




DEEP LEARNING FOR PATIENT FLOW

MALCOLM PRADHAN, CMO

- ▶ Why are smart machines are important for health care
- ▶ The emergence of deep learning
- ▶ Deep learning vs existing methods
- ▶ Some early results
- ▶ Practical tips on getting started
- ▶ Future directions



- ▶ Health informatics company with products in
 - ▶ Patient flow & bed management
 - ▶ Emergency Department
 - ▶ Outpatient and referrals management
 - ▶ SmartForms, clinical decision support (CDS)
- ▶ A health informatics approach
 - ▶ Computers should play a more active role in health care
 - ▶ Assist clinical staff so that the right thing to do is the easier thing to do
- ▶ We want to turbo charge our products using advances in smart machines (AI)

- ▶ 448 nodes, > 900 connections, > 90m probabilities



- ▶ With an aging population the demand for health care is increasing rapidly

Number of People Over 85 yo in Australia



- ▶ Problems of safety, productivity, variation
- ▶ How else do we scale health care?

- ▶ Increasingly complex patients
 - ▶ Increased referrals to allied health and other specialties
 - ▶ Higher resource utilization, difficult to predict ahead of time
- ▶ Current models
 - ▶ Predict ED admissions and future admissions
- ▶ Challenges
 - ▶ Predicting detailed resource needs ahead of time
 - ▶ Early detection of variation
 - ▶ Logistical support for bed management
 - ▶ Using clinical context to better understand patient needs
- ▶ AI has been around for ages, why hasn't it helped us?




ALCIDION A (VERY) BRIEF HISTORY OF AI IN HEALTH



- ▶ Expert level performance since 1970's
- ▶ Clinical Decision Support for
 - ▶ Diagnosis
 - ▶ Management
 - ▶ Safety

AAPHelp, Internist-1, Mycin, Casnet, PIP, Oncocin, DxPlain, QMR



- ▶ Difficult to integrate into workflows
- ▶ Not integrated with IT systems
- ▶ Brittle
- ▶ Difficult to maintain over time
- ▶ Not easy to localize
- ▶ Descriptive, based on expert opinion

Traditional 'AI'

Expertise



Model & structure
(Knowledge base)

Algorithms to
update model


Neural Networks

Data*


Learning
algorithms, simple
update



1. Technical improvements to deal with deep networks
2. Large data sets e.g. ImageNet has >14m annotated images
3. GPU






- ▶ Image description
- ▶ Text generation
- ▶ Self-driving cars
- ▶ Playing complex games
- ▶ Early work in demonstrating
 - ▶ Radiology
 - ▶ Histopathology
 - ▶ Risk detection



A close up of a child holding a stuffed animal.

Deep neural networks learn hierarchical feature representations





- ▶ Data generated from anonymised HL7 feeds
 - ▶ Admissions, readmissions, ICDs, DRGs
 - ▶ Later adding labs, radiology reports, OPD
- ▶ Multilayer perceptrons (MLP)
 - ▶ Tested 2-5 layers, many variations
 - ▶ Significant work in encoding categorical variables such as DRGs and ICDs
 - ▶ Approx 8,000 x 200,000 matrix of data
- ▶ Hardware & software
 - ▶ Linux based PC with Nvidia 980 Ti
 - ▶ Keras using Theano back-end
 - ▶ Python for running experiments and data transformation

► Readmission

- Detection of readmission < 30 days against
- Usually for COPD, Heart failure, Pneumonia, AMI, Total hip/knee arthroplasty
- On par with best published algorithms based on ICDs and LACE tool approx. 0.65 AUC (better than traditional models on this data set)
- 3 layers seems to work well, 4-5 have increased training times with minor performance benefits
- We are exploring an ensemble approach with multiple disease based models for major conditions


► General hospital demand prediction

- Similar performance to time series models (approx 9% error 1 month prediction) compared to about 15-20% error from historical average methods
- Require much more data to learn
- Good performance from ARIMA models

Futoma, K. et. al. A comparison of models for predicting early hospital readmissions. Jof Bio Inform 56 (2015) 229–238

Kim, K. et. al. Predicting Patient Volumes in Hospital Medicine: A Comparative Study of Different Time Series Forecasting Methods

- ▶ Very data hungry
 - ▶ Can't provide hints about the domain
- ▶ Confidently make errors
 - ▶ Recognised with $> 99.6\%$ probability
- ▶ Not explainable, hard to debug
- ▶ Embed probabilities, making them less portable to new settings
- ▶ Still descriptive, learn from what has been done



- ▶ Large volumes of data required
 - ▶ 100,000+ cases
- ▶ Non-image methods are still being developed
 - ▶ Representation of categorical data
- ▶ Research moves ahead fast using a direct publish model (arxiv)
- ▶ Many experiments required over hours – days
 - ▶ Configurations, epochs, batch sizes, etc.
- ▶ Hardware & software
 - ▶ You will need a powerful Nvidia GPU and configured Linux machine
 - ▶ Open source software is available: Theano, Tensorflow (Google), Torch (Facebook), CNTK (Microsoft), Keras

▶ Collaboration

- ▶ Aim to create Health Informatics Machine Learning group, currently discussing with UniSA, University of Adelaide
- ▶ Working to integrate more data sources e.g. community

▶ Improved models

- ▶ Further incorporation of clinical data such as labs
- ▶ Research to incorporate histopathology reports, radiology reports using unsupervised techniques

▶ Building smart machines

- ▶ Make software part of the health care team, rather than a barrier to productivity
- ▶ Allow patients to monitor their own health and navigate health care