Machine Learning

Past, Present, and Future

Luke Heffernan, Australian Institute for Machine Learning (AIML)

The IoT Connection, Hyperconverged Hybrid Networks: Space. Stratosphere. Terrestrial.

About Me

- Managed AI/ML projects in:
  - Space, defence, art, agriculture, traffic optimisation, communications, more
- Founded and assisted companies implementing AI
- Researching alongside world leaders in research and development
- AIML globally 6th best in computer vision

What To Expect

- Brief history of AI
  - From ideas to reality
  - Concepts
- General does and don’ts
  - With examples
- Brainstormed use cases
- Blend of technical and analogy
  - Further reading
  - Some parts will be skimmed
  - Slides and links available on Linktree:

Contact details

Slides (pdf)

History

- Ancient ideas of automaton
- Beginning of computation
  - Expert systems, basic neural networks
- Deep learning
  - Computational scale, latest AI boom

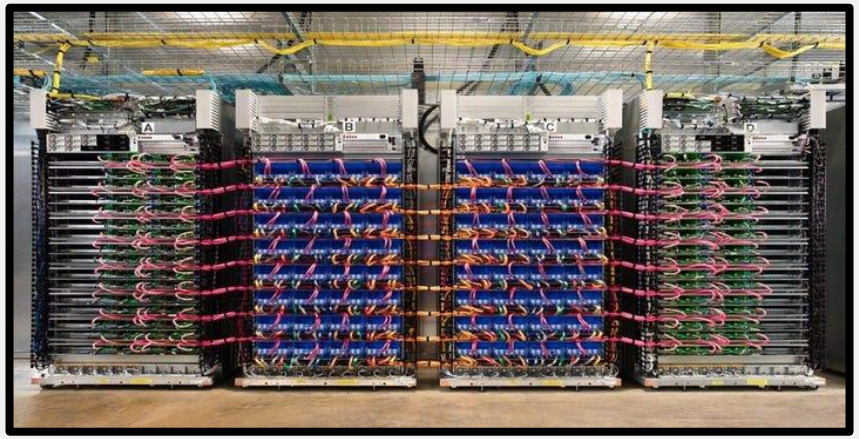

Enablers

- Better processors, smarter processing
  - CPUs vs GPU, TPUs
- More effective, efficient, tailored algorithms
- Investment, education, and expertise

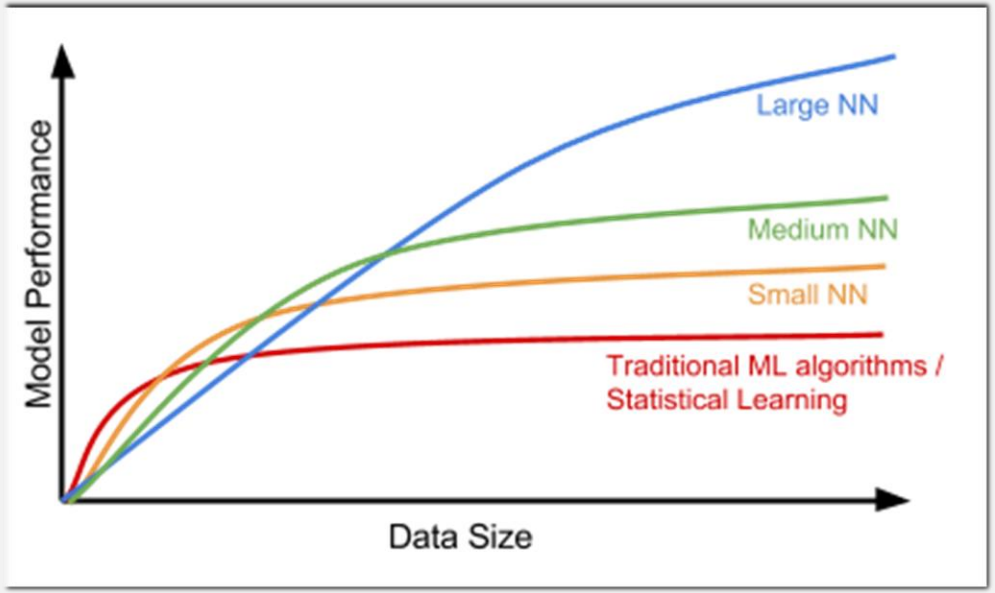

'Recent' Big News

- Transformers (2017, ancient)
  - Massively increased size -> memory
- Large language models
  - GPT-2/-3/-4
  - ChatGPT

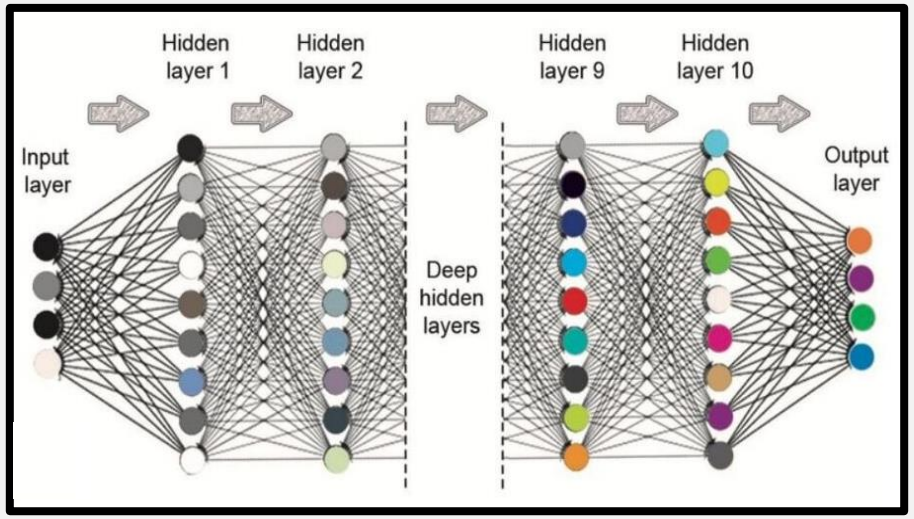

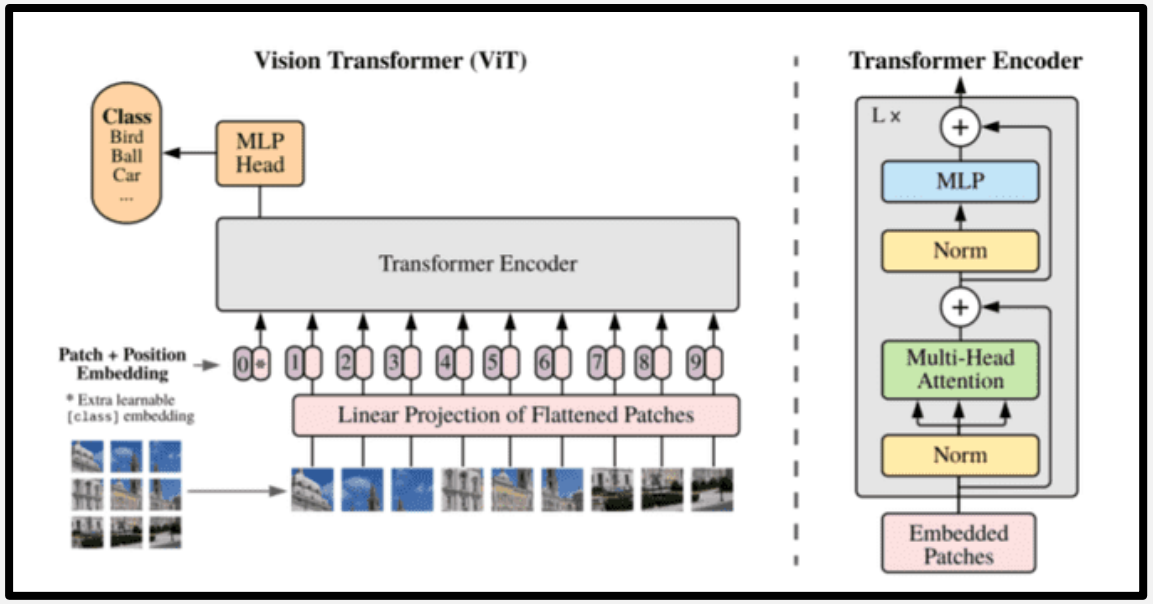

Concepts

- Some understand analogies to pattern recognition or compression
- Different algorithms for:
  - Text
  - Images
  - Audio
  - Long vs short term memory
  - Training speed
  - Data volume

Concepts

- Learning Types
  - **Supervised learning**
    - **Fully**
    - Semi-
    - Un-
    - Weakly
    - Self-
  - **Reinforcement learning**
  - Generative
    - Adversarial
    - Diffusion

- Data Types
  - Audio:
    - Recurrent NNs (RNN), Long Short Term Memory Units (LSTMs), Gated Recurrent Units (GRUS), Transformers
  - Text:
    - RNNs, LSTMs, GRUs, Transformers
  - Images:
    - Convolutional NNs (CNNS), Visual Geometry Group (VGG), Visual Transformers (ViT)
  - Games/Interactions:
    - Reinforcement systems (Q learning, actor-critic, model-based, etc)

Rules/Suggestions of Thumb

All general rules have exceptions, often many

- Input = output
- Bigger = better
- Expertise needed
  - Machine learning
  - Subject matter
  - Soft skills

Input = Output

- Don’t ask a vet for personal knee surgery
- Trends toward generality, but not reached

Bigger = better

- True for layers/data volumes/compute, with the right model
- False, with irrelevant data or poor model architecture

Expertise Needed

- Machine learning
  - Easy to trial, hard to perfect
  - Scale matters and multiplies skill*
- Subject matter
  - Measures of success
  - Input = output
- Soft skills

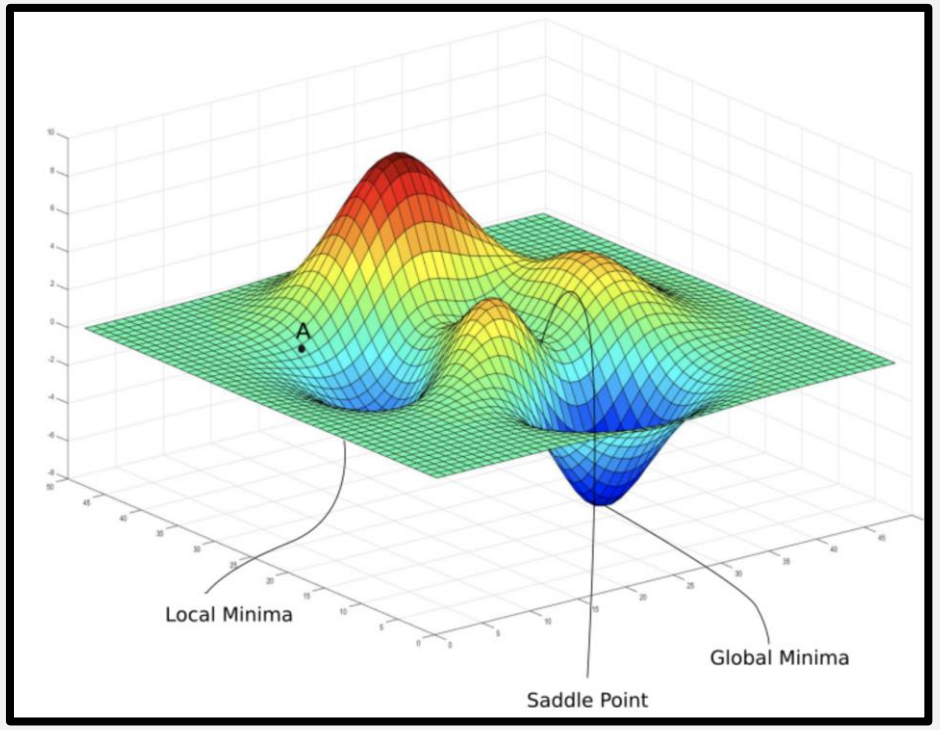

**Concepts to familiarise with:* encoding, data types and formats, activation functions, catastrophic forgetting, batch sizes, learning rates, training types, real-world transferability, prediction types, layer number/size, dropout, feature, training augmentation, tests and assessments, data splits, model type and architecture, data pipelines, optimisers, iterations/epochs, entropy, memory allocation, warm-up, fine-tuning, convergence, synthetic vs real data, collection and verification, data processing, hallucination recognition, learning types, reward functions, shortcut mitigation, psychology and interpretation, bias mitigation, filtering…

Common Mistakes/Misconceptions

- Overconfidence
- A good model can fix bad data
  - Hidden patterns
  - Mess and noise
- *Do you really need it?*

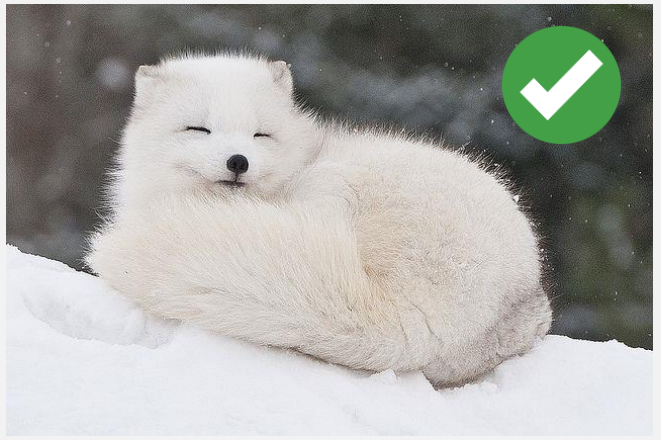

Study #1: Burning Money 101

- Knight Capital auto-trading
  - KC traded 11% of all US stocks in early 2012
  - Lost $440mil USD in 44 minutes ($10mil/minute)
- If they can do it, anyone can

Study #1: Burning Money 101

- Implementation at KC
  - As simple as copying and pasting wrong
  - Errors in implementation and confidence

- Popularity in Financial Trading
  - Dangers of a little knowledge
  - Errors in expertise

Study #2: Meta's Missing You

- Social networks
- Recommender systems in games, feeds, events, markets, groups, and more

Study #2: Meta's Missing You

- Comprehensive approach
- Advertisers
  - Buyers and suppliers
- Proximity in network effects
  - The best data is your friends' data

Study #3: Emergency Response

- Natural disasters are universal and increasing
  - Omdena's 2020 project to mitigate human impact due to earthquake in Istanbul, Türkiye

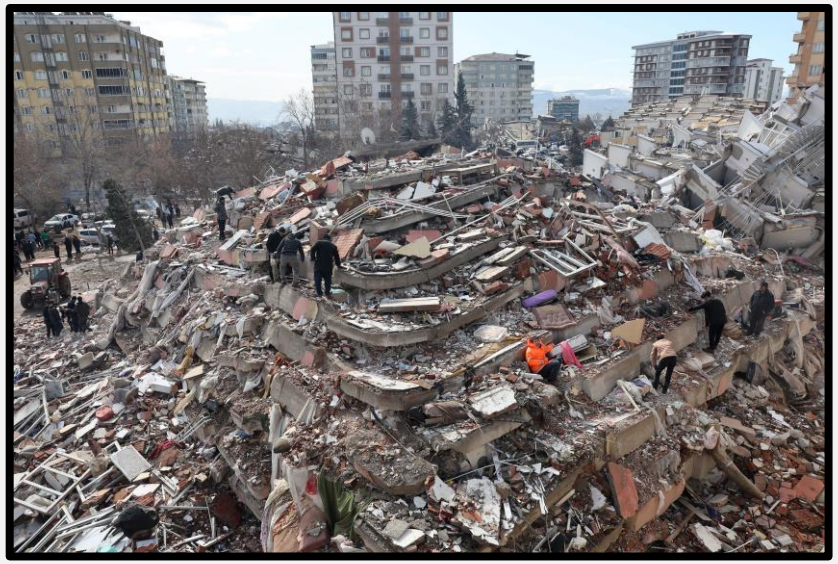

Study #3: Emergency Response

- Omenda had AI expertise and relevant SMEs
  - SMEs advised reuniting families is priority #1
  - Created risk heatmap and path generator
- Well-defined problem
  - Intuitive interface allows easy use
  - Explored the root causes and related effects
- Similar approaches in fires and floods
  - AIML allowed firespread prediction and resource allocation for first responders

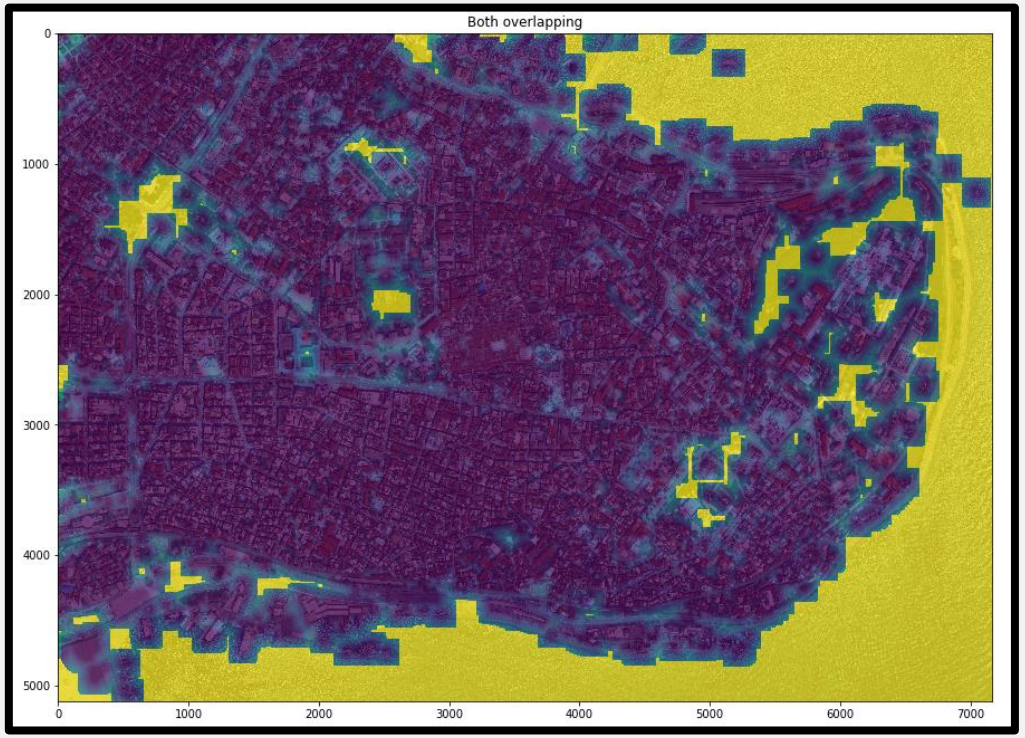

Potential use cases

- Space Vehicle Conjunction
  - Pose estimation enables satellite docking, refuelling, resupply, in-space SDA
  - Northrup Grumman/SpaceLogistics flew the Mission Extension Vehicle 2 (MEV)
  - NASA, ESA, OrbitFab, Space Machines Co., HEO Robotics

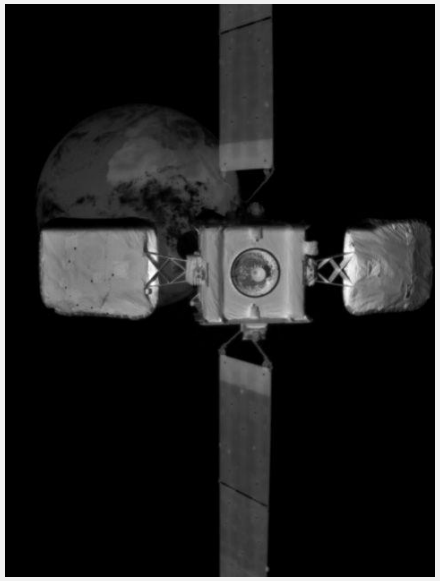

An image of Intelsat 10-02 taken by MEV-2's infrared wide field of view camera at 15m away.

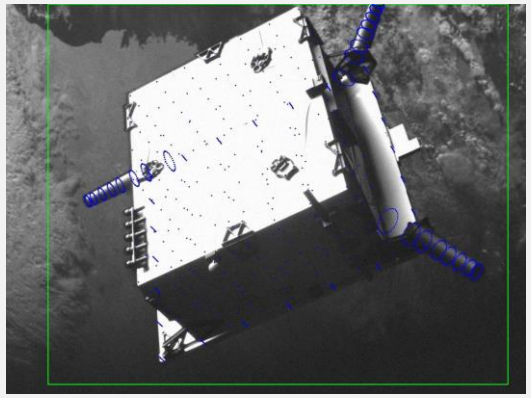

AIML's 1st place entry into ESA's Pose Estimation Challenge 2021

Potential use cases

- Life Monitoring and Support
  - Excels at balancing multi-sensor inputs
  - Many cases of detecting issues before they're issues
  - Optimise not just detection, but response and resources
- Non-destructive testing
  - Defect detection in manufactured items
- Low light imaging
  - X-ray results with 99% less radiation
- Mental health warning systems
  - Changes in message content and tone
- Consumer trend monitoring
  - Purchasing patterns
  - Right ways and wrong ways, e.g. Target

Potential use cases

- Life Monitoring and Support
  - Excels at balancing multi-sensor inputs
  - Many cases of detecting issues before they're issues
  - Optimise not just detection, but response and resources

- Non-destructive testing
  - Defect detection in manufactured items
- Low light imaging
  - X-ray results with 99% less radiation
- Mental health warning systems
  - Changes in message content and tone
- Consumer trend monitoring
  - Purchasing patterns
  - Right ways and wrong ways, e.g. Target

Future Directions (Disclaimer)

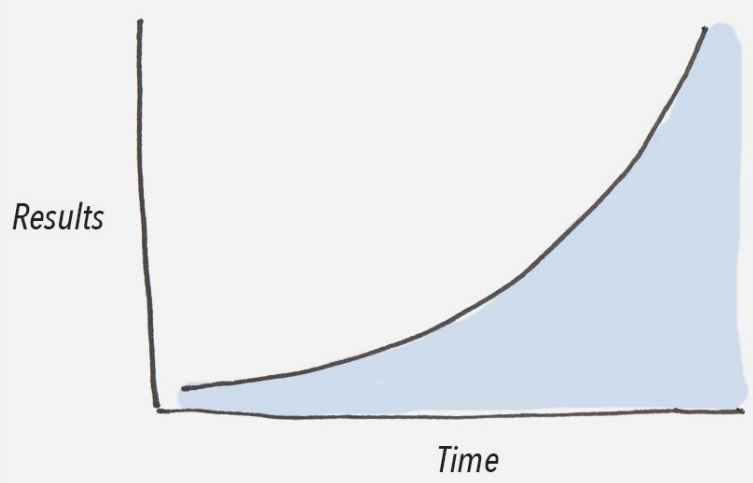

- Every time we guess, we're wrong
- Experts can't easily account for acceleration

- Hard to predict future of prediction

Future Directions

- Current focuses:
  - Reinforcement learning to interact with humans
  - Reliable integration of LLMs with other tools/programs
  - ML in a post-LLM world
  - Vision-language crossovers

- Suggestions:
  - Change will be fast, this is disruptive
  - For every interest or anxiety, endless answers, very few solutions
  - Adopt, be sceptical, engage experts

Future Directions

- Current focuses:
  - Reinforcement learning to interact with humans
  - Reliable integration of LLMs with other tools/programs
  - ML in a post-LLM world
  - Vision-language crossovers

- Suggestions:
  - Change will be fast, this is disruptive
  - For every interest or anxiety, endless answers, very few solutions
  - Adopt, be sceptical, engage experts

+ session slides

Luke Heffernan

luke.heffernan@adelaide.edu.au