



Using Matlab to Develop Respiratory Disease Diagnostic Tools

Vesa Peltonen
Senior Software Engineer
vesa@resapphealth.com.au

Matlab Conference Brisbane
4 May 2017

Acknowledgements

The technology used by ResApp Health has been and is continuously developed by Associate Professor **Udantha Abeyratne** at The University of Queensland.


Dr Abeyratne's team has been engaged in the research and development of this technology since 2009 and has been funded by the Bill and Melinda Gates Foundation, The University of Queensland and UniQuest.

Digital healthcare for respiratory disease

- ResApp Health is developing the world's first clinically-tested, regulatory- cleared respiratory disease diagnostic test and management tools for smartphones
- Diagnosis of respiratory disease is the most common outcome from a visit to the doctor
- Huge global market, 700M+ doctor visits annually for respiratory disease
 - Unique opportunity to integrate into telehealth providers' existing platforms
 - Strong demand also seen within clinics, emergency rooms and outpatient facilities
 - Currently diagnosed using stethoscope, imaging (x-ray, CT), blood and/or sputum tests
- → Time consuming, expensive and not very accurate

Digital healthcare for respiratory disease

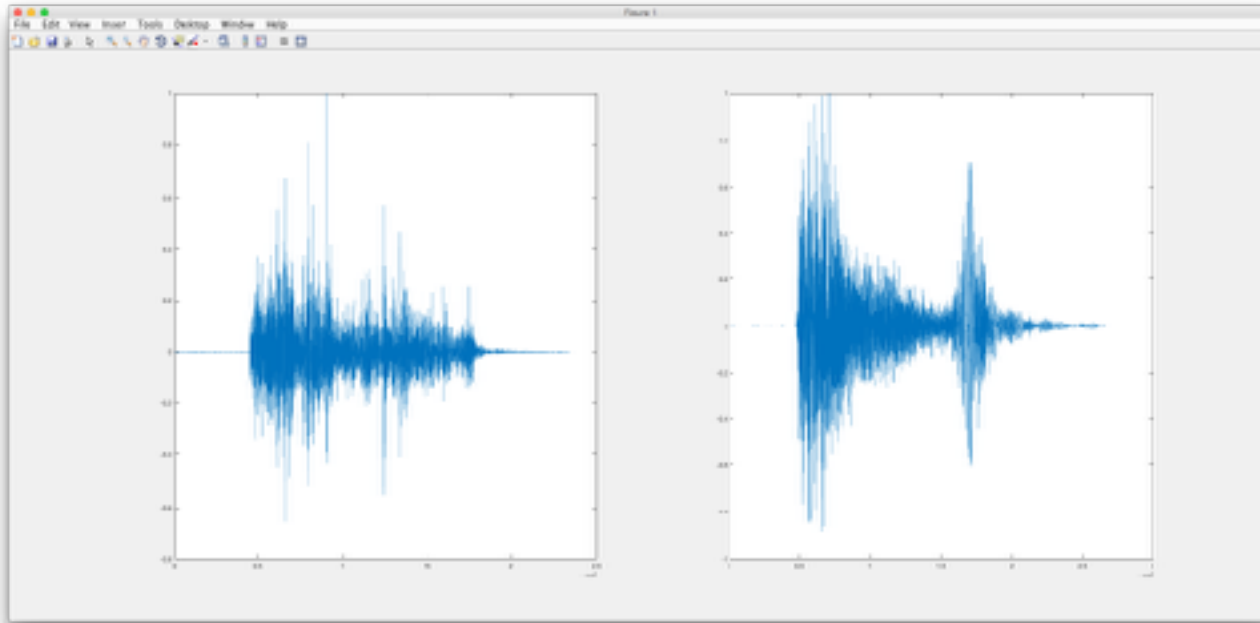
<https://www.resapphealth.com.au/solutions/>



Respiratory disease diagnosis using only the sound of a patient's cough

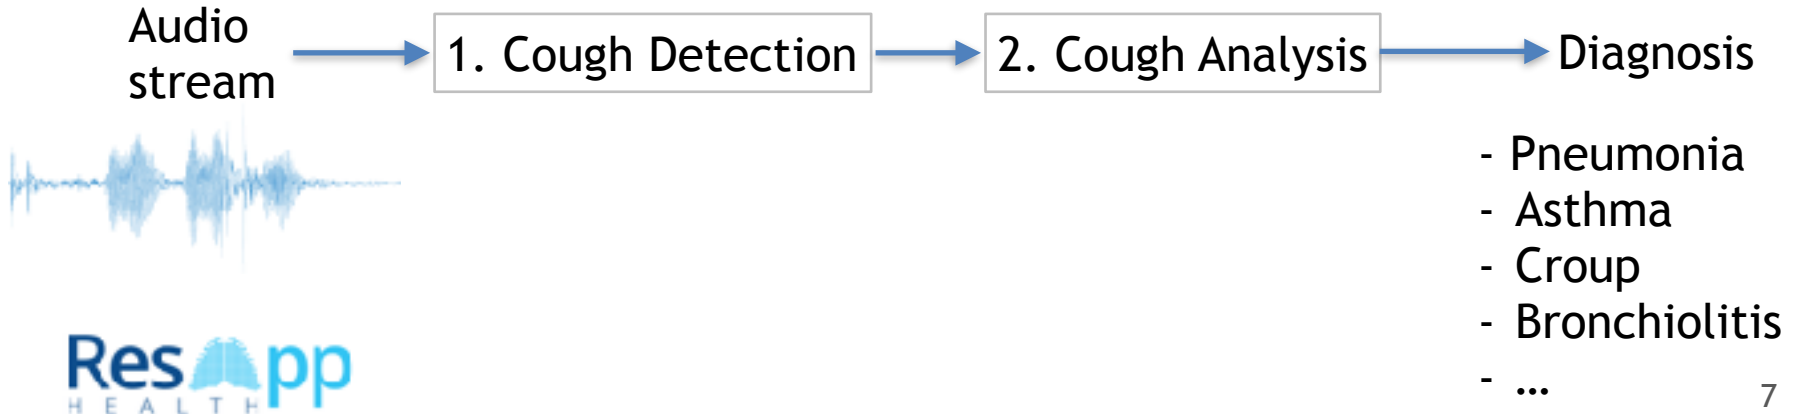
- The technology is based on the premise that cough sounds carry **vital information on the state of the respiratory tract**
- We have taken a *supervised machine learning* approach to develop highly accurate algorithms which diagnose disease from cough sounds
- In the next slide there is an example of a healthy cough and a pneumonia cough
 - Can you guess which one is healthy and which one is pneumonia?

Two different coughs



Automatic respiratory disease diagnosis

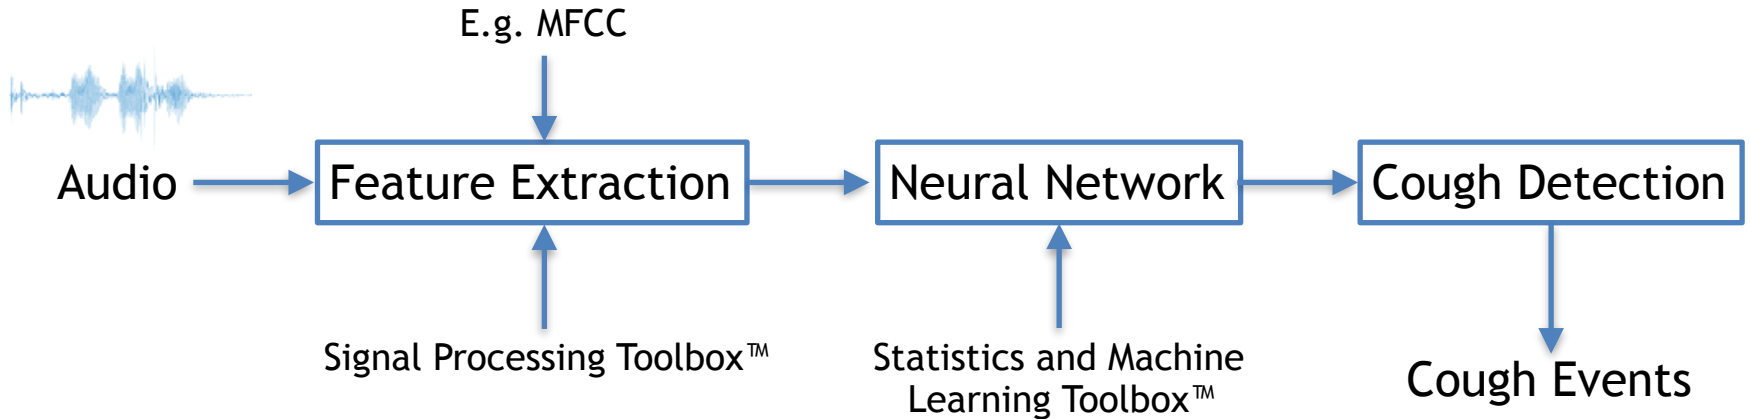
- The system consists of two main parts:
- Front-end: finds cough sounds in a continuous audio stream
- Cough Analysis: provides a diagnosis based on learned cough signatures for various respiratory diseases



Part 1: Detecting the coughs

- The front-end of the system finds the cough sounds in a continuous audio stream
- The cough detection system is able to filter out speech and other background noises
- The model has been trained with many hours of cough and non-cough sounds

Part 1: Detecting the coughs




Audio features

- Audio features are calculated from the raw audio signal in processing windows
- The audio features can be roughly divided into two categories:
 - Time domain features. E.g. Zero-Crossing Rate (ZCR)
 - Frequency domain features. E.g. Mel-Frequency Cepstral Coefficients (MFCC)
- Traditionally feature vectors are used instead of raw audio as input for ML algorithms
 - Features reduce the dimensionality of the input
 - Some features, like MFCC, approximates the human auditory system's response

Audio features

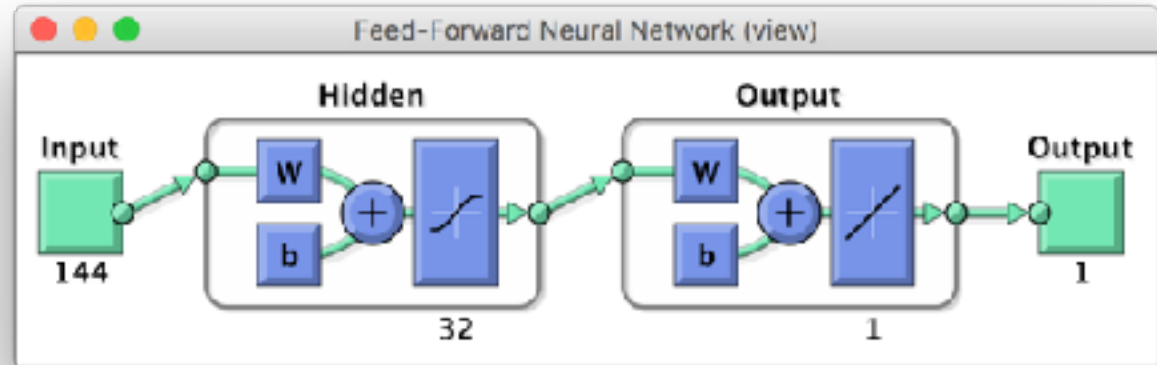
E.g:
ZCR
MFCC



Feature Vectors

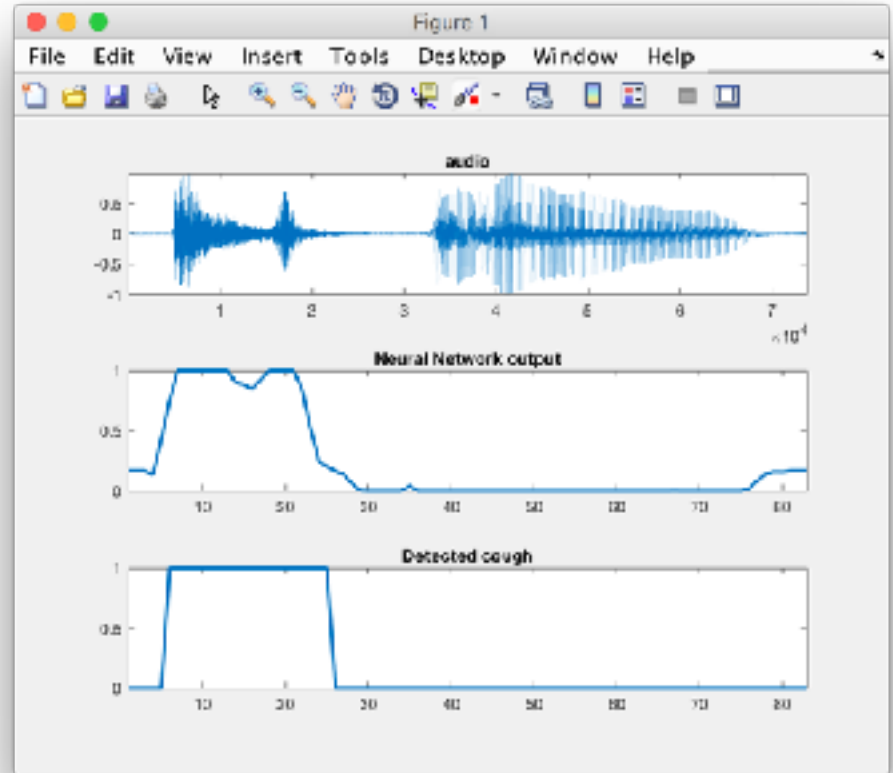
Feed-forward neural network

- Neural Network is a collection of artificial neurons
- Each neuron is a function of weighted sum of its inputs (+a bias).
- Number of the inputs equals to the number of features in the feature vector
- One can have any number of hidden layers and any number of nodes in each layer



Example of detecting a cough

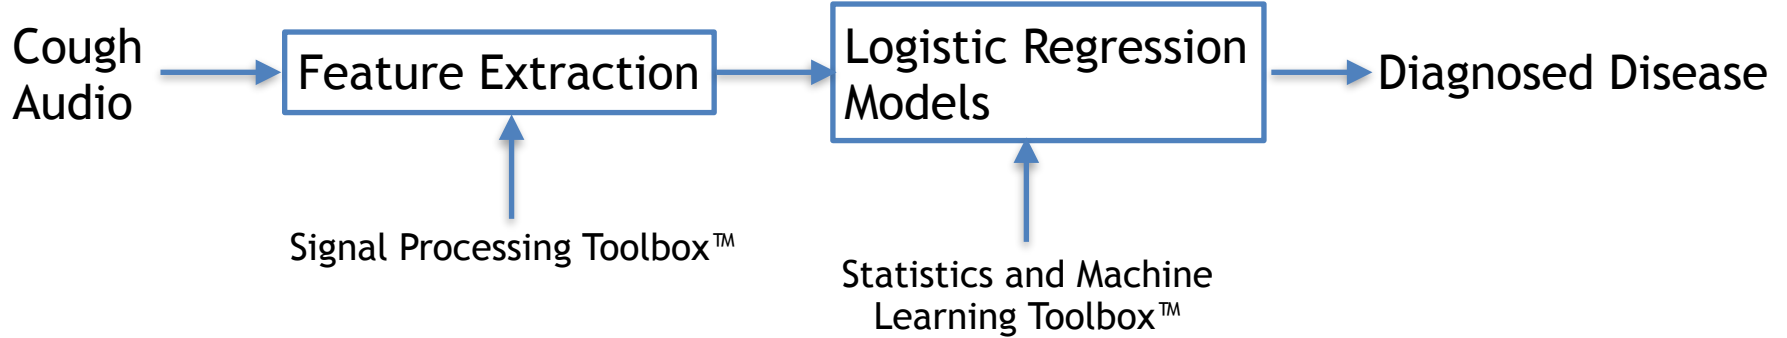
1. Audio Signal
2. Probability output of a neural network
3. Detected coughs



Part 2: Analysing the coughs

- The cough sounds are fed into the analysis part
- It extracts various audio features for each cough
- The features are then fed into previously trained logistic regression models and classified into diseases or control classes
- Each model is a binary classifier (control vs. disease). I.e. the approach is so called one vs. many
- Number of coughs are used per diagnosis to increase the accuracy of the diagnosis

Part 2: Analysing the coughs



Easy to use, instant diagnosis using only a smartphone

- The algorithms are developed using Matlab™
- The final algorithms are deployed on various smartphone platforms
- The end-to-end diagnosis runs real time on a phone
- Easy to use



Verified by compelling pediatric clinical evidence

Breathe-Easy Study Pediatric Results



| | Sensitivity | Specificity | Accuracy |
|---|-------------|-------------|---------------|
| Pneumonia vs. no respiratory | 100% | 95% | 97% |
| Asthma vs. no respiratory | 97% | 92% | 95% |
| Bronchiolitis vs. no respiratory | 100% | 100% | 100% |
| Croup vs. no respiratory | 94% | 100% | 99% |
| URTI vs. no respiratory | 100% | 95% | 96% |
| Pneumonia, croup or bronchiolitis vs. URTI⁴ | 89-100% | 90-95% | 89-98% |
| Differential diagnosis of pneumonia, croup, URTI and bronchiolitis | 91-99% | 89-98% | 89-98% |

Building strong clinical evidence in adults

| Breathe-Easy Study Adult Results | Sensitivity | Specificity | Accuracy |
|-------------------------------------|-------------|-------------|----------------|
| <i>COPD vs. no respiratory</i> | 100% | 96-100% | 98-100% |
| <i>Asthma vs. no respiratory</i> | 91% | 91-93% | 91-92% |
| <i>Pneumonia vs. no respiratory</i> | 97-100% | 100% | 98-100% |
| <i>URTI vs. no respiratory</i> | 100% | 100% | 100% |
| Asthma vs. COPD | 93% | 96% | 94% |
| Pneumonia vs. Asthma | 92% | 81% | 88% |
| Pneumonia vs. COPD | 92% | 92% | 92% |

Machine learning: Lessons learned

- The **simplest** model that fits the data is also the most probable model (Occam's razor)
- Why is simple better:
 - There are fewer simple hypotheses than complex ones. If we manage to fit a less likely model (i.e. a simple one) it is more significant than fitting a complex model
 - If our in-sample error rate is very small, but out-of-sample error rate is big, our model is too complex for the given data. The model has learned the data too much (overfit)



Machine learning: Lessons learned

- Partition your data into **three** sets
 - **Training set** (>70%)
 - train your model with this data
 - **Dev set** (<15%)
 - tune your parameters with this data
 - **Test set** (<15%)
 - estimate your accuracy with this set
 - Don't over use the test set. For example, if you run 100 different experiments and test them with the same test and then pick the best result. This result is not a good estimate anymore for the out-of-sample error rate

Machine learning: Lessons learned

- Use single number evaluation metric:
 - Helps you quickly evaluate algorithms and models, and therefore iterate faster
 - For example use F1-score instead of recall and precision.
- Analyse error systematically
 - Categorise the errors and find the single reason that contributes most to the error rate
 - Start addressing the error sources from the most significant one
 - For example: with the cough detection we found that cries were a prominent source of false positives