

LISTEN TO YOUR INHALER, IT MIGHT BE TELLING YOU SOMETHING

In this article, Kian Min Lim, PhD, Data Scientist, and David Harris, Head of Respiratory Drug Delivery, both of PA Consulting Group, discuss how tailored algorithms can accurately detect inhalers' acoustic signatures, effectively adding connectivity without having to add a chip to the inhaler or, indeed, modify it in any way.

Machine learning technology has advanced in leaps and bounds, spearheaded by the rapid progression of deep learning. Deep learning – widely considered as narrow artificial intelligence (AI) – has been shown to perform a well-defined task at, or beyond, the human expert level.¹ The deep learning algorithm can be designed to perform a singular well-defined scope or task effectively and accurately. In fact, there are many deep learning algorithm implementations in the consumer sector, for example, Apple's Siri, Samsung's face unlock, Google's voice assistant, and many others.

Leveraging these advancements, we can transfer some of these deep learning

“Our group has demonstrated that an image recognition algorithm can be adapted to listen to an inhaler during normal use, and produce a remarkable 98% confidence level of correctly identifying the inhaler's acoustic signature.”

techniques to innovate current inhaler offerings. Our group has demonstrated that an image recognition algorithm can be adapted to listen to an inhaler during normal use, and produce a remarkable 98% confidence level of correctly identifying the inhaler's acoustic signature.²

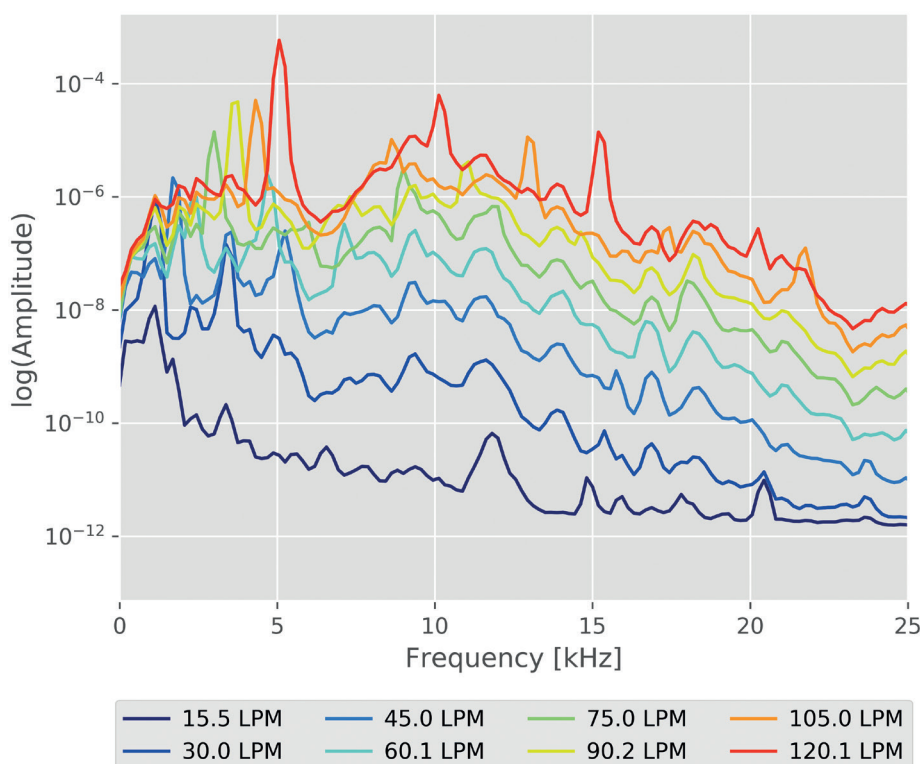


Figure 1: Frequency response plots for Sun Pharma's Starhaler DPI.



Dr Kian Min Lim
Data Scientist
T: +44 7970 690 859
E: kianmin.lim@paconsulting.com



David Harris
Head of Respiratory Drug Delivery
T: +44 1763 267 608
E: david.harris@paconsulting.com



PA Consulting Group
Cambridge Technology Centre
Melbourn
Herts
SG8 6DP
United Kingdom

www.paconsulting.com

SOUND EMISSION FROM THE INHALER

An airflow impingement will inevitably generate noise, and each small impingement or change in direction of the airflow, will contribute accumulatively to the overall acoustic signature. These flow impingements and subtle changes of direction are specific to the particular inhaler type. For example, swirl-based inhalers have a strong tonal (frequency) response to the flowrate.

Understanding the underlying physics of how the sound is generated influences the type of parameters and performance metrics that can be derived from the emitted sound data. In a swirl-dominated flow regime, common in many different dry powder inhalers, swirl number will be a function of flowrate, with higher swirl numbers producing higher dominant frequencies (Figure 1). Thus, by listening to the sound produced and analysing the frequency spectrum, we can infer the flowrate with high confidence.

DEEP LEARNING CAN DETECT PERCEPTUAL DIFFERENCES... & MORE

The wealth of information that is contained within the emitted sound can be used to infer various characteristics of the inhalation event, including flowrate, event actuation and even drug emission. These events often require the algorithm to detect subtle acoustic differences that – to the human ear – are inaudible. The advantage of deep learning is that it allows these acoustic signature patterns to be accurately detected and recognised, even beyond what the human ear and brain is capable of perceiving (Figure 2).

A second advantage of deep learning in this application is its robustness (98% accuracy even in initial studies), which is achieved by considering the full spectrum of the acoustic event. Through a large number of data, the algorithm automatically tunes in to a set of statistically relevant signatures that lead to the desired result. This is often referred to as automatic feature detection.

SMARTPHONES & DEEP LEARNING

Using sound recognition technology the microphone(s) contained within a smartphone can record a sound event – and the ever-increasing processing power of the smartphone means that as time goes on it is more able to perform deep

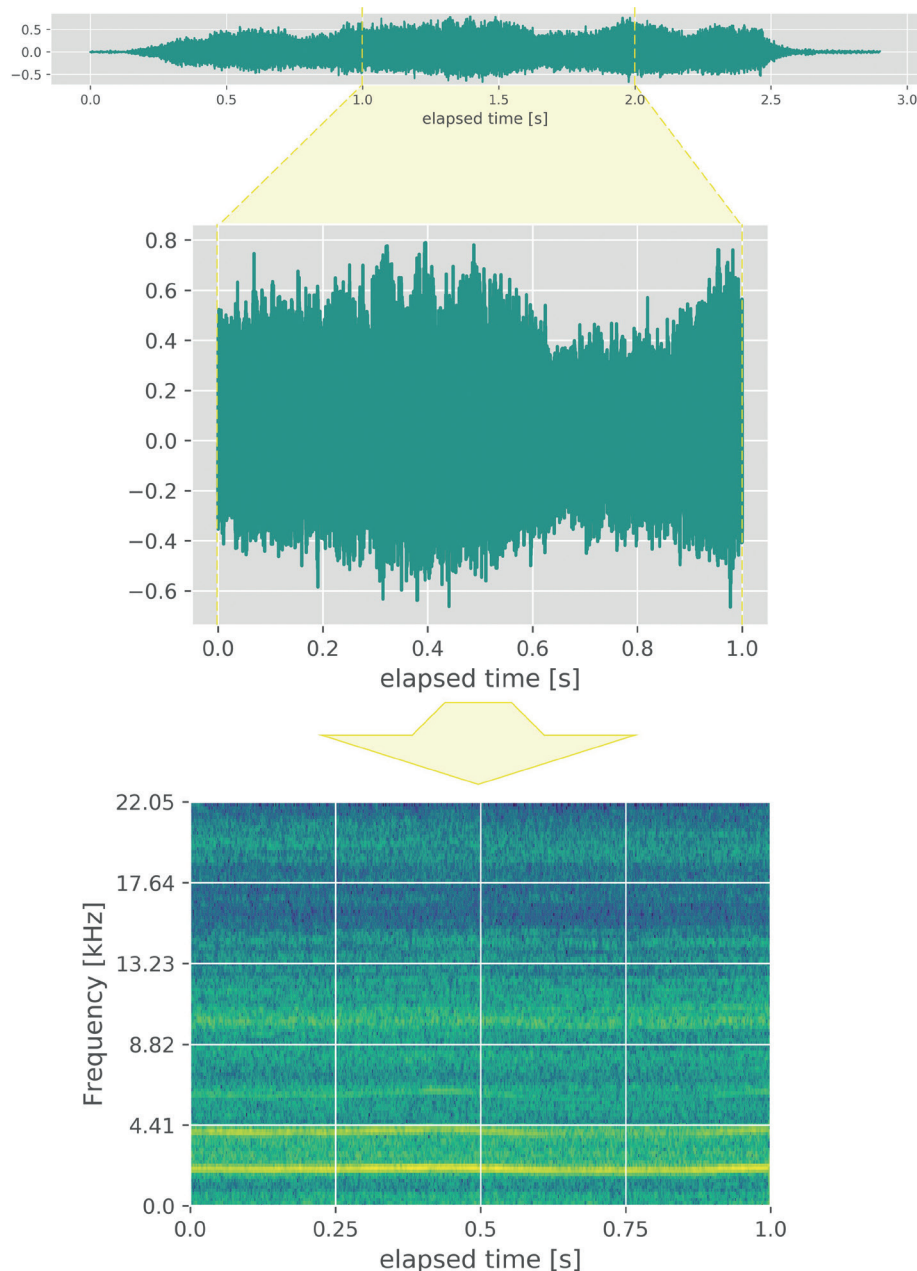


Figure 2: Audio sampling process to create spectrograms of the sound data.

learning analysis. The advancement in silicon and machine learning technology are the cornerstones of this concept becoming feasible now and their combination enables the implementation of inhaler sound recognition using a smartphone, which would not have been possible only a decade ago.

Deep Learning Explained

Deep learning acquired its name from its architecture, in which typical deep learning algorithms encompass multiple layers of neural networks. Each layer of the neural network is composed of multiple neurons, the design of which was inspired by the human brain cell neuron. Each of these neurons is tasked to identify a single subset

or decision. Combining multiple neurons into a layer produces an algorithm that can make more complex decisions. Deep learning goes a few steps further by joining up multiple layers of neurons and other mathematical transforms together to make increasingly complex predictions possible (Figure 3).

A complex neural network (CNN) that has multiple layers of neurons and mathematical transforms enables the deep learning algorithm to be trained to predict a complex action with high accuracy. A deeper network would often be capable of producing more complex predictions, providing the information gradient across high numbers of layers has been properly addressed.³

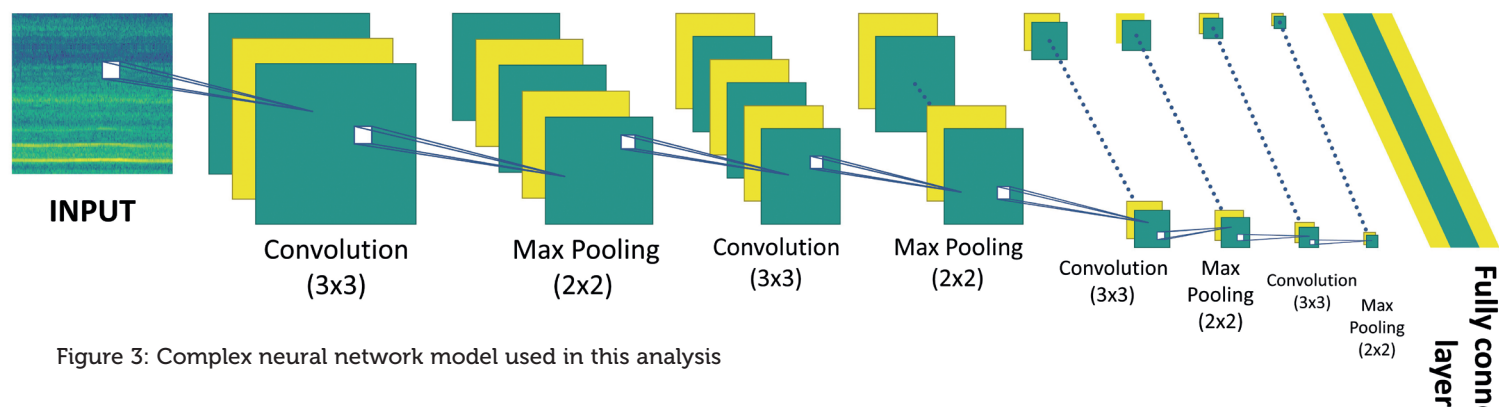


Figure 3: Complex neural network model used in this analysis

“Flowrate detection, for example, will be improved by hearing the sequence of flowrate changes, whereas intermittent events such as breath actuation or the opening and closing click of the inhaler, can be used to identify the beginning and end of the manoeuvre.”

The depth of deep learning layers is a double-edged sword because often a deeper network will require more samples and more data to be trained adequately. So it becomes a trade-off between the quantity of data versus the required level of confidence in the prediction.

IMPROVING ALGORITHM ACCURACY

Synthetic Data Handles Enormous Data Requirements

One shortcoming of using a deep learning algorithm is that it can require a large number of samples to achieve the statistical significance that leads to robust predictions. In fact, a deep learning algorithm trained using only small samples sizes (say hundreds to low thousands), is likely to have limited effectiveness and accuracy.

One method to address the necessity for large datasets is to augment the original data with additional reconditioned data. Data augmentation has frequently been used in the deep learning community to generate additional data through data reconditioning. For example, an image can be flipped to generate additional mirrored data for image recognition training.

Another method to increase the

quantity of data available to train the deep learning algorithm is to use synthetic data. These are data that have been generated through artificial synthesis based upon an underlying model. This technique requires understanding of the underlying algorithm and architecture. Both augmentation methods are useful to close the gap on the data requirement rapidly. Care must be taken with the method of augmentation and synthesis given that deep learning is essentially a pattern recognition algorithm. The augmentation and synthesis of data should be generated to represent the actual event in a way that is statistically relevant and acceptable. Nonetheless, when performed properly, this technique enables the production of a highly accurate and robust deep learning algorithm.

In the context of inhaler sound data, an example of data augmentation would be intentionally to superimpose various background noises onto a cleanly recorded inhaler sound. Background noises, such as a coffee shop or noisy canteen, will help the deep learning algorithm to be more robust in a real-world environment.

The robustness is achieved by feeding the deep learning training with additional “unexpected” data, containing an underlying actual inhaler flow noise. This will prepare the algorithm to anticipate these corruptions, and make the algorithm robust across different acoustic background environments.

Using Sequence to Improve Accuracy

The unique sequence of events that is observed during an inhalation manoeuvre can be used to reinforce the deep learning algorithm. Flowrate detection, for example, will be improved by hearing the sequence of flowrate changes,

whereas intermittent events such as breath actuation or the opening and closing click of the inhaler, can be used to identify the beginning and end of the manoeuvre.

ECONOMICS: SOUND COMPARED WITH CHIP-ON-INHALER

The elegance of simply listening to the sound is that it does not require any modifications to the inhaler whatsoever. Given that each type of inhaler produces a unique acoustic signature, which results from its physical design, this technology will be able to detect and identify that signature, and infer useful information about the way in which an inhaler was used through analysis of the sound data.

The primary advantages of this technique are that:

1. There is zero additional cost to the device
2. There are no modifications to the device

It effectively achieves connectivity for free. Thus, the general principle of using a microphone to listen to the sound of an inhaler opens up the benefits of connectivity in drug delivery to a wider market, specifically low income economy countries where adding a chip to an inhaler is simply too expensive. Using only a smartphone, with its built-in microphone as the primary sensor, provides previously impossible access to connected delivery device technology, potentially improving drug efficacy, adherence and inhalation technique.

“There is zero additional cost to the device. There are no modifications to the device. It effectively achieves connectivity for free.”

CLOSING REMARKS

The acoustic signature emitted from an inhaler provides a tremendous quantity of information. Flowrate and device actuation, for example, can be inferred through the deep learning algorithm.

The deep learning algorithm can robustly detect subtle differences in the

sound data, as well as producing an accurate prediction of flowrate profile based upon the sequence within the recording.

Combining the deep learning flowrate detection and the metadata that is readily available from the smartphone, a training program can be developed to improve patient inhalation technique, inhalation

effectiveness and adherence through a well-designed application that provides visual feedback that is appropriate for the target patient group.

Importantly, once the algorithm has been developed it adds zero additional cost, meaning the associated training program can be made widely available, even in developing countries.

ABOUT THE AUTHORS

Kian Min Lim is a Data Scientist and machine learning specialist at PA Consulting. He is passionate about innovation through cross-discipline technology implementation, such as developing continuous optimisation algorithms on pharmaceutical processes, implementing deep learning on consumer and medical devices and supporting medical device research using patient data. Dr Lim holds a PhD in Engineering from the University of Cambridge, UK.

David Harris leads the Respiratory Drug Delivery sector at PA Consulting. He is a physicist and has been working in the field of medical device development since 1994, where he started his career in the Respiratory Physics group at Fisons. He specialises in respiratory drug delivery and enjoys applying solid aerosol science and fluid dynamics to improve the efficacy of inhaler technology. Mr Harris has numerous patents and publications in this area and regularly presents at conferences.

REFERENCES

1. Gibney E, "Google AI algorithm masters ancient game of Go". *Nature News*, 2016, Vol 529(7587), p 445.
2. Lim KM, Lee SM, Harris DS, Seeney P, "Robust Characterization of Inhalation Information Using a Deep Neural Network and Smartphone". *Proceedings of Respiratory Drug Delivery 2018 Conference, April 22-26, 2018, Tucson, AZ, US*.
3. Szegedy C, "Going deeper with convolutions". *Proceedings of the IEEE conference on computer vision and pattern recognition, June 7-12, 2015, Boston, MA, US*.



Creating Space for Your Ideas

With the World's Smallest Flow Sensor

Sensirion consolidates incomparable accuracy, long-term stability and excellent repeatability in a tiny new device called the SDP3x - the world's smallest flow and differential pressure sensor. It's the perfect choice for cost-sensitive mass production, opening up new possibilities in integration and application.

www.sensirion.com/sdp3x

SENSIRION
THE SENSOR COMPANY