CLORTEX

Machine Intelligence based on Jeff Hawkins' HTM Theory

Fergal Byrne @fergbyrne

A GUY FROM KRAKOW

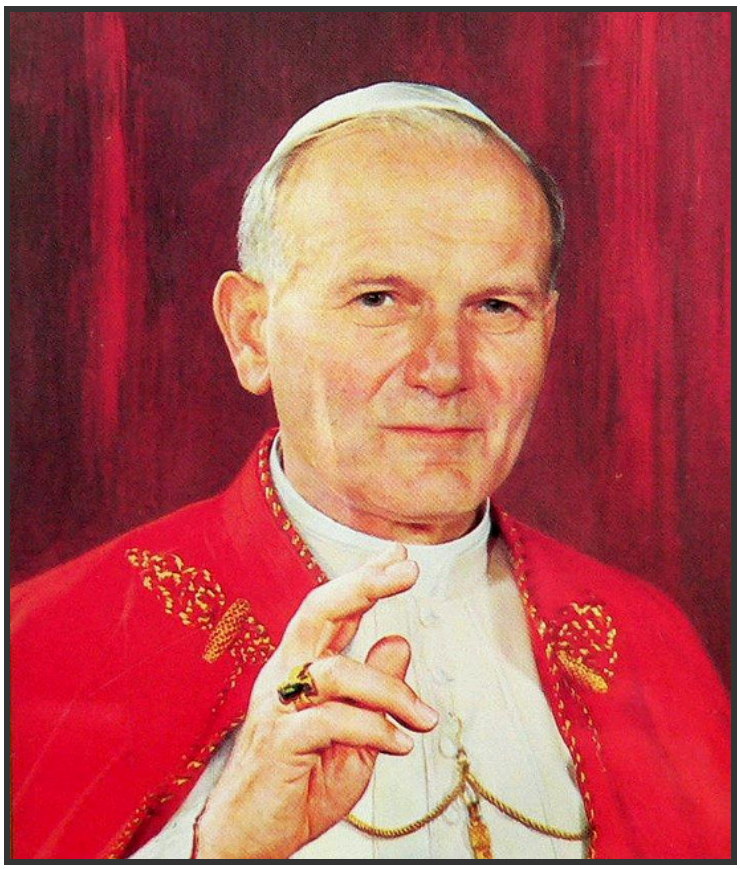

A GUY FROM DUBLIN

- Aged 9, joined 1.25m people to see Pope in Dublin, 1979
- Aged 12, got a 16K Sinclair Spectrum for Christmas..

MOTIVATION

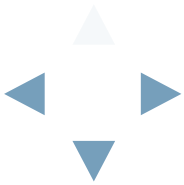

BIG DATA

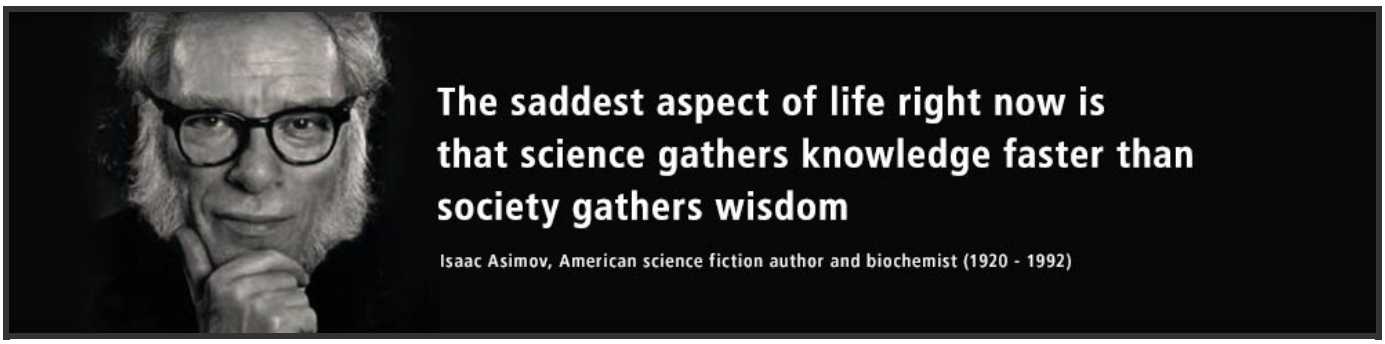

asimovfan.wordpress.com

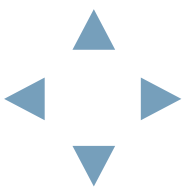

BIG DATA IS LIKE TEENAGE SEX:

- everyone talks about it
- nobody really knows how to do it
- everyone thinks everyone else is doing it
- so everyone claims they are doing it.

Dan Ariely, Center for Advanced Hindsight at Duke University

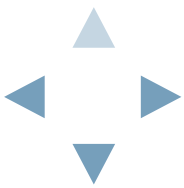

MACHINE LEARNING NEEDS HUMAN INTELLIGENCE

- Running an algorithm is often the easy part
- Most of the work is in preparing the data
- The rest of the work is about humans finding good models
- The machine cranks the wheel...
- Oh, and the rest of the rest of the work is interpreting results

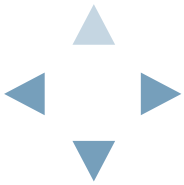

A LIMIT THEOREM

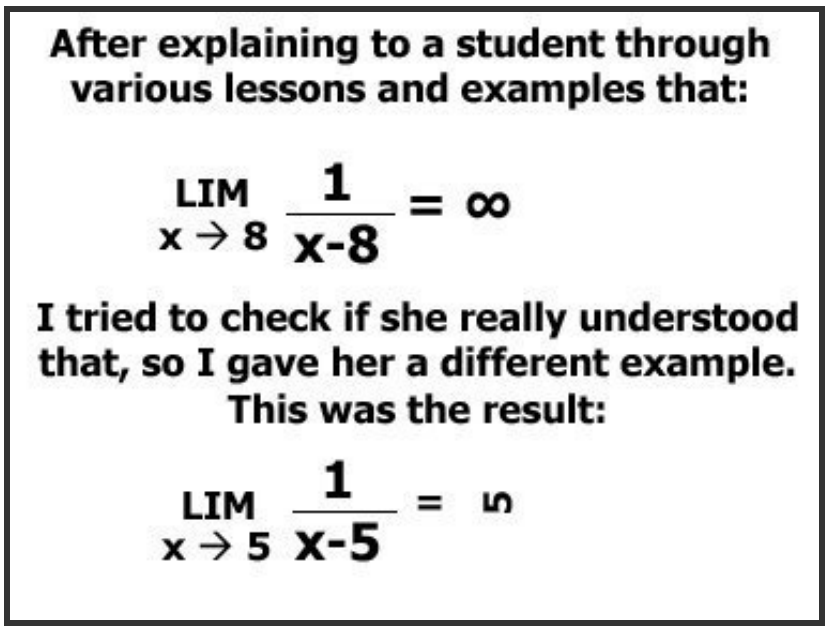

Guillaume & Jennifer Dargaud's website, gdargaud.net

THIS IS IMPORTANT

- Every detail of the financial world
- Security, Military decisions
- Education, Recruitment, Compensation
- Climate change and environmental regulation
- Medicine, viability of transplants, treatment
- <insert more things affecting your life here>

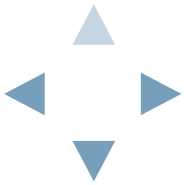

REMEMBER WHAT ISAAC SAID?

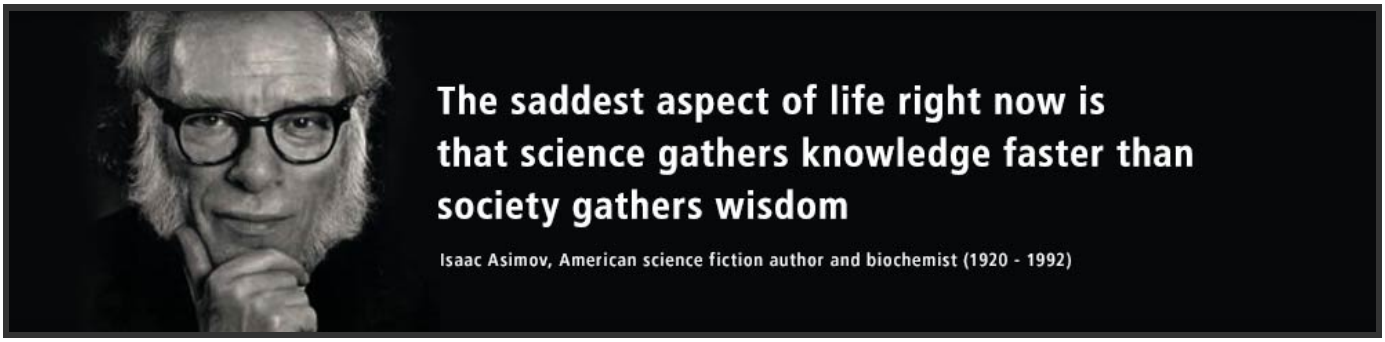

This is about the gap between *human* knowledge and wisdom. When our knowledge is based on unwise interpretation of machine-augmented data science, we are looking at some 'challenges'.

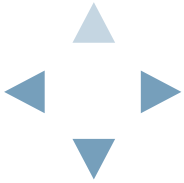

SOME POTENTIAL AVENUES

- Incrementally better machine learning: just a better hammer?
- Can we improve or augment human intelligence?
- Can we identify the mechanisms of intelligence in humans?
- Might intelligent machines save us from our own irrationality?
- Can a machine-human joint approach achieve something new?
- What can neuroscience teach us about intelligence?

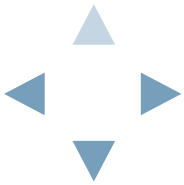

SADLY, ONE BIG BIG DATA PROBLEM IS..

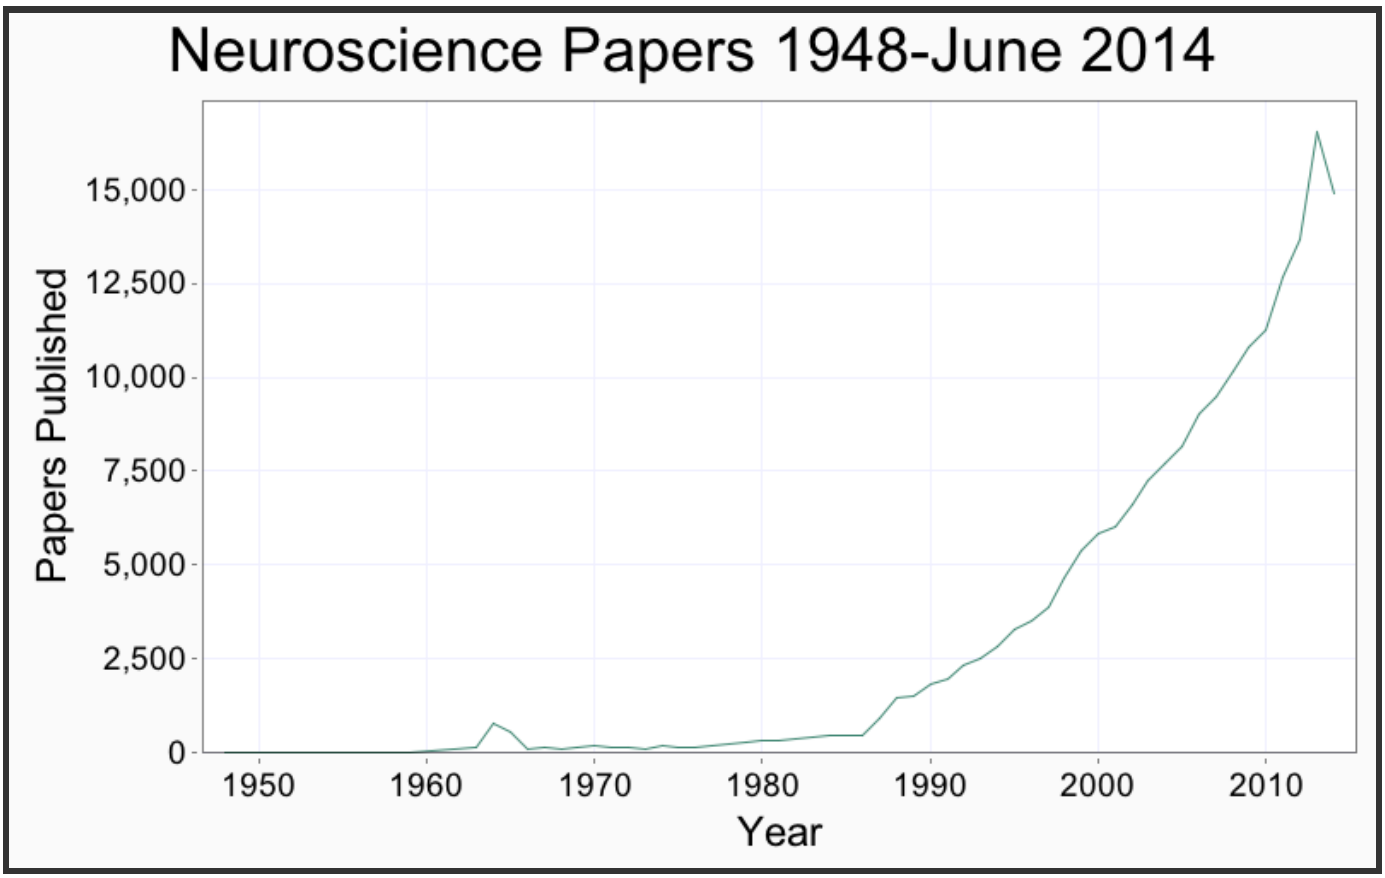

PubMed, June 25th 2014

One every 32 minutes (2013), 17 minutes (2014)

THE BRAIN

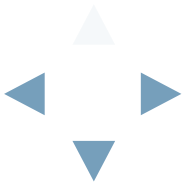

BRAIN, N.:

an apparatus with which we think we think.

Ambrose Bierce, The Devil's Dictionary

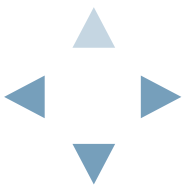

WHY STUDY THE BRAIN?

- Traditional symbolic AI doesn't seem to work
- Perhaps the brain holds the secret to intelligence
- We can learn a lot about ourselves too!

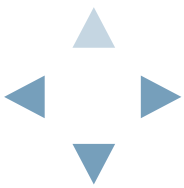

JEFF HAWKINS' GOALS IN HTM

Jeff Hawkins, co-founder of Palm and Numenta

- Study the neocortex and establish its principles
- Build intelligent machines based on these principles
- Wrote *On Intelligence* in 2003, founded Numenta in 2005
- *Numenta Platform for Intelligent Computing (NuPIC)* developed in Python and C++
- Open Sourced NuPIC in 2013

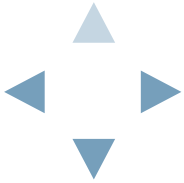

THE NEOCORTEX

The neocortex is the wrinkly part covering our old brain.

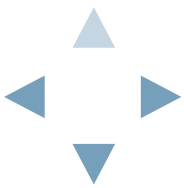

THE NEOCORTEX: 70% OF THE BRAIN

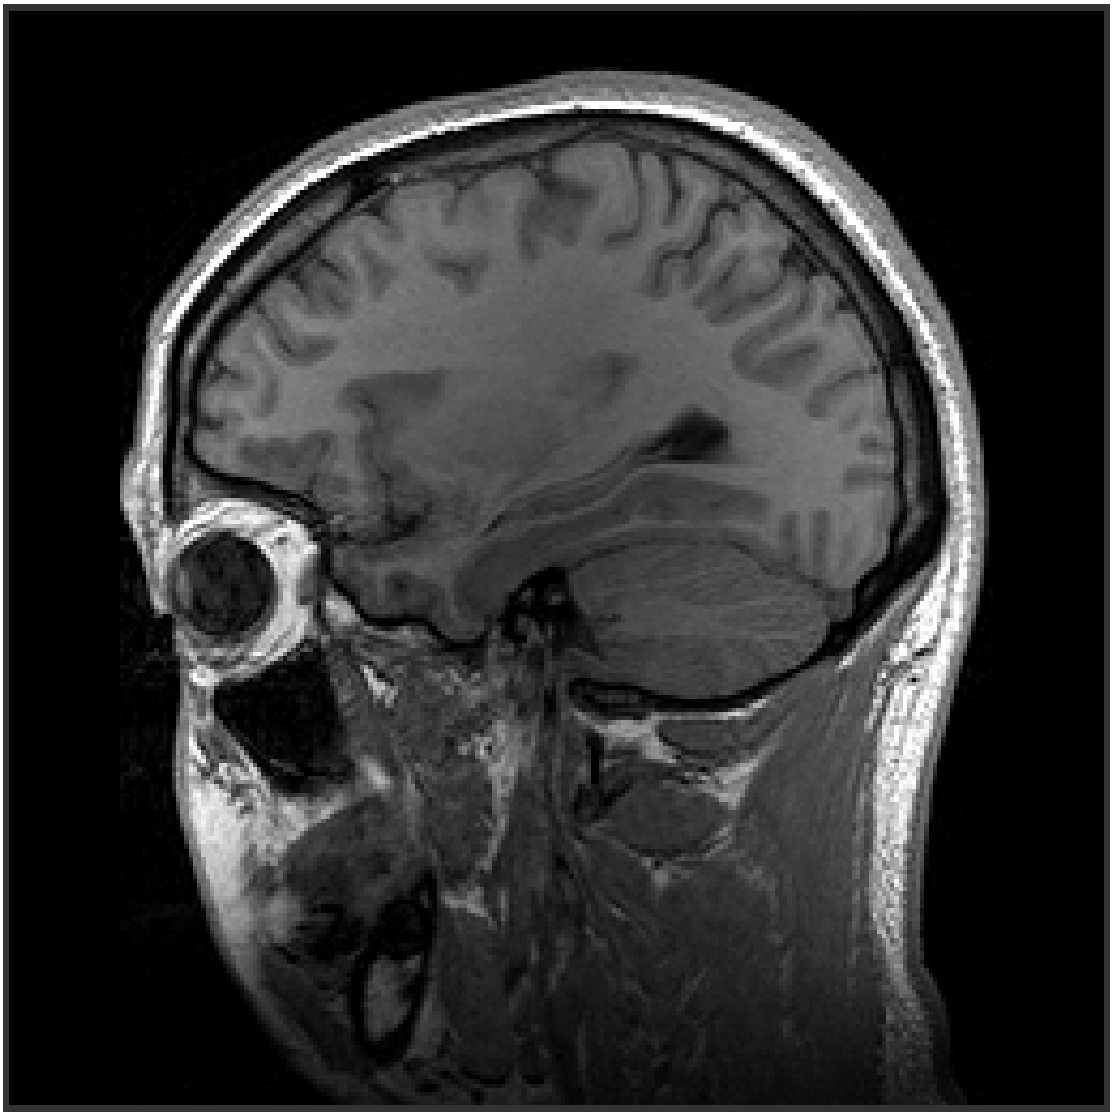

LOBES IN THE NEOCORTEX: SURFACE

Lateral Surface Lobes in Neocortex

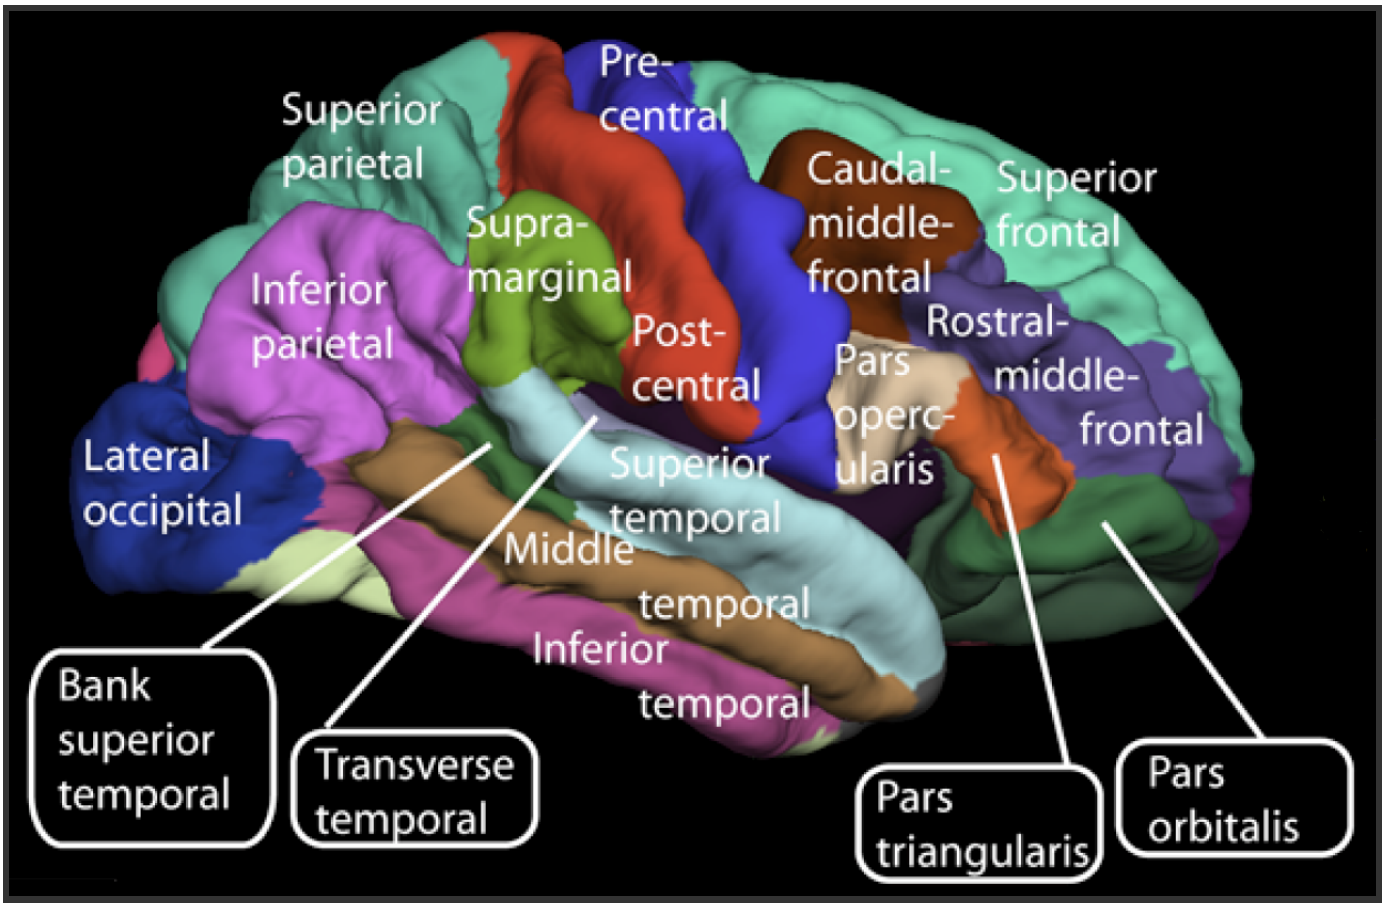

LOBES IN THE NEOCORTEX: MEDIAL

Medial Lobes in Neocortex

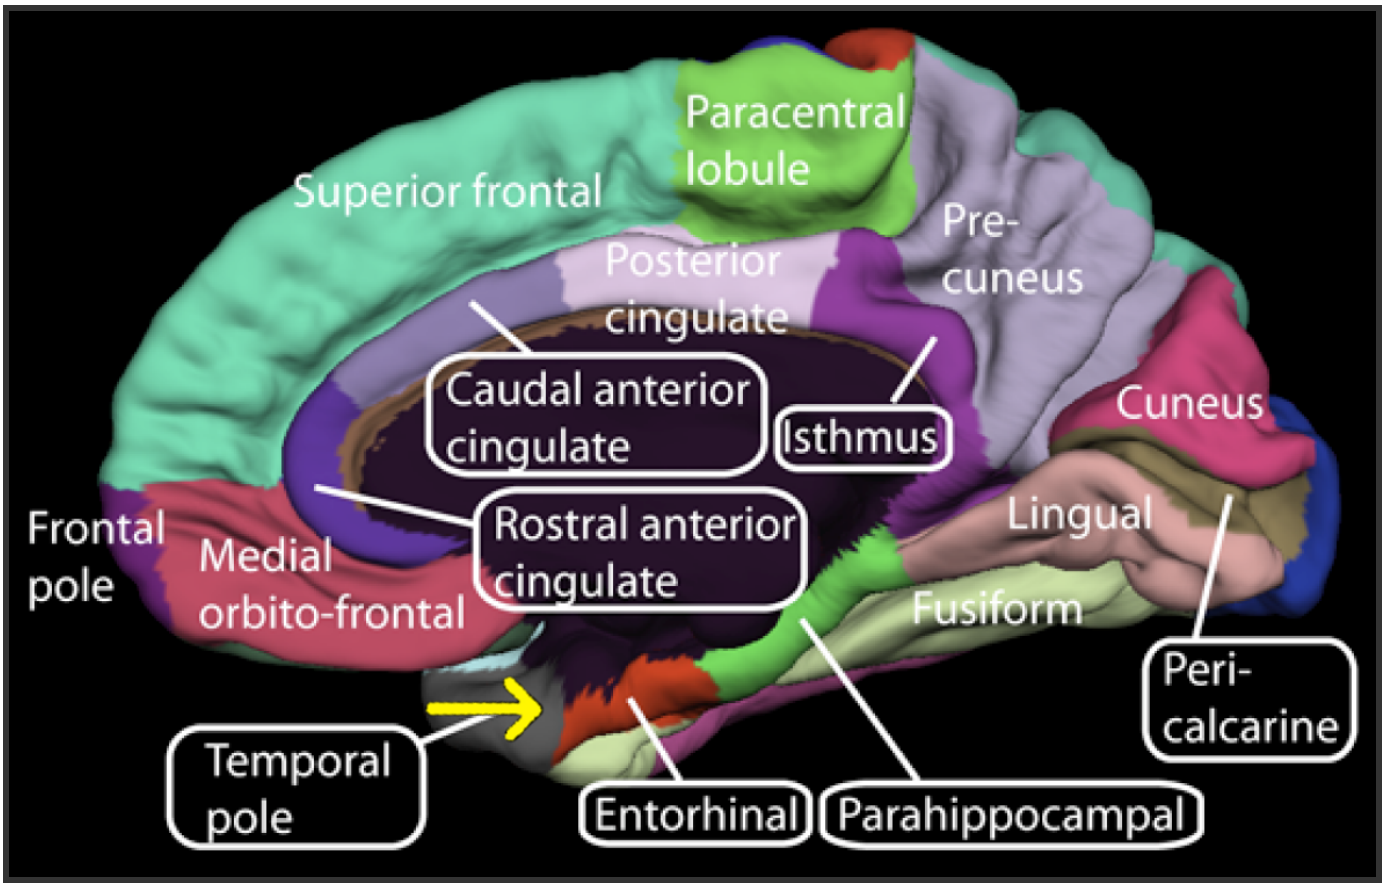

THE NEOCORTEX: SOME FACTS

- About 2mm thick, 100cm^2 in area (or about the size of a dinner napkin!)
- 30-50 Billion neurons (grey matter)
- More than 1 Trillion connections (white matter)
- The seat of intelligence
- The neocortex is **hierarchical and uniform**

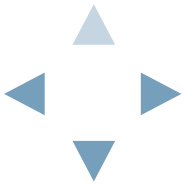

THE NEOCORTEX: HIERARCHY

- The neocortex is divided up into many regions
- Regions form hierarchies
- Every region looks like every other (almost)
- Each region is doing the same thing
- HTM says **all regions have the same algorithm**
- Jeff calls this the **Cortical Learning Algorithm**

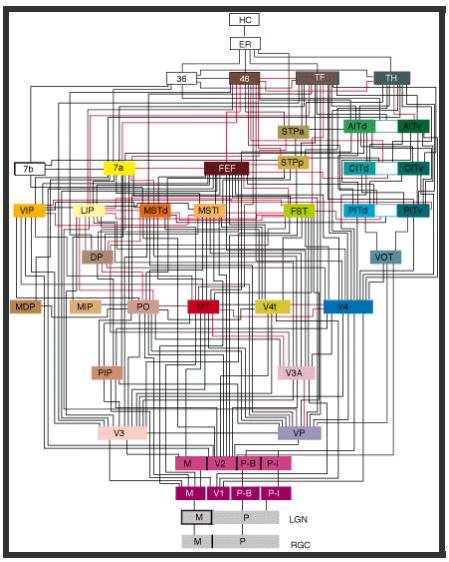

Visual System of the Macaque

THEORY: HIERARCHICAL TEMPORAL MEMORY

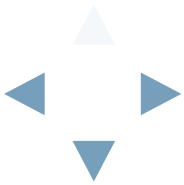

SIX KEY PRINCIPLES

- On-line Learning from Streaming Data
- Hierarchy of Memory Regions
- Sequence Memory
- Sparse Distributed Representations
- All Regions are both Sensory and Motor
- Attention

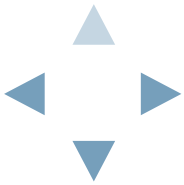

ON-LINE LEARNING FROM STREAMING DATA

- Up to 10 million senses feed the brain
- We don't (can't) store this data
- We build models from live data
- Models are updated with new data

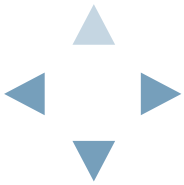

HIERARCHY OF MEMORY REGIONS

- Sensory data enters at the bottom
- Models are built in every region
- Things change more slowly as you go up
- Hierarchy enables sequences of sequences
- The hierarchy works upwards and downwards

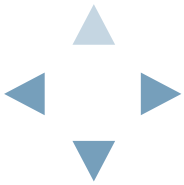

SEQUENCE MEMORY

- All sensory data involves time
- Sequence memory allows predictions
- Structure in data is elaborated over time
- Sequences can be composed in hierarchy

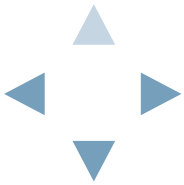

SPARSE DISTRIBUTED REPRESENTATIONS

- In each region, many neurons, few are active
- SDRs represent spatial patterns
- SDRs have many useful properties:
- Fault-tolerant, semantic operations, high-capacity
- Key to understanding and building intelligent systems

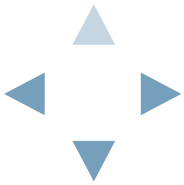

ALL REGIONS ARE BOTH SENSORY AND MOTOR

- Every region processes sensory data and produces behaviour
- Behaviour provides context for sensory data
- Structure in the model is navigated via behaviour
- The neocortex learns to control the old brain
- A sequence memory is a sensorimotor model of the world

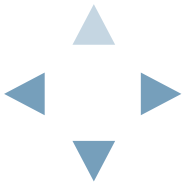

ATTENTION

- We use attention to manage the neocortex
- Attention allows for planning and previsualisation
- Novel data or anomalies can 'demand attention'
- Whole sub-hierarchies can be switched on or off

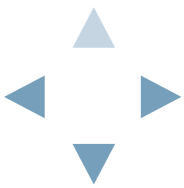

THE NEOCORTEX - LAYERS

Drawing by Ramon y Cahal, 1911

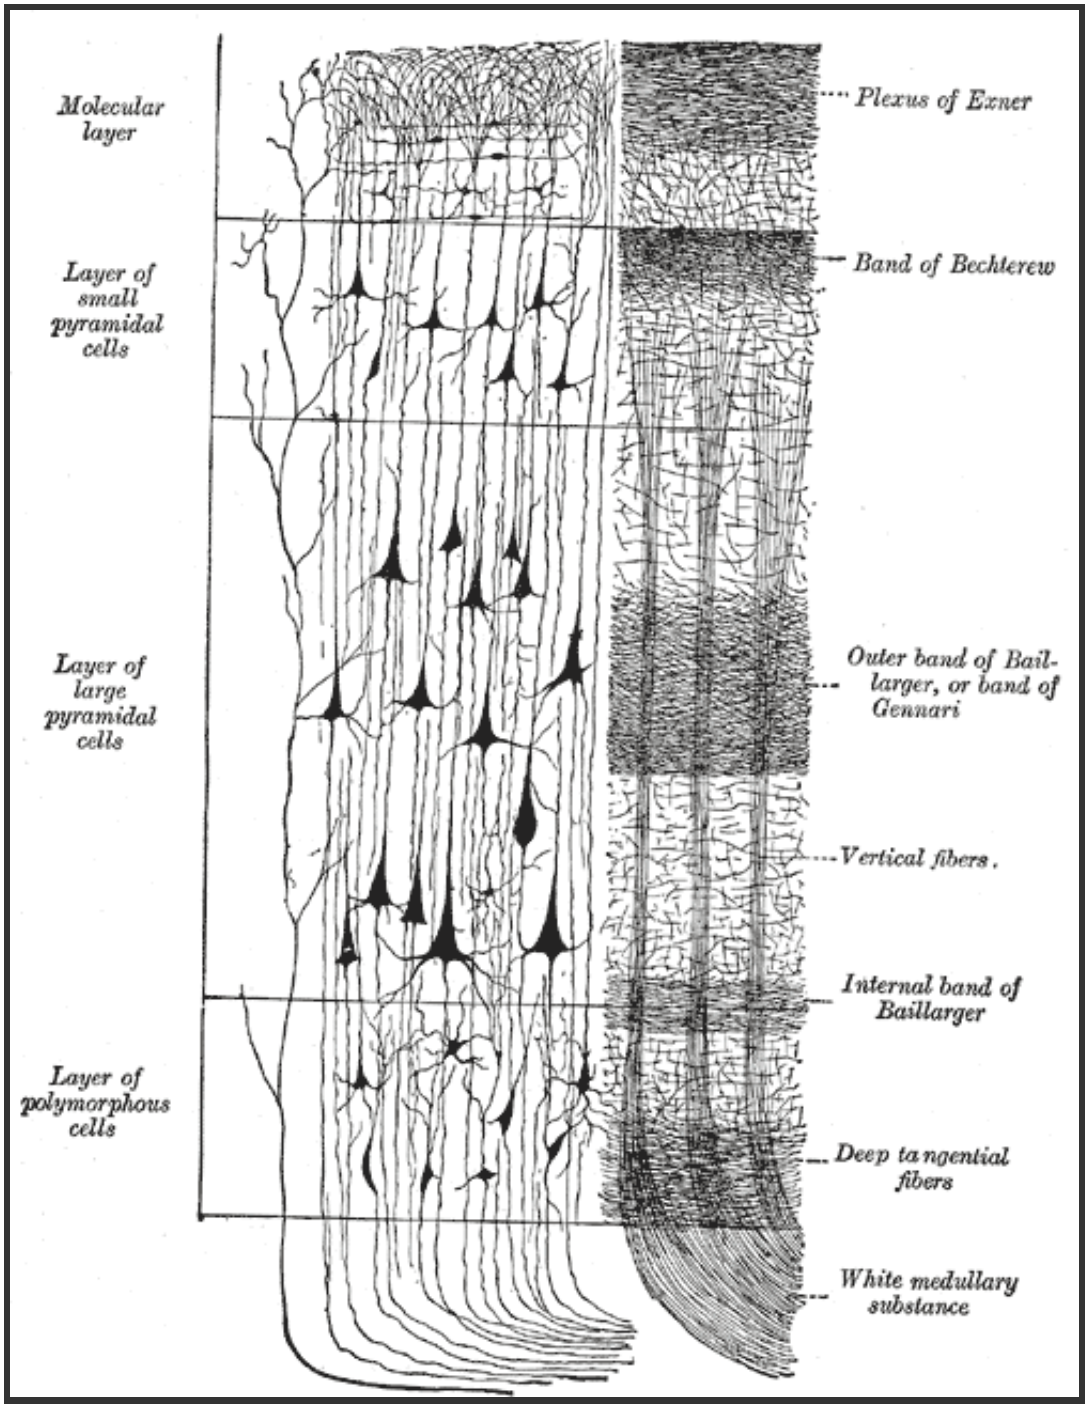

THE NEOCORTEX - NEURONS

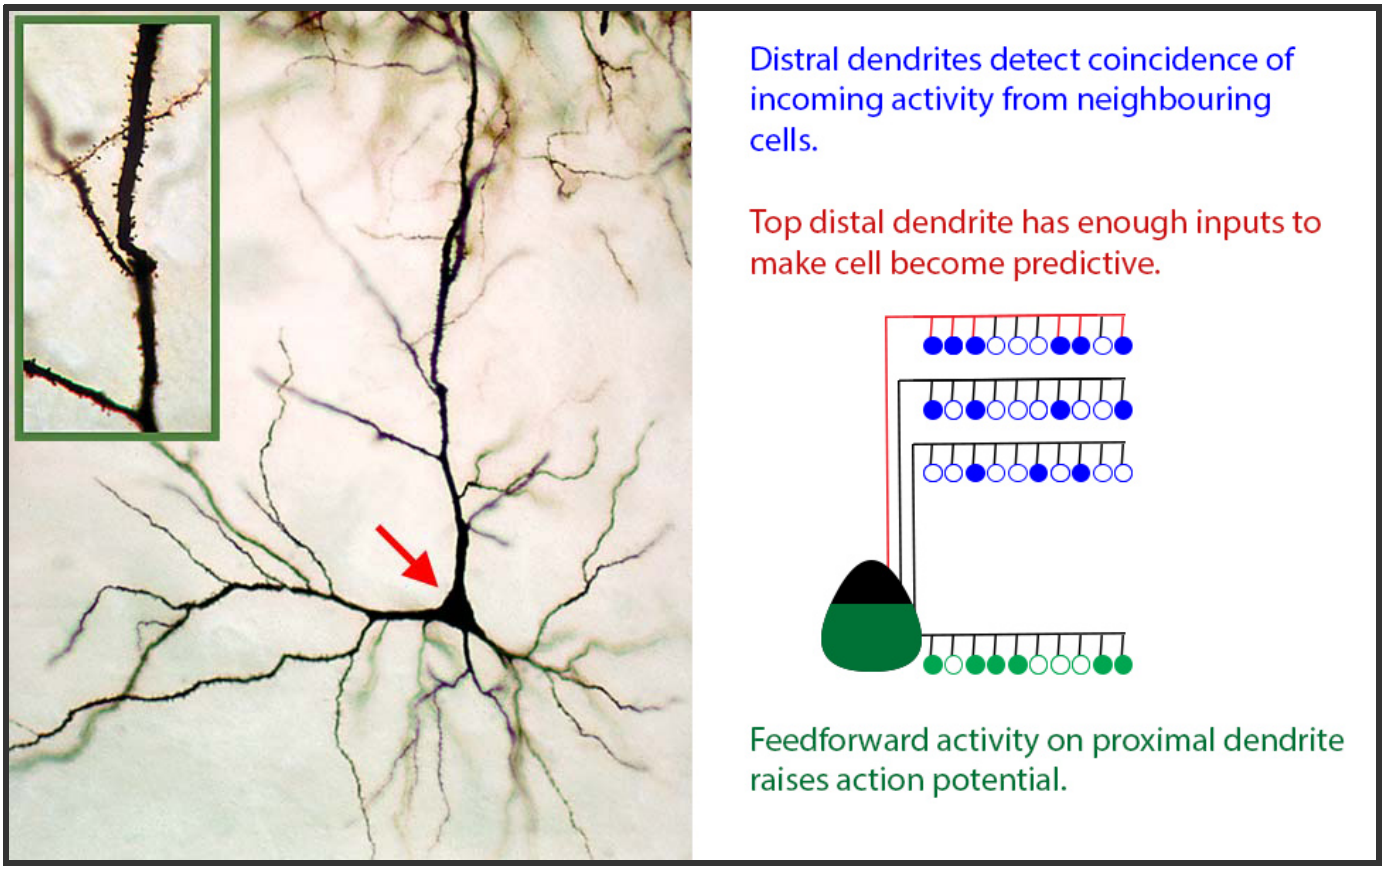

CLORTEX

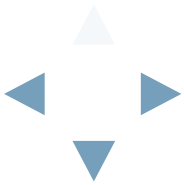

NUMENTA'S NUPIC

- In development since 2005
- Partially implements HTM/CLA
- Written in Python and C++
- Open Source - see Numenta.org

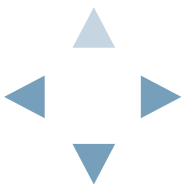

NUMENTA'S NUPIC: STRENGTHS

- Skilled dev team at Numenta with Jeff leading
- Numenta eat their own dog food - Grok uses NuPIC
- Operates using a subset of HTM/CLA principles
- Tunable using swarming on your data
- Works well on streaming scalar (e.g. machine-generated) data
- Great community - join us at Numenta.org

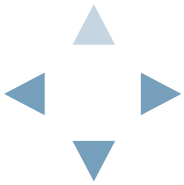

LIMITATIONS

- Codebase evolved over time, not built in one go
- Difficult/scary to rewrite for flexibility
- Uses OO with large, coupled classes (~1500 LOC per class)
- Need to swarm to find parameters, no real-time control
- Not easy to extend beyond streaming scalar use case

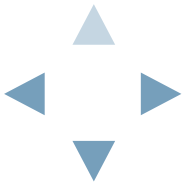

ARCHITECT, N.:

One who drafts a plan of your house, and plans a draft of your money.

Ambrose Bierce, The Devil's Dictionary

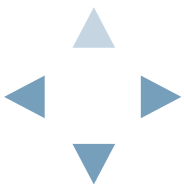

CLORTEX: REQUIREMENTS

- Directly Analogous to HTM/CLA Theory
- Transparently Understandable Source Code
- Directly Observable Data
- Sufficiently Performant
- Useful Metrics
- Appropriate Platform

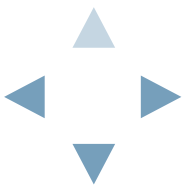

DIRECT ANALOGY TO THEORY

- Each element of the theory appears in the software
- Regions contain Layers of Neurons in Columns
- Neurons have Proximal and Feedforward Dendrites
- Synapses connect/disconnect based on Permanence
- Strategies for Connection, Inhibition, Topology etc.

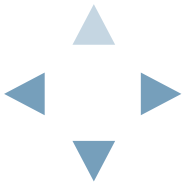

RUSS MILES: AXIOMS FOR ARCHITECTURAL SIMPLICITY

- Your Software's First Role is to be Useful
- The best software is that which is not needed at all
- Human Comprehension is King
- Machine Sympathy is Queen
- Software is a Process of R&D
- Software Development is an Extremely Challenging Intellectual Pursuit

CLORTEX DESIGN DECISIONS

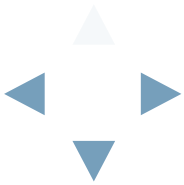

#1 JUST USE DATA!

- Everything modelled using simple maps, vectors and sets
- Layers are vectors of columns (vectors of neurons)
- Neurons are a map of `:proximal-dendrite` and `:distal-dendrites`
- Dendrites are vectors of synapses
- Synapses are maps with `:permanence` and `:pre-synaptic-neuron` (a ref)
- All done in a one-page Datomic (adi) schema

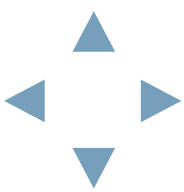

ADI SCHEMA FOR CLORTEX

```
(def clortex-schema
  {:patch   {:type     [{:type :keyword}]
             :name     [{:type :string}]
             :uuid     [{:type :uuid}]
             :timestep [{:type :long :default 0}]
             :columns [{:type :ref
                        :ref  {:ns   :column
                               :rval :patch}
                        :cardinality :many}]
             :neurons [{:type :ref
                        :ref  {:ns   :neuron
                               :rval :patch}
                        :cardinality :many}]
             :inputs   [{:type :ref
                         :ref  {:ns   :dendrite
                                :rval :patch}
                         :cardinality :one}]}
   :column  {:type     [{:type :keyword}]
```

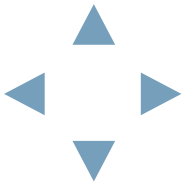

#2 CLOJURE & ITS ECOSYSTEM

- Clojure Data instead of 'Domain Objects'
- Algorithms just functions of data
- 'Components' just look at data
- Composable, swappable, scaleable
- Use 'standard' libraries in simple combinations

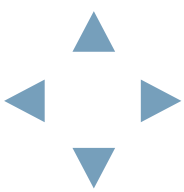

#3 APPLY RUSS MILES' LIFE PRESERVER

- Answers questions about 'where does this go?'
- Everything's either 'core' or 'integration'
- Core: a datomic database for the neocortex
- Core: each 'patch' of neurons is a graph (a map)
- Integration: algorithms, encoders, classifiers, SDRs
- + visualisers, management, metrics...

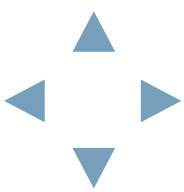

KEY CLOJURE LIBRARIES & TOOLS

- datomic for the core (+adi in places)
- quil/Processing libs for visualisation and GUI
- incanter for exploratory data science
- lein-midje-doc for literate documentation/test
- hoplon-reveal-js for presentations
- LightTable for a lovely editing experience

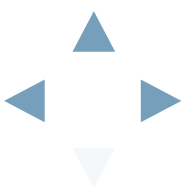

SUMMING UP

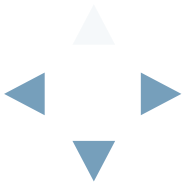

REVIEW

- Big Data is not just a Machine Intelligence Problem
- Need to understand and augment human intelligence too
- HTM is an exciting R&D project
- Using Clojure's thinking and tools can change the game
- I would love to interest Clojure's great community in HTM

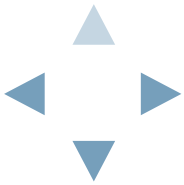

RESOURCES

- http://numenta.org/
- http://inbits.com/
- http://github.com/fergalbyrne/clortex

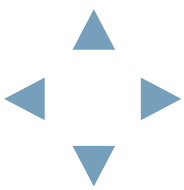

THANKS & ADIOS

- Jeff Hawkins, Matt & many friends at Numenta and NuPIC
- Rich Hickey, Stuart and many more in the Clojure world
- Alex Miller, Carin Meier for their friendly advice