation?
  - Better more complete ontologies?
  - Deeper belief networks?
  - Bigger training sets?
  - Faster machines?
  - New algorithms?
  - New paradigms?

Current Status

- AI Spring and AI Winter
- First wave systems were not so much failures as “limited”: the scale-up problem.
- Cursed by definition.
- $8.5b in 2015 by Google, Microsoft, Amazon etc. Up four fold since 2010.
- We seem to be in a new AI Spring.
- Perhaps winter is coming.

Successes

- Considerable success of deep learning
  - Image tagging
  - Voice recognition
- Decent heuristics for NP-Complete problems
- Beat grandmasters in Chess and Go
- Watson won Jeopardy
- Siri is somewhat useable
- Route planning (eg Tom Tom, Google maps) is routine.
- Machine translation is somewhat useable.

Limitations

- Learning from limited examples.
- Generalising.
- Creative solutions.
- Analogical thinking.
- Transferring learning to new domains.
- Dealing with the unexpected.
- Natural language understanding.

Biological Inspiration

- If we want human level AI, why not look to human intelligence for inspiration?
- Need to be selective:
  - No flapping wing aircraft or walking cars.
  - Human level intelligence might come with human biases, neuroses etc. And motivation!
- How do we determine which aspects to mimic and which to ignore?

Human Visual Perception

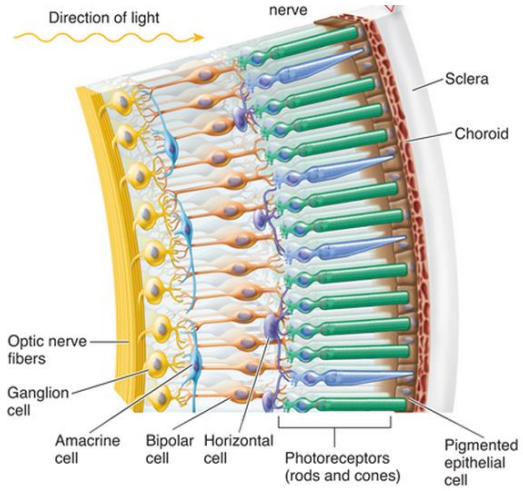

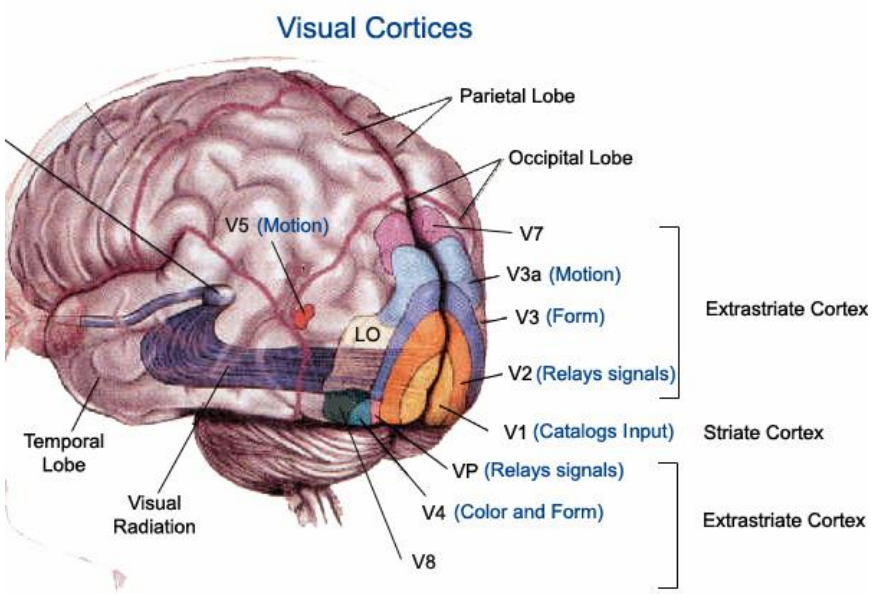

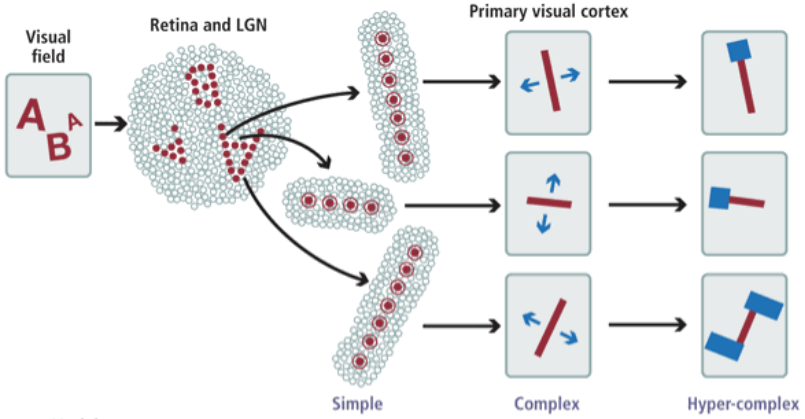

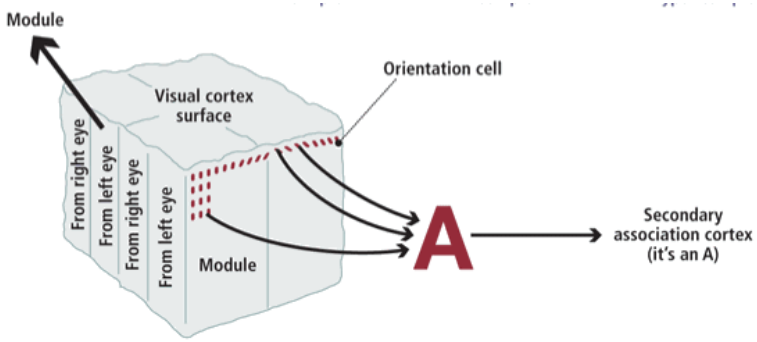

Vertical and Horizontal Specialisation

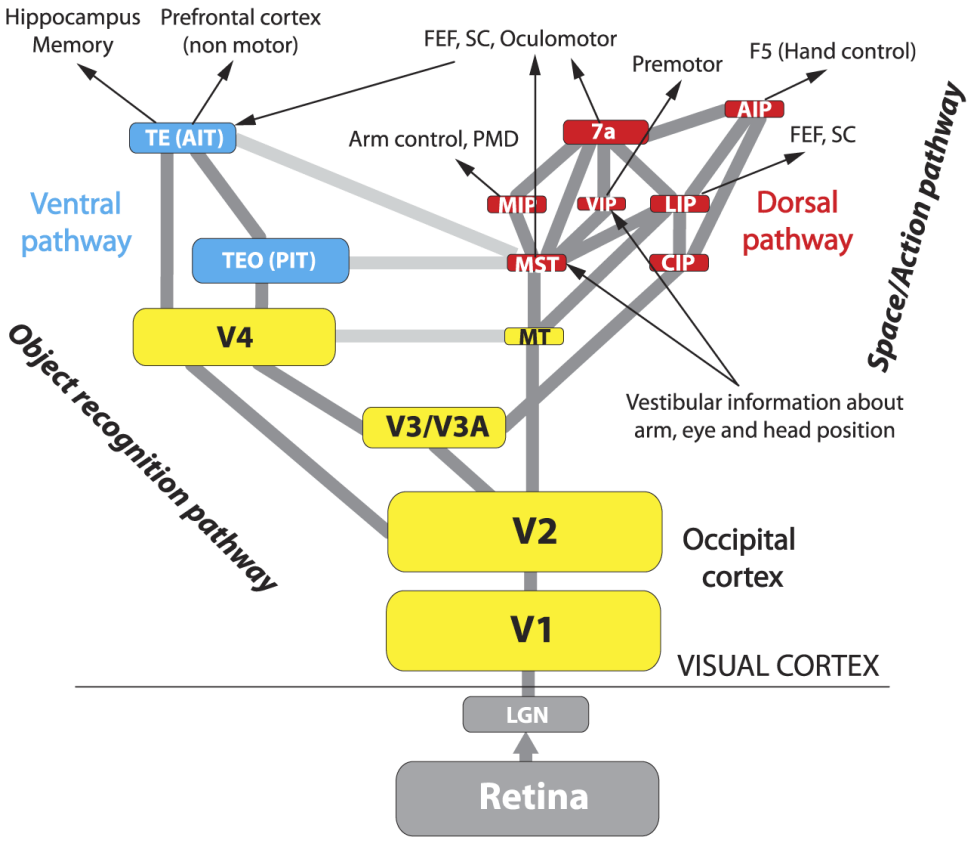

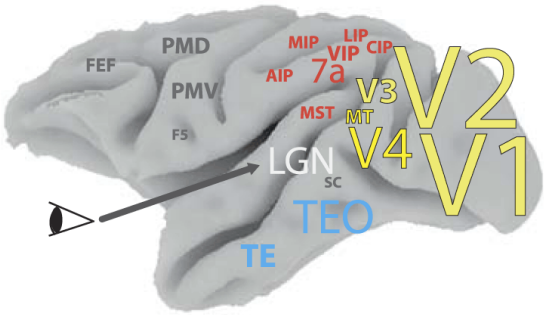

Motivations for Machine Learning

- Adaption to changing environments.
- Modelling complex difficult to define functions, eg mapping arrays of pixels to symbols.
  - Slow.
  - Machine learning inevitably results in specialisation to the characteristics of the training set.

Grow intelligence rather than design it.

- Perhaps if we can't design AI we can grow it.
- Mimic the process by which humans evolved rather than the result.

Fast and Slow Learning

Biologically:

- Fast learning – don't touch the hot plate.
- Slow learning – 10k hrs to become an expert.
- Very slow learning – evolution.
- Peer learning and Cultural learning (education).

Both slow and fast learning have their place.

Machine:

- Fast Learning eg Inductive Logic learning, Hierarchical Bayesian Networks, AGI.
- Slow Learning eg Deep Belief Networks.
- Very slow: evolution of structures and algorithms?

Hybrid Approaches

- Use slow learning for stable aspects of the environment:
  - Laws of physics
  - Human nature
  - Common (reusable) cognitive routines
  - Reusable primitives cf retinal ganglion.
- Use fast learning for more ephemeral knowledge

Pragmatic Approaches

- Use crafted algorithms when they work
  - Search, classical planning, minimax games, SLAM etc.
- Use neural networks when they work.
- Use logic, logic programming, induction, deduction, abduction when they work.
- Incorporate engineered knowledge when it works.

Deciding when to use is metacognition.

Requires a suitable architecture – horizontal specialisation cf human perception.

Questions to think about

In the context of achieving autonomy:

- What are the outstanding challenges in machine cognition.
- How can we build on the substantial achievements in AI so far.