Advanced Machine Learning

Andrew H. Fagg

General Takeaways

- Understand the nature of your data and the prediction problem you are trying to solve **before** implementation
- Understand the distribution of your data
  - Are there biases?
  - Stratification is often necessary
- Be clear up front how you will use your data
  - Training/validation/testing sets all have specific function. Don’t confuse them
  - Multiple, independent data sets are ideal
  - Cross-validation may be necessary if data are at a premium

General Takeaways

Quick implementations are a good way to explore an idea quickly:

- Strive to write and reuse generalizable code
- Generalizable code makes it easy to build large architectures
- Command-line interfaces make it really easy to try many different variations of models / architectures

Big Steps in the Last 2 Years+

- Text prediction (e.g., GPTo)
  - Automated writing or writing assistance
  - Automated coding or coding assistance
- Connecting text to images
  - Diffusion methods: Photorealistic image generation and in-painting
- Physics simulations
  - Alpha-fold

Challenges: Data

Many parameters -> expressive models. But:

- Need a lot of data to support the selection of parameters
- Examples need to be labeled

Also:

- Training set sample can be heavily biased, which can dramatically affect the quality of the models
- Big models / data require a lot of computation and energy

Challenges: Data

Reducing the reliance on training data with quality labels

- Unsupervised approaches for dimensionality reduction
- Semi-supervised learning approaches
  - Mix of examples that are labeled / unlabeled
  - Mix of examples in which only a subset of labels are vetted

Moving Beyond Labels: Reinforcement Learning

- Agents making sequences of action choices in some environment
- Feedback is evaluative: how good is each action or outcome?
- Games, control systems, robots …
- Agent generates its own experience for training by trying many action choices in many different situations
- Also very data/experience hungry

Doing More with Less

- Smaller models -> less training and less energy use
- Efficient use of available data

Analog Implementations

- Now seeing a lot of silicon-based implementations of NN architectures
  - Speed and power efficiency
- Learning in silicon seems to be a long way off

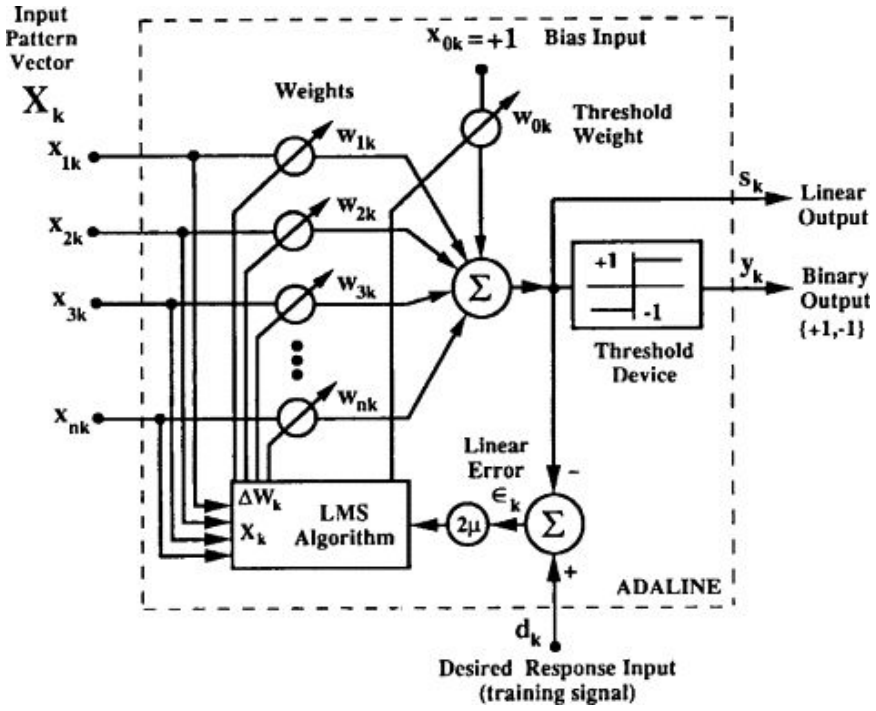

Widrow and Hoff (1959)

General AI Problem

- While Deep Learning has produced amazing results, it is not a solution to the General AI problem
- But, DL will make up critical components of larger AI architectures

General AI Problem

Need to solve a lot problems including:

- …

General AI Problem

Need to solve a lot problems including:

- Creating and managing goals
- Connecting the continuous world (images & other sensor data) to symbols / objects
  - For communication and for generalization across many situations
- Modeling the physics of the world
- Short term and long term memory
- Efficient, long-term planning
- Modeling of other agents (“theory of mind”)
- Being able to explain reasoning processes

Implications of High-Performance AI Systems

Implications of High-Performance AI Systems

- Enhance our ability to model complex, dynamical systems: improved prediction and decision making
  - Medicine, weather/climate, energy, manufacturing, …
- Moore’s Law -> widespread availability of AI products
- Potential for democratized access to the big benefits of AI

Existential Threats

Existential Threats

- Centralized control of AI systems
  - Computation, outcomes and data
- Dramatic changes in the human component of work
- Automating large-scale decision making processes
- Optimizing decision making processes that do not take into account individual/community/social/human priorities

How Does our Educational System Change?