31YB: Final Lecture 15: Thurs, 2 Dec
Neuromorphic Systems

Overview

- Current Trends - Way forward in ANN research
- Intro. to Neuromorphic Systems: History
- Real & Neuromorphic Neurons
- Silicon implementation techniques
- Applications of Neuromorphic Systems
- Timeliness
- Way Forward in Neuromorphic Systems Research
    - For your interest! (NOT Examinable part!)

Current Trends - Way Forward in Neural Networks Research

•More recently, over the last few years, neural networks (NN) have started to be seen as part of the larger field of soft computing or natural computing, and there has been real movement of the NN discipline in three directions:

•**Neural Networks, statistics, generative models & Bayesian inference**

–There is a sense in which these fields are coalescing. The real problem is making conclusions from incomplete, noisy data, and all of these fields offer something in this area. Developments in the mathematics underlying these fields have shown that there are real similarities in the techniques used.

•**Biological Neural Networks**

–There is real interest in how neural network research and neurophysiology can come together. The pattern recognition aspects of Artificial Neural Networks don't really explain too much about how real brains actually work. The field called Computational Neuroscience has taken inspiration from both artificial neural networks & neurophysiology& tries to put two together.

•**Neuromorphic Systems**

–Existing neural network (and indeed other soft computing) systems are generally software models for solving static problems on PCs. But why not free the concept from the workstation? The area of neuromorphic systems is concerned with real-time implementations of neurally inspired systems, generally implemented directly in silicon, for sensory and motor tasks. Another aspect is direct implementation of detailed aspects of neurons (Biological Neural Networks) in silicon.

–MORE IN THIS LECTURE !! - NEXT

Intro. to Neuromorphic Systems: History

•Research into neuromorphic systems part of larger field of computational neuroscience

•Term ‘neuromorphic’ started off meaning ‘neuron-like’ (late80’s)

•another definition: ‘mimic(ing) specific neurobiological functions’

•Neuromorphic systems are implementations in silicon of systems whose architecture and design are based on neurobiology.

•Earliest neuromorphic systems were concerned with providing an engineering approximation to some aspects of sensory systems, such as, the detection of sound in the auditory system, or detection of light in the visual system

•More recently, work on robot control systems, on modeling various types of neurons, and on including adaptation in hardware systems.

•Also, research not restricted to one specific implementation technology

Real & Neuromorphic Neurons

• Real neurons - remarkably complex entities

•May be modelled in many way - emphasizing learning, or neural geometry, or ionic channels; - and using differential equations or simple discrete time

•Considerable success has been shown, with very simple neuron models, such as threshold units (McCulloch-Pitts neuron model) or logistic units (Perceptrons), for performing interesting computation - However, these units are unable to directly model any of the behaviour of the neuron over time - hence, too simple for use as models for systems which need to process time varying environmental data, or for controlling movement over time

•Recently, considerable interest in more complex neural models, such as integrate & fire neurons or spike response neurons, which permit modelling of the processing of time-varying data

– [NOT Examinable] e.g. in such networks, the neural elements can act as oscillators, and communicate using spikes. These offer the possibility of more sophisticated forms of representation of data than simple McCulloch-Pitts neurons because they can take advantage of the temporal domain. This could be through synchronized oscillation, or through coincidence or near-coincidence of spikes. Recent work on antidromic spikes suggests mechanisms whereby Hebbian rules are implemented

Real & Neuromorphic Neurons

• However, real neurons much more complex - at most basic level, they have many different charge carriers (such as Na^+ , Ca^+ , K^+ ions), and many varieties of channels whose precise properties depend on both local depolarisation & on presence or absence of other neuromodulator chemicals.

•Further, real neurons have spatially extended processes for both input (dendrites) & output (axons).

•*Question*: How can one hope to produce any useful model based on ‘real’ neurons using electronic circuitry (whose only charge carrier is the electron & the only active element is the transistor)?

Real & Neuromorphic Neurons

• *Answer 1*: Each technology must make use of whatever active processes it can support, there are a plethora of these in biological systems but only a few in electronics. However, in electronics, one can achieve complex behaviour by combining no. of elements into circuits, which can run in real-time.

•*Answer 2*: One should consider modelling *what* the neural system is doing rather than *how* it is achieving its results. That is, one can model how the overall architecture achieves its results (thus retaining e.g. massive parallelism typical of neurobiological systems) without modeling ionic channels, or even synapses. Indeed even the neurons might disappear, replaced by functions approximating their apparent behaviour.

•*Jury is still out on exactly what requires to be simulated in order to capture important aspects of real neural computation - mimicing neurobiological function doesnot need to entail precise emulation of original neural system.*

Silicon implementation technologies

• CMOS (Complementary Metal-Oxide Semiconductor) is the dominant electronic technology for implementation of neuromorphic systems.

•This technology allows integration of large numbers of transistors (also known as Very Large Scale Integration, VLSI) onto a silicon device (single chip).

•CMOS VLSI may be operated in either analogue or digital modes, or combination of both

•Most neuromorphic research uses analogue VLSI techniques (e.g. subthreshold analogue CMOS) – to achieve greater parallelism, very low power dissipation and by careful use of analogue characteristics of transistors, arithmetic functions such as addition, multiplication, exponential, logarithmic & *tanh* functions may be implemented using relatively few transistors

  –This is a great advantage over digital VLSI implementations of these functions in which very large circuit areas & relatively high power dissipation may be required

Timeliness?

• NS not entire now, so why recent increased interest?

•*Answer*: There has been convergence of work in neuroscience, neural networks and chip design.

•In Neuroscience, we now better understand both operation of neuron (at spike/axon/dendrite & synapse/neuro-transmitter levels) as well as importance of particular specialized structures (e.g. retina or organ of corti) at sensory neural surfaces.

•Research on neural networks has shown that one can use networks of relatively simple neural models to implement complex data transforms - also one can use VLSI techniques to implement some of these networks, but this has also clarified limitations all NNs (artificial or biological) suffer from, in terms of interconnections of large no. of active elements.

•Chip design is now a mature subject, and some of concepts (such as design rules & re-use of cells) that made digital VLSI straightforward can now be applied to analogue VLSI.

•*Convergence in above three areas* has allowed engineers to produce circuits to embody some of the design strategy which neurobiology has evolved

Where might NS be applied

- Neuromorphic systems (NS) offer possibility of small low-power devices which can process real data directly: that is, data which has been sensed directly, perhaps using a transducers (also part of the chip)
- Since, no human intervention between energy source & transducer, these systems can be applied to sensing for independent robots.
- In addition, they can be applied to prosthetics -e.g in the design of hearing aids
- Also, in motor control, permitting control of robot manipulators
  - Bio-robotics : The intersection between biology and robotics is a growing area in neuromorphic systems. Biorobotics aims to investigate biological sensorimotor control systems by building robot models of them e.g. development of novel sensors and actuators, hardware and software emulation of neural control systems etc. At Stirling and elsewhere, a variety of animal systems are currently being investigated using this methodology.
- For both above applications, real-time & low-power characteristics of NS are crucial
- Current PCs are static, their inputs mediated by humans via KB etc
- NS represent part of vehicle for moving PCs off the desk & out into the environment, for processing real 'dirty' data on-line

Summary – Neuromorphic Systems

- Neuromorphic systems are implementations in silicon of systems whose architecture and design are based on neurobiology.
- This growing area proffers exciting possibilities such as sensory systems which can compete with human senses, pattern recognition systems that can run in real-time and neuron models that can truly emulate living neurons.
- Neuromorphic systems are at the intersection of neuroscience, computer science and electrical engineering
- The main centres worldwide are at the <u>Institute for neuroinformatics</u> at Zurich, and at <u>Center for Neuromorphic Systems Engineering</u> at Caltech, together with emerging centres in the UK (e.g. Edinburgh and Stirling)

END OF COURSE!

- *REMAINING MATERIAL IN THIS LECTURE IS FOR YOUR INTEREST* (NOT Examinable!): In the next few slides, we discuss in more detail, an application area for neuromorphic systems, which is subject of considerable current interest – synthetic sensory systems

Synthetic Sensory Systems – Way forward for Neuromorphic Systems Research

Aims: To develop systems which provide senses for (e.g.) computers, or autonomous robotic systems.

What is a *sensory system*?

Sensory systems are what animals (and even plants) use to find out about their environment. A *sense* is the capability to interpret some variable flux: from a biological viewpoint, the flux being interpreted must impact on the sensing entities survival (whether through finding food, or a mate, or avoiding prey, or whatever). A flux is simply something which varies (or fluctuates).

Examples of biological systems and their fluxes are given below:

Sense	**Flux**
Sight	Electromagnetic waves in the visible spectrum
Hearing	Pressure waves between about 30 and 20000Hz
Smell	Concentration of various airborne molecules
Taste	Concentration of various ions etc. in saliva
Touch	Pattern of pressure at points on animal's surface
Proprioception	Pattern of nerve impulses from muscles

With the exception of proprioception, these senses all tell the animal something about the external world: proprioception tells the animal about relative position of parts of its body. Note that some, but not all, of the fluxes are patterns of incident energy.

Synthetic Sensory Systems – Way forward for Neuromorphic Systems Research

What is a *synthetic sensory system*?

•Detects some (probably external) flux, translating it into internal signals, and interpreting these signals in some way.

•The interpretation should have some relevance to the reason for the machine's existence, although this really only applies if the system which the sensing is part of is autonomous.

•One can consider everyday computers to have a limited set of senses: they can interpret key depressions and mouse movements.

•One aim of the work implied here could be to increase this, possibly by providing hearing and sight for computer systems.

•This could certainly increase the range of inputs available to the machines, and conceivably make them easier to use.

•Clearly, more autonomous machines have more need of senses, if only to permit their (fragile) systems to survive in a hostile environment.

Synthetic Sensory Systems – Way forward for Neuromorphic Systems Research

Why do research into synthetic sensory systems

- understand the biological system better: by building a working model of a biological sensing systems we gain insight into the biological system.
- the other primary motivations are the possibility of producing better sensing systems for robots or sessile computer systems
- developing prosthesis for humans who have impaired sensing faculties (e.g. hearing aids)

- The possible gains from any of these research motivations are huge: by understanding animal sensing we will both improve our understanding of the way in which animals fit their ecological niches and improve the design of artificial prostheses for impaired humans.
- Better sensing systems for completely synthetic systems (like mobile robots or desktop computers) will allow them to interact much more effectively with their environment.
- For an autonomous robot, this is the difference between being able to work usefully independently and not being able to do so.

Synthetic Sensory Systems – Way forward for Neuromorphic Systems Research

Why do research into synthetic sensory systems

- At first sight, one might imagine that adding senses to a desktop computer would not be useful?
- however, if one compares the sophistication of the display with that of the input devices (keyboard and mouse) one rapidly realises that the input devices lag way behind.
- In the last 10 years, screens have improved enormously in quality, yet the last improvement in input devices was the mouse.
- Keyboards have not altered materially in many years.
- Input based on sensing includes sound input (and that, in turn, includes speech), visual input, and even intelligent usage of keyboard and mouse input (e.g. are key-depressions frequent or infrequent, often incorrect or always right, is mouse usage smooth or jerky, etc.).
- Indeed, one can imagine the desktop machine merging with the mobile robot to produce a synthesis in which the static computer becomes a thing of the past. The limitations on these machines are imposed primarily by our imaginations!

Synthetic Sensory Systems – Way forward for Neuromorphic Research

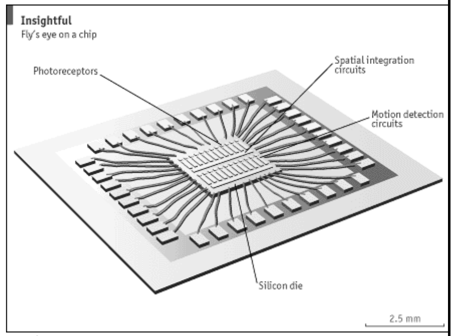

Example: Robot Going straight - 1

One of the many projects demonstrating this sensory application of neuromorphic systems at the annual Telluride meeting last year, was a robot that could drive in straight lines—thanks to electronics modelled on the optic lobe in a fly's brain. The vision chip, built by Reid Harrison at the University of Utah, is a "pixellated" light sensor that reads an image using an array of individual cells, with additional circuitry built locally into each cell to process the incoming signals. The fact that these processing circuits are local and analog is crucial to the device's operation —and is a feature that is borrowed from the biological model.

Dr Harrison and his supervisor identified the various processes taking place in the so-called lamina, medulla and lobular-plate cells in a fly's brain as being worth implementing in silicon. These cells form a system that allows the fly to detect motion throughout most of its visual field—letting the insect avoid obstacles and predators while compensating for its own motion.

Synthetic Sensory Systems – Way forward for Neuromorphic Research

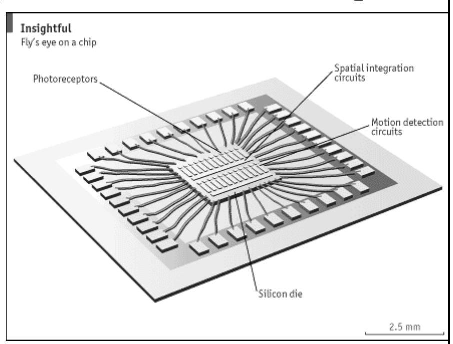

Example: Robot Going straight -2

To prove that the chip not only worked, but could be useful, Mr Harrison attached it to a robot that had one of its wheels replaced by a larger-than-normal one, making it move in circles. When instructed to move in a straight line, feedback from the vision chip—as it computed the unexpected sideways motion of the scenery—was fed into the robot's drive mechanism, causing the larger wheel to compensate by turning more slowly. The result was a robot that could move in a straight line, thanks to a vision chip that consumed a mere *five millionths of a watt of power*.

For comparison, the imaging device on NASA's little Sojourner Rover that explored a few square metres of the Martian surface in 1997 consumed three-quarters of a watt—a sizeable fraction of the robot's total power.

The image system that helps make the "Marble" trackball developed by Logitech of Fremont, California, a handy replacement for a conventional computer mouse, takes its cue from a fly's vision system.

Way Forward in Neuromorphic Systems Research: Conclusions:-1

- Classical artificial intelligence paradigms have failed to deliver their early promise, and one growing belief is that, for the next generation of IT systems to approach human performance, we will need to utilize **hybrid hardware-software brain-like architectures**, where the hardware (e.g. silicon) guarantees the speed and the software (neuronal computational intelligence) guarantees the (robustness of human) performance – approach known as neuromorphic systems.
- Neuromorphic engineers have a wholly different vision of AI. Rather than write a computer program from the top down to simulate brain functions, such as object recognition or navigation, this new breed of “neuromorphic engineers” builds machines that work in the same way as the brain.
- Neuromorphic engineers look at brain structures such as the retina and the cortex, and then devise chips that contain neurons and a primitive rendition of brain chemistry. Also, unlike conventional AI, the intelligence of many neuromorphic systems comes from the physical properties of the analog devices that are used inside them, and not from the manipulation of 1s and 0s according to some modelling formula.
- Neuromorphic engineering is likely to change the face of artificial intelligence because it seeks to mimic what nature does well rather than badly – through millions of years of evolution, nature has come up with some ext remely efficient ways of extracting information from the environment.

Way Forward in Neuromorphic Systems Research - Conclusions:-2

- For centuries, engineers have concentrated on developing machines that were stronger, faster and more precise than people. Whether tractors, sewing machines or computer accounting software, the automata have been simply tools for overcoming some human weakness. But the essential thing has been that they always needed human intelligence to function.
- What neuromorphic engineering seeks to do is build tools that think for themselves—making decisions the way humans do.
- But the neuromorphic route will not be an easy one. The highly efficient analog systems (e.g. the robot going straight) described above are far more difficult to design than their conventional counterparts. Also, billions of dollars have been invested in digital technology—especially in CAD (computer-aided design) tools
- More troubling still, almost all neuromorphic chips developed to date have been designed to do one job, albeit remarkably well. It has not been possible to reprogram them (like a digital device) to do many things even adequately
- However, as work advances, neuromorphic chips will doubtless evolve to be general purpose in a different sense. Instead of using, say, a camera or a microphone to give a machine some limited sense of sight and hearing, tool makers of tomorrow will be buying silicon retinas or cochleas off the shelf and plugging them into their circuit boards !
- At the other extreme, the combination of biological short-cuts and efficient processing could lead to a whole family of extremely cheap—albeit limited—smart sensors that do anything from detecting changes in the sound of a car engine to seeing when toast is the right colour !

Way Forward in Neuromorphic Systems Research - Conclusions:-3

- In fact, the neuromorphic approach may be the only way of achieving the goal that has eluded engineers trying to build efficient “adaptive intelligent” control systems for years. Neuromorphic chips are going to have enormous implications, especially in applications where compactness and power consumption are at a premium—as, say, for replacement parts within the human body.

- This is slowly being recognised. In the biggest annual event on neuromorphic research world-wide, the Telluride workshops, one of the participants recently was a venture capitalist. After genomics, perhaps the next stockmarket buzz will be neuromorphics!