Pilot 4: COPD

1. Key Information

1.1 Involved Partners

- University of Southampton
- My MHealth Limited
- Astrazeneca UK Limited

1.2 Involved Countries

- United Kingdom

1.3 Keywords

- Risk prediction
- Predict acute exacerbations of COPD

1.4 Task Description

This pilot utilises data captured using mobile and web enabled platforms MY COPD and MY Asthma (MyMhealth data), which is used to develop predictive models of acute exacerbations of COPD. These models will enable a move from a reactive to proactive approach to care. The pilot utilises data captured on the platforms to create models using daily data on symptoms, treatment, and environmental observation data including temperature, humidity, pollen counts and air pollution to create risk models which are personalized to the patient's own disease state and environment.

2. Building Blocks

2.1 Architecture

2.1.1 System Architecture

The pilot inherits existing secure infrastructure as initially deployed by one of our partners (MYM):

Existing Infrastructure - currently used

- MYM Patient Data Store – Cassandra cluster running on AWS London stored on Encrypted volumes
- MYM Analytics Data Store – PostgreSQL cluster running on AWS London stored on Encrypted volumes with authentication and authorisation on database server.

Existing Infrastructure - used during first half of the project

- MYM Research Data Store – Mongo DB cluster running on AWS London stored on Encrypted volumes
- MYM BigMedilytics Study Data – Mongo DB cluster running on AWS London stored on Encrypted volumes with proper authentication and authorisation on database server.

-

Further data connectors to important data sources under the existing infrastructure at MYN are the following ones:

- UK atmospheric Air Quality data from BreezoMeter
- UK Pollen data from BreezoMeter
- UK meteorological data from DarkSky (https://darksky.net)

In the first half of the project the research work was carried out using environmental data from Copernicus but it was concluded that the resolution of the data is not sufficiently high.

Big Data Connectors / API's

Other Big Data Connectors with specifically secured API's are also in place. These include:

- Primary care EHR – FIHR integration from research sites (Yet to be confirmed on how this is to be integrated working with research partner in Glasgow UK)
- MYM Patient – Currently we have a secure manual transfer working towards exposing data through a user interface.
- Wearables / Smart Inhalers – We are integrating with devices to capture data and use open API to store them in MYM back end store.

MYM Core Applications built on AWS London instances which includes a mobile and web interface for patients, nurses and clinicians.

In addition to the above, a dedicated MYM data store will be securely deployed by MYM to experiment on new proof of concept applications on COPD exacerbation in patients at multiple scales in the Bigmedilytics project.

2.1.2 Data Flow & Interoperability of services

MYM m-health solution is deployed on Kubernetes utilizing services hosted behind an API gateway. These services are all documented using OpenAPI and can easily be exposed to third parties as required. OAuth 2.0 and OpenId connect is used for authorization and authentication of the services (JWE token). The use of Kubernetes makes it easy for us to add new services or scale existing ones to meet changing demands.

The MYM m-health platform is hosted as software as a service on AWS in London with three availability zones (data centres) to provide high availability (HA). The infrastructure is resilient with failure in one data centre not impacting on operations in another and can easily be scaled to meet demand by adding more nodes (machines) to the cluster. There is no hard limit in terms of the number of users we can handle. The following diagram depicts the MYM m-health platform information/data flow.

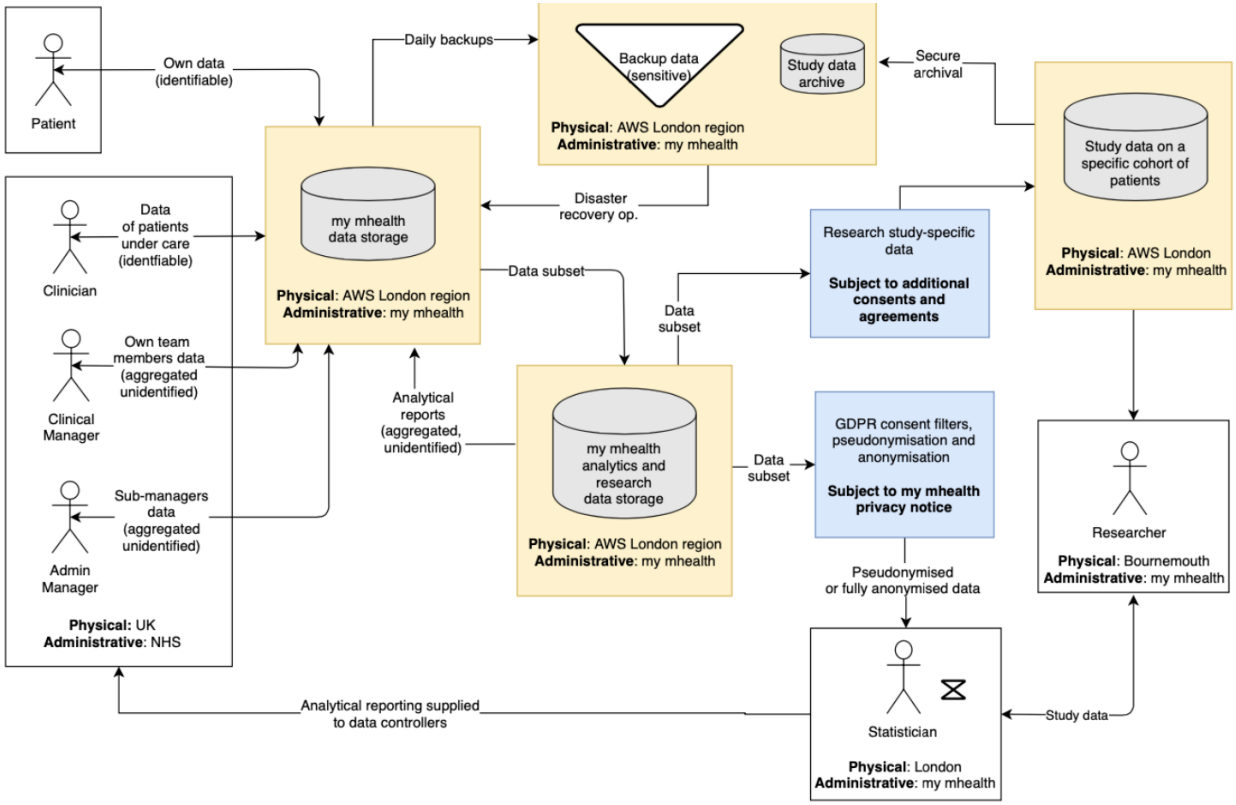

Patient data is collected is directly from patients using the service. This is entered via an individual account controlled by log in credentials chosen by the user. Clinicians are also able to add data such as observations and medicine changes following an appointment with the patient.

Healthcare professional data is collected by their commission group/trust and clinical manager accounts. Healthcare professionals are also provided with in an individual account as part of the clinical dashboard, accessed via their email address and their chosen password (aligning to my mhealth password policy).

Data is stored within Amazon Web Services London Regions only. A cloud service database cluster over 3 separate locations for fewer down time hours. Each region of our infrastructure is fully partitioned/isolated with availability zones (AZ), to better isolate any issues and achieve high availability. Each AZ (London) has its own power infrastructure and is connected with a fast, private fibre-optic network. Amazon Web Services London are made up of a cluster of Tier-4 connected data centres.

2.1.3 Necessary Hardware

The MYM m-health platform is hosted as software as a service on AWS

2.1.4 Software Components

The MYM m-health platform's components are virtualised using the Docker technology and the Docker components are managed by Kubernetes.

2.3 Data Processing

2.3.1 Processing of large structured / unstructured data sources

2.3.1.1 Data Sources

The aim in this pilot is to predict the risk of a patient experiencing a COPD/Asthma exacerbation event within a specified timeframe, based on the known influences of the event: a medical time series from the MyCOPD app, meteorological data and atmospheric pollution exposures at a patient's given location.

Data Source	**Description**	**Acquisition**	**Characteristic (Size, Patients, Years, Origin/Region)**

MYM patient data	The medical time series data will be provided by MYM (MYM patient data). This includes information about the symptom score and CAT score (patient inputs), as well as demographic features for each patient. The environmental data will be derived at the patient's given location and timestamp. Data has been segregated based on functional usage and secured techniques which involves authentication and authorisation for users accessing the data. Data has been segregated into three zones: 1. Application Zone – This is a MYM application data which is not allowed to access without proper consent from patients for any research purposes 2. Research Zone – Data extracted from application zone applying rules on consented patients, hence can be used for research purposes. 3. Study Zone – Study specific data which includes a limited set of agreed attributes and any filtering techniques specific to the study. Also includes an anonymisation or pseudo anonymisation process. Data extraction was performed on MYM data zone, applying filters on consented patients using the myCOPD app. The process is as follows: -anonymised extracted data, based on their user keys -extracted all agreed non-identifiable data in to BigMedilytics study zone. -apply pseudo anonymisation on patient postal code to help in extracting weather, pollution and pollen information for study purpose. -data transferred securely using an encrypted and password protected email only for intended recipients for modelling with our partner UNIS. -the data retention process will be agreed as per GDPR. -research team at MYM to work with Information Governance Team to have audit a governance on the data transfer process. E.g. Data Transfer Agreement.	Data is acquired via the MyCOPD app	This data is a stream of individual messages with daily to weekly frequency on average. The messages have a patient identified and a timestamp and a symptom score or a CAT score.
Atmospheric	The weather data covering the UK was retrieved from the European Centre for Medium-Range Weather	Research data is acquired	The data is available at the

meteorological and Pollution data	Forecasts (ECMWF). The ECMWF-Interim dataset(1) is based upon a large set of varying observations, from satellite data to land data. The results are based upon the results of the Integrated Forecasting System (IFS) model CY31R2 (2)..	via FTP. Data for MyCOPD app's operation is acquired via data connectors	spatial resolution of 0.1 degrees, and the latitude-longitude grid covers the whole of the UK. It includes 1. Atmospheric surface pressure 2. Atmospheric surface temperature 3. Total cloud cover 4. Surface dew point temperature. These values are recorded for the four non-overlapping 6-hour time periods each day within each grid section. The pollution data cover the UK was retrieved from the Copernicus Atmospheric Monitoring Service, via the European Air Quality service (3). Daily analysed data is

			available at a resolution of 0.1 degrees, with a measurement available at each hour in the day. This data includes 10 chemical species: 1. O3 (Ozone) 2. NO2 (Nitrogen Dioxide) 3. SO2 (Sulphur Dioxide) 4. CO (Carbon Monoxide) 5. PM10 (particulate matter < 10 μm in diameter) 6. PM2.5 (particulate matter < 2.5 μm in diameter) 7. NH3 (Ammonia) 8. NO (Nitrogen Monoxide) 9. NMVOC (Non-methane volatile organic compound) 10. PANs (Peroxyacytyl nitrates)

			For both datasets, we have obtained data from the dates 2016-04-01 to 2018-08-31; we expect to gather data as further medical data becomes available to the pilot.

multiple sources	**integration to data warehouse**	**data access**	**data stored in cloud**	**multi-party architecture**	**secure environment**	**transform raw / unstructured data**
yes	yes	Transfer/processing of de-identified data	yes	yes	yes	no

df2.3.1.2 De-Identification and anonymisation

All identifiable data has been removed from MYM data extracts.

- An anonymised patient id.
- Demographics (anon_id, created_date, gender, age, partial post code)
    - Age is used as opposed to the date of birth

Partial post code (first 3 chars)

2.3.1.3 Acquisition

The medical time series data will be provided by MYM. This includes information about the symptom score and CAT score (patient inputs), as well as demographic features for each patient.

The environmental data will be derived at patient's given location and timestamp. This data was previously retrieved from the Copernicus Atmosphere Monitoring Service but now MYM have implemented connectors to BreezoMeter, and DarkSky.

2.3.1.4 Cleansing

We will require that any patient used in the production of the prediction data-driven model(s) has the following requirements:
1. The patient has an associated location
2. They have at least one symptom score measurements.

As the timestamps in the medical time series are daily timestamps, when using meteorological data, it must be aggregated within the same time interval, as the data is usually given at smaller intervals, e.g. 6-hour intervals.
We also expect to eliminate anomalies by utilising known correlation between the symptom score and CAT score.

2.3.1.5 Data Integration

Data sources are being integrated at patient level using patient ID, timestamps and geo-location information (when analyses include environmental data). The eHRs have associated ID as well as every point in the time series. This allows us to construct a consistent record of medical events for each patient. The environmental information can be extracted based upon the patient's postal outcode of their home address, as well as the relevant timestamps associated with the respective measurements in the time series.

2.3.5 Complex real-time event detection

1. **Type of notifications**

Table 7. Types of notifications and alerts to be issued

Need for notification services			
Notification	**Warning**	**Alarm (automated / manual reaction)**	**Other**
Prediction of patient symptoms is performed. A 4-tiered score shows the severity of the COPD symptoms.	If the predicted patient symptoms are in the two top severe states (red and black), with a probability lower than the alarm threshold (currently set as 0.4)	If the predicted patient symptoms are in the two top severe states (red and black), with a probability higher than the alarm threshold (currently set as 0.4)	

2. **Type of situations of interest**

Table 8. Types of complex events to react on

Type of situations of interest			
Simple	**Trends (time-window / frequency based)**	**Complex (multiparameter / historical context)**	**Other**
-	-	If the predicted patient symptoms are in the two top severe states (red and black), with a probability higher than the alarm threshold	

3. **Type of event processing**

Table 9. Types of event processing actions

Filter	**Transform**	**Other**
-	The input patient symptoms are enriched with the weather, pollution and pollen data and then projected into the future using the learned models.	-

4. Event sources

Table 10. Event source during complex-event processing

Stream name	Contents of stream	Stream velocity	Description of the stream
Patient self-recorded symptoms	Symptom Score	Every time the patient inputs their symptom score; between daily and weekly.	This is a single score taking the value 1, 2, 3, 4 in order of the severity of the COPD symptoms. It is retrieved from the patient inputting their symptom score via the MyCOPD app. This is the primary stream.
Patient self-recorded detailed evaluation score	CAT Score	Every time the patient completes the questionnaire; between weekly and monthly.	This is a single score taking any integer value between 0 and 40, and representing the self-assessment of the patient's recent symptoms and wellbeing. It is retrieved from the patient completing the questionnaire via the MyCOPD app.
Weather and pollution data	weather, atmospheric pollution and pollen exposures	Variable (e.g. 6-hourly, daily)	For each patient symptom score or CAT score event, corresponding weather, atmospheric pollution and pollen exposures will be extracted from the data provided by Copernicus Atmosphere Monitoring Service and Met Office. These will be retrieved at patient's geospatial location, using the event time stamp.

2.3.5.1 Notifications

Push notification software providers to communicate medication reminders and updates from healthcare teams. This functionality is to assist the patient to ensure adherence to their medication plans and for clinicians to communicate via the in-app functions. Also, SMS messaging services are used for communicating with patients communicating information relevant to their condition.

2.3.5.2 Situations of Interest

A situation of interest arises when a patient sends a symptom score of 3 or above or an exacerbation event is predicted by the ML prediction algorithm.

2.3.5.3 Event Processing

Two types of events are sent by the MyCOPD app to the backend MyCOPD platform services: CAT scores and symptom scores. CAT scores are stored in the platform for subsequent use by the clinicians or for use by the ML algorithm when performing exacerbation event prediction.

Symptom scores are stored in the platform but also are processed by the ML algorithm for predicting an exacerbation event.

2.3.5.4 Event Sources

Events are generated by the MyCOPD app while interacting with the patient as an app user.

2.3.5.5 Evaluation

2.4 AI Components

2.4.2 Prediction Algorithms

2.4.2.1 Task

1. Predict acute exacerbations of COPD

The primary aim of the COPD/Asthma pilot is to understand and be able to predict the likelihood of COPD sufferers encountering an exacerbation of their symptoms, hereafter referred to as an “exacerbation event”, based upon the known correlations between exacerbation events and smoking status as a social determinant. In data analysis terms, this task translates to predicting exacerbation events and further understanding the reasoning behind that prediction. To this end, the pilot shall utilise two data types: MYM patient data generated from the MyCOPD app, and CAT score eHR data to model and predict exacerbation events. We would like to note that in the first half of the project we attempted linking environmental data and modelling it as an environmental determinant of COPD exacerbation events. This activity encountered the following main challenges. First, the location information from the MyCOPD app is not consistently available and also the exceserbation can be reported with a delay and the recorded location information might not match the location of the actual exacerbation event. The second and more important challenge is that COPD events are influenced by the micro-level environment while patients are outdoors or indoors and this very high resolution for environmental data is not currently available.

The analysis to produce the model(s) comes in the following steps:

1. Pre-processing of the data to ensure it is clean and consistent: investigating for errors and outliers in the values of the different attributes in the datasets, analysis of the level of missing data and selecting an ML approach that appropriately deal with missing data; selecting appropriate encoding of nominal and ordinal variables.
2. Linking of the MyCOPD app and CAT score eHR data
3. Building the model based upon the cleaned, linked data: dataset is split into training and testing dataset; model hyper-paprameters selection and fitting is based on k-folds cross validation using the training data and then the accuracy is validated on the testing dataset.

We shall now describe the process in more detail below.

2.4.2.2 Data, Data Modelling

Pre-processing

MYM patient data

The MYM patient information, symptom scores and CAT scores CSV files went through the following pre-processing procedures:

1. Ensured that each column have entries that are relevant to the variable name, in particular:
    a. Checking that categorical variables in the CAT and symptom scores take correct values.
    b. Checking that the dates are in the ISO 8601 format and, if not, converting to this standard.
2. Standardised the missing data indicator.
3. Manually processed the postcode parts in patient information to remove invalid codes and cleared unnecessary whitespace.

Environmental Exposures data (studied and used in the first half of the project but dropped due to the above challenges)

The meteorological data has been checked for consistency: ensuring that all variables have suitable values.

The pollution data was cleaned in the following way:

1. Missing data has been appropriately labelled.
2. Aggregating the hourly measurements into 6 hourly measurements, computing the mean and standard deviation of the measurements respectively with the 6 hour time slices matched with the weather data.

Linking/integration

The patient information, symptom scores and CAT scores CSV files contain anonymised identifiers, which allow for the linking between these files. The linking of the weather and pollution data to the other information occurs through the postcode part contained within the patient information CSV file. However, this level of linking was found to be insufficient during the COPD modelling work. The procedure for this linking is the following:

1. The patient information file links the anonymised identifier to their postcode part.
2. The postcode part is linked to a set of latitude and longitude coordinates. To simplify the process, we take the geometric centre of the set.
3. The single latitude and longitude coordinate is then matched with a grid section in the weather and pollution data.
4. The anonymised identifier is then attached to the weather and pollution data time series and outputted to a separate CSV file.

2.4.2.3 Features

- CAT score info (CAT scores range from 0-50):
    - Average CAT score
    - Last 1 score
- Symptom score info (Symptom scores range from 1-4):
    - Average symptom score
    - Last 1 score
- Exacerbation count
- Smoking status:
    - How many years the patient smoked
    - Whether the patient is an ex-smoker, non-smoker, or a smoker
- Demographic information:
    - Gender
    - Age

-

2.4.2.4 Model

There are two branches of analysis on this data: time series classification analysis, and predictive statistical modelling. The goal of these analyses is to construct a predictor that indicates when a patient will submit a symptom score equal to three or four. This, by definition, is an exacerbation event, as patients are advised to only select a symptom score of three when they have had to use their COPD rescue pack. We describe both analyses below.

Time series classification

To formulate the problem of detecting exacerbation events as a time series classification problem, we follow a sliding window analysis to produce suitable training and testing sets (1), and transform the data into a suitable dataset for supervised machine learning. The process for producing the model is given in Figure 2. The overall process for producing an exacerbation event predictor model from the MyCOPD patient, meteorological and pollution exposure data

To begin the process, we first select three characteristic times for the windows: the extraction time, the prediction time, and the shift time in days. For this analysis, we have chosen the extraction time to be 14 days, the prediction to be 7 days, and the shift time to be 7 days. For any particular date, we shall refer to the 14-day time period prior to the date as the “extraction period”, the 7-day time period after the date as the “prediction period”, and the time period relating to the shifting of the window as the “shift period”. For a visual representation of these periods, see Figure 1. The time series feature extraction process visualised

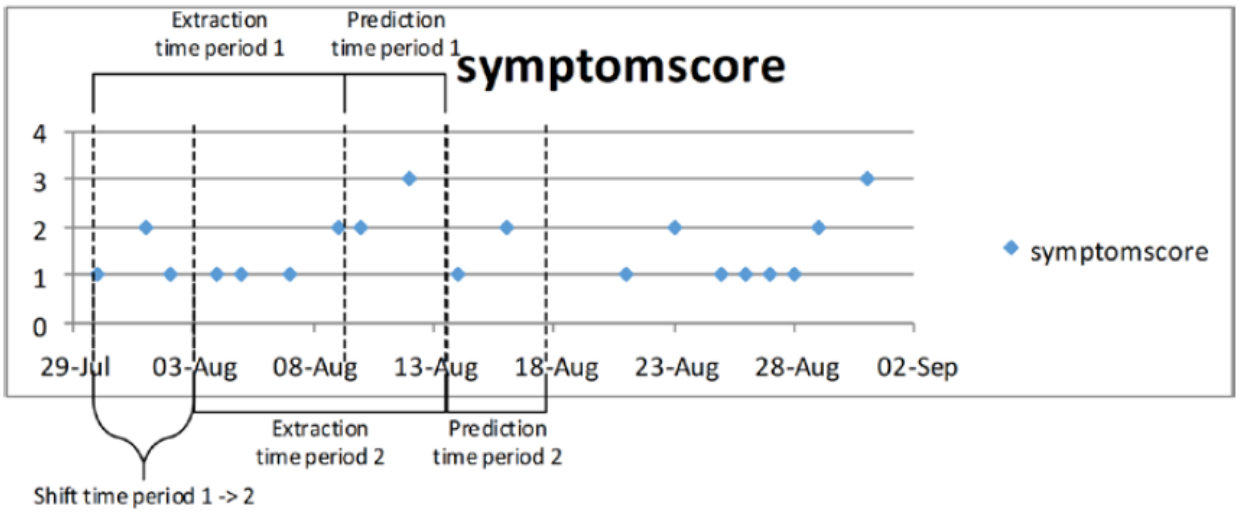

Figure 1: Time windowing for extracting features from the symptom score time series

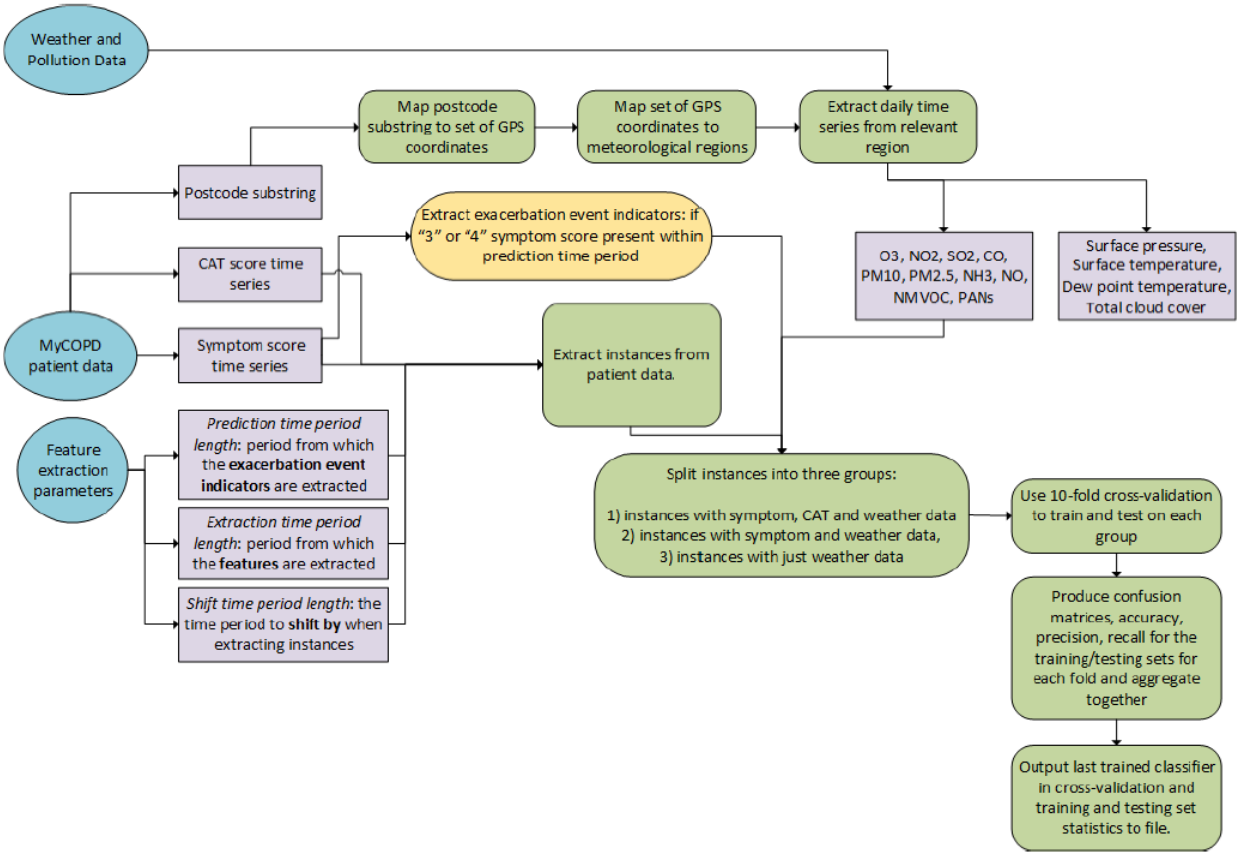

Figure 2: Data flow for training and validating the predictive model

For each patient, we start from the date equal to the first date that contains any medical data, plus the extraction time. Relative to this date, we then extract time series features from the data contained within the extraction period. In the case of symptom scores and CAT scores, we extract some basic ordinal time series features representing the temporal evolution within the

extraction period. Some examples of such features are: the number of times certain categories were recorded, the most recent entry, as well as more complex features based on the ordinal nature of the time series. In our initial work, from the weather and pollution data, we computed a standard range of features, including the average, the minimum, the maximum, as well as the maximum change (as well as others contained in the *tsfresh* Python package. These values are all placed into a feature vector. We then move onto the prediction period, which we use to label the resulting feature vector. The label in this case is a binary label, where a 1 indicates that a 3 or 4 category record is present in the prediction period and 0 otherwise.

The above procedure is then iterated over all patients and all records, with the window shifted by the shift time. The result of this procedure is the production of a set of feature vectors, all of which are labelled with a binary label, which is suitable for supervised machine learning. In our initial work this dataset was then used to train a number of shallow machine learning classifiers, such as random forest and support vector machines, using 10-fold cross-validation to select the best overall classifier. In our work in the second half of the project we evaluated and selected for use the XGBoost decision-trees based ensemble learning algorithm which produced better accuracy and is also robust to missing data.

Statistical modelling

A patient's lung function, as affected by COPD and asthma conditions, can be considered a stochastic process which responds to outside stimuli depending on its inner state. Given this, a generic lung function can be modelled using a dataset of patient symptom scores and the hypothesized drivers.

Statistical models for sequential data were chosen for this task. The outside derivers were considered of environmental nature such as weather, pollution and pollen exposures. The specific parameters from each category were the same as what was discussed in the previous section.

Here, the prediction of exacerbation is tied directly to the internal states of the models and the outputs at these states. In this, the latency and speed in which the drivers affect the internal states of the models are not explicitly defined but they are partially learned by the model and are partially guided by the structure of the model.

The predictions for the state of each patient can be made at each time step of the model, given forecasts of the inputs or drivers (weather, pollution and pollen) of the model. The time step is defined at the modelling stage this will most likely be a daily time step. These results can be aggregated to produce exacerbation likelihoods in a given time frame.

In respect to this work we would like to note that we experienced the technical and phenomenological challenges for utilising the environmental data as discussed earlier, namely reliable geo-lincing and need for very high resolution of the environmental data.

2.4.2.5 Evaluation

2.4.3 Data Analytics

2.5 Security and privacy of data access and processing

2.5.1 Access Control

2.5.1.1 Authentication

Individual accounts are controlled by log in credentials chosen by the user. This is currently a single factor authentication with an email address and password.

Secure cookies and HTTP-only cookies are enforced in HTTP communications. Authentication cookies are encrypted and salted. Passwords are hashed utilising PBKDF2. Tokens sent to users expire in 3 hours or when utilised a single time.

2.5.1.2 Authorization

Patients are able to access their own data. Clinicians are able to access data of patients under their direct care. Clinical Managers and Top Level are able to access anonymised aggregated data and also data input about the clinicians.

2.5.2 Data Protection

2.5.2.1 Data at rest

Data is not stored outside of the UK boundaries. Data transferred to AWS is encrypted at rest and AWS have a series of recognised international standards such as ISO 27001.

2.5.2.2 Data in transit

Restriction to TLS v1.2 only, using updated, secure ciphers (AES 256 where possible). Known insecure protocols, ciphers and configurations are disabled, e.g. RC4, SSL3, non-perfect-forward secrecy, client re-negotiation.

Ciphers utilised for data in transit are:
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_WITH_A ES_256_GCM_SHA384
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384

2.5.3 Auditory and logs

2.5.3.1 System Auditory

System auditory activities include:

- internal and external security audits
- software quality assurance process
- application security updates
- policies on network configuration
- security advisory reviews covering full stack software components
- IT staff training on security

2.5.3.2 Services Auditory

Access is logged in the database. Entry length of time and activity and the database is backed up to an encrypted back up provider - AWS.

2.5.4 Privacy measurements

2.5.4.1 Data Privacy Impact Assessment (DPIA)

The development of the machine-learning model, as well as other research activities associated with the responses to and adoption of the big data analytics components, was covered by a DPIA, and reviewed both by my mhealth as clinical lead as well as the University of Southampton DPIA panel.

2.5.4.2 Legal/Ethical process

As well as the DPIA, research activities were reviewed and approved by the Faculty of Engineering and Physical Science (where research work was located). As data custodians, the my mhealth database was covered by the HRA research database exemption[1]. A data sharing agreement was agreed between my mhealth and the University of Southampton partners covering the specific data fields used and the constraints under which the data could be used (i.e., limited to the direct use of the BML project).

---

[1] https://www.hra.nhs.uk/planning-and-improving-research/policies-standards-legislation/research-tissue-banks-and-research-databases/

During ethics review, discussions were initiated in regard to the status of data subjects for the pilot. Typically, app users who download and use the myCOPD app would do so as patients of the NHS. As such, HRA oversight would be required to process and exploit the data. However, since for the purposes of the trial itself – notwithstanding any future exploitation plans – the app users were effectively the same as any voluntary app users, rather than patients involved in a clinical trial. This affects the governance landscape. As part of a follow-on, where alerts would become part of a feasibility trial, then their status would change, not least in respect of how alerts are routed. Nevertheless, careful consideration should be given, therefore, in similar projects as to the practical status of trial participants.

Recommendations: look for ways to conceptualise different parts of a project in different was as they relate to different research and regulatory frameworks.

2.5.4.3 Processes for complying with the current legislation

Since all of the activities for Pilot 4 have been run and managed in the UK, the overarching governance has been informed by the UK Data Protection Act (2018), which was the UK implementation of the GDPR (2016) prior to the UK leaving the EU as well as after. As the COVID-19 pandemic spread, additional regulatory instruments were relevant, including the COPI regulations. These were relevant and important for the clinical partners, but had limited effect on the University of Southampton data scientists.

With that in mind, activities in the project were visible to the University of Southampton DPO (via the DPIA process) and the my mhealth DPO in addition to the BML project DPO. Specifically for Pilot 4, the University of Southampton partner included the chair of the Faculty Research Ethics Committee, who is also a member of the DPIA panel.

This level of expertise is not necessarily required, we believe, however given the dynamic status of the project in response to external factors (like the UK leaving the EU, and the unprecedented pandemic) and internal factors (app users become potential patients etc.), it was useful to be able to involve these colleagues in discussions.

Recommendation: consider project-wide engagement of both specialist ethicists as well as a DPO

2.6 Trustworthy AI

The COPD Pilot has been designed and executed with reference to a number of different instruments and guidance documents, including:

- The Toronto Declaration on the use of (personal) data[2]
- The UK Government Digital Services[3]
- The EU (HLEG) Trustworthy AI guidelines[4]
- The UK Government Health and Social Care Code of Conduct[5]
- The CARE Principles[6]

Although these statements share similar principles (fairness, transparency, accountability, explainability, etc), it was important to ensure that the research and innovation activities complied with the essence of the different perspectives. In regard to the CARE Principles, originally defined in regard to indigenous populations involved in research, the focus of this work was the direct link back to the community most affected by the condition, namely COPD sufferers. Their engagement would therefore be repaid in concrete benefit within the self-reporting app myCOPD; would benefit other apps under development by my mhealth; and potentially would also contribute to local, regional and international databases for respiratory.

Recommendation: big data healthcare projects should continually monitor relevant guidelines and initiatives.

2.6.1 Technology/user adoption and establishing trust

Pilot 4 partners have experience in causal models for technology adoption, as well as models and frameworks for the adoption of healthcare interventions. It was clear, therefore, that specific issues (such as risk awareness, likelihood, and self-efficacy[7]) would need to be brought to the fore in communication with the app-user cohort and understanding their responses to the technology; whilst issues of complexity and the broader deployment context (e.g., coherence and cognitive participation[8]) would be major indicators with the clinicians. Although the COVID-19 pandemic forced a significant change in the evaluation phase of the work (equally because of patient / app-user vulnerability as sufferers of severe respiratory disorders and because of respiratory clinician reassignment to critical care duties), we have been able to

[2] Amnesty International and AccessNow, The Toronto Declaration: Protecting the right to equality and non-discrimination in machine learning systems. 2018.
[3] UK Government, Data Ethics Framework, Government Digital Services, Editor. 2020
[4] European Commission (HLEG on AI). Ethics guidelines for trustworthy AI. 2019; Available from: https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai.
[5] UK Government, A guide to good practice for digital and data-driven health technologies, Department of Health & Social Care, Editor. 2021.
[6] https://www.gida-global.org/care
[7] These are constructs from the Health Belief Model (and Prevention Motivation Theory); these function as the external variables of traditional technology adoption models (such as TAM).
[8] From Normalisation Process Theory

engage remotely both with clinicians and app users to validate their responses in the context of our focus on the potential for a big-data-enabled solution[9].

Recommendations: especially for healthcare technologies, the healthcare context (as it applies to different stakeholders) is more important than specifics of the technology as predicted by traditional technology acceptance causal models.

2.6.2 Ethical principles

As part of the relevance both of PPI (Patient and Public Involvement) and establishing trustworthiness as a team *vis-à-vis* app users / research participants, we have positioned both app-users and supporting clinicians centre-stage in the thinking around basic ethical principles, including:

- respect for human autonomy ~~authority~~

  The whole aim of Pilot 4 is to empower patients to engage with and become involved with their own care regime: it gives them a level of independence and feeling of self-efficacy which would ultimately benefit their long-term prognosis. As well as the self-reports, it allows them increase their knowledge of their own condition, even 'ask the questions they're scared to ask'.

  In addition, through modelling indicators within the self-reports of app-users, this enables clinicians to plan better (potentially triage patient lists) and to manage the sheer volume of data available; further, it provides an overview of patient well-being over time to supplement (and make more effective) consultations: it allows them to improve their own clinical practice in interaction with their patients.

- ~~prevention of harm~~ benevolence / non-malevolence

  The introduction of the data analytics is intended in the first instance to predict exacerbation events in COPD patients. These events are known to contribute to the deterioration in the well-being of the patient. Being able to make such predictions would not only empower the patient with a feeling of greater control over their condition, but also improve their quality of life longer term. Pilot 4 is therefore very much about patient-directed benevolence (improving the course of the illness) and non-malevolence (avoiding negative health outcomes).

- fairness

[9] Note: much of the thinking here has been submitted to an MDPI Future Internet Special Edition, acknowledging the contribution of BML (titled: Trust, but Verify: Informed Consent, AI Technologies, and Public Health Emergencies)

The patient takes control of what and when they report. There is no bias towards any particular patient cohort. The big data model itself will raise an alert based on the data irrespective of any other information (such as an identifier) of individual patients. This lack of bias and equal access to all emphasises the fairness associated with the project.

- explicability

  As far as the myCOPD app itself is concerned, the self-reports are intuitive and easy to manage (based on a simple 4-choice visual presentation); supporting educational materials about the condition are also well-received.

  As far as the big data component is concerned, Pilot 4 has progressed over repeated cycles of data scientist exploration of the data, based on initial intuitions by the clinicians, and then discussions with the clinicians to make sense of the findings. There is a plan to formalise such iterations to try to understand how complex data analytics can be rendered 'understandable' to experts in other domains (such as clinicians, in this case).

-

2.6.3 Key requirements

- Human agency and oversight

  The design and implementation of the myCOPD app and integrated ML-component is very much dependent on a human-in-the-loop approach: in particular, no automatic decisions are made on the basis of the outputs from the ML-component. Instead, it is the clinician who currently receives the alert and decides what to do with it.

- Technical Robustness and safety

  The app itself must comply with MHRA standards.

- Privacy and data governance

  All data is transferred via secure connection; the health database is encrypted at rest; and any data exchange between partners is handled via a secure, partner owned version of "dropbox", whereby the data are encrypted in transit and at rest.

- Transparency

  The function of the myCOPD app is entirely transparent to all actors (app-users/patients as well as clinicians); as mentioned above, we continue to work directly with clinicians to establish a communication mechanisms about the ML-components to ensure some degree of transparency at that level too.

- Diversity, non-discrimination and fairness

  There are no constraints on who can interact with the myCOPD app, which has been designed with accessibility at its heart (https://mymhealth.com/accessibility); most information is provided visually to support ease of use.

- Societal and environmental well-being

  As evidenced during the COVID-19 pandemic, the myCOPD app supports virtual, remote interaction. In predicting exacerbation events, the aim is to reduce hospitalisations and thus the environmental effects of multiple journeys and better use of hospital resource.

- Accountability

  All operations are traceable and auditable within the overall solution. The app development, support and deployment is subject to review by appropriate authorities (e.g., MHRA).

-

2.7 System-Interaction

2.7.1 Human-Machine Interface / GUI

The MyCOPD condition management solution provides interfaces for the patient through the MyCOPD app and for the clinicians through the clinical dashboard.

The functionalities provided by the MyCOPD app are as follows:

- Easy-to-follow educational videos to learn how to manage their condition
- Complete online education such as pulmonary rehabilitation courses
- Reports can be generated to show changes in symptoms over a period of time
- Weather and pollution forecasting - Receive an accurate forecast daily to understand how the weather and air pollution in local areas can impact health. Plan the day with confidence
- Notifications to inform patients of medication reminders, to advise of any changes made by their clinician or if their clinician has sent them a message.
- Self-management plan and diary- Know when, and how to take your medication with the online, self-management plan. The person can also record when they have taken their treatment in the medication diary. This is real time user contributed data that can be viewed in the clinical portal.
- Upload information / photos to support shared decision making e.g. diabetes eyes, kidney and foot care

The functionalities provided by the clinical dashboard are as follows:

- The clinical dashboard enables clinicians to deliver self-management, education, inhaler technique training and education courses e.g., pulmonary rehabilitation course on any smartphone or tablet. Each intervention has been shown to deliver the same outcomes as access to a face-to-face education e.g., rehabilitation class and correct 98% of inhaler errors and enables you to manage your patients like never before.
- Real-time patient symptom tracking
- View prescriptions against national guidelines, check medication conflicts and
- assess overall monthly cost of prescriptions.
- The videos e.g. inhaler videos can be used to update own education, or use the video button to deliver education to the user at their community or clinic visit.

2.7.2 Education

3. Learnings

3.1 Challenges & Barriers

Regarding the technical aspects of the innovation, the main challenge comes from the fact that the local environment is a determinant of the COPD exacerbation events. In the first part of the project we attempted linking environmental data and modelling it as an environmental determinant of COPD exacerbation events. This activity encountered the following main challenges. First, the location information from the MyCOPD app is not consistently available and also the exceserbation can be reported with a delay and the recorded location information might not match the location of the actual exacerbation event. The second and more important challenge is that COPD events are influenced by the micro-level environment while patients are outdoors or indoors and this very high resolution for environmental data is not currently available.

In respect to the technology acceptance aspect of our work, as previously stated (see section 2.6 above), Pilot 4 has continued to monitor activities by various groups (Amnesty International, the EU HLEG for AI, various UK Government departments) to ensure a dynamic understanding and approach to concerns and challenges for Trustworthy AI. This has also been formalised in a research paper[9].

Main challenges:

Engagement of all relevant stakeholders; it's essential to understand who the main 'players' are (i.e., the actors in the network) to be able to understand what service delivery is really about (i.e., what their respective expectations are). This may lead to *tussles* (in the Clark et al (2005) sense), but it's better to understand these up front.

Communication and understanding between data scientists and clinicians; these two significant actors are from different disciplines and have different priorities day-to-day. Time taken to understand and facilitate cross-disciplinary communication is well worth it.

AI-enabled technology which is disruptive within a well-established human-to-human (trust-based) interaction; guidance and suggestions from, e.g., the Toronto Declaration, and HLEG AI's Trustworthy AI, need to be mapped onto the plans for a given healthcare / data analytics project.

Demonstrating and Maintaining trust across the whole eco-system. Healthcare is largely dependent on a trust relationship between patient and clinician. As demonstrated with contact tracing during the pandemic (cf. Rowe et al., 2020), the socio-political context can have very negative effects and these need to be understood and managed.

Main barriers:

Ethics is often too readily assumed to be a sub-bullet for data protection regulation: privacy is not the only factor of importance to patients – as other studies during the pandemic have discovered. So, *doing the right thing* (which is more what ethics is about), involves understanding patient needs and expectations.

Euro- (individualist) centric focus on individual rights versus community benefit: data sharing in general is especially difficult in healthcare. However, for the full potential from big data to be realised, there needs to be a concerted effort to explore and emphasise ways in which datasets can be shared for the benefit of all stakeholders and with their understanding (cf. the South African construct: *ubuntu*).

Regulation is often too onerous (i.e., misunderstood or causes a nervousness about innovation and data sharing): (1) the COVID-19 pandemic has highlighted that there needs to be a quick response in such times to enable not hinder appropriate research (cf. the UK COPI Regulations); (2) Consent is not the only (or even most appropriate) basis for personal data processing.

-

3.2 Lessons Learned

Predicting and measuring outcomes in a traditional sense is not always possible: *a priori* KPI setting can be disruptive. Projects, especially those involving direct engagement with humans (!), may need to revise metrics (including dropping some and introducing others) as the project evolves. A straightforward process for this needs to be devised.

Different actors speak different languages: it's been helpful to step back and re-visit our understanding of "trust". This is not just a reaction to trustworthiness; *and* needs to be separated from reliance (where agency is not present) and real trust. For AI-enabled technologies (cf. the Turing test), where human actors' propensity to anthropomorphise "intelligent" machines, the first principle from the HLEG AI's guidelines (Human In The Loop) is particularly important: (a) to avoid unreasonable and inappropriate expectation; and (b) to 'train'

Clinical expertise and limitations: there needs to be a careful, and independent 'brainstorming' session with clinicians to understand what their jobs (job families) is *really* about – they are used to providing an authoritative voice within a high-stakes interaction; it's not clear that an exploratory ML-model fits completely within this paradigm.

Medical ethics: whatever the approach (Beauchamp & Childress 4 Principles *versus* Boyd et al), it's time to take stock of expectations and the realities of big data and data analytics. This is a significant project in itself but is needed to provide a context within which to implement the HLEG AI type recommendations.

Ethics "theories": especially in research and innovation projects, from the outset it is worth considering existing ethical frameworks and which best meets the needs of project objectives. My experience is leaning more towards the Floridi "infosphere" approach. But again, this requires stakeholder buy-in **and** an appreciation that ethics is separate, albeit parallel, to data protection.

-
-

3.3 Main (quantifiable) achievements

The pilot partners are actively looking to scale up the developed innovation from the current local level to a UK national level. UoS IT Innovation and MyMHealth are leading the commercialisation and adoption of BigMedilytics' outcomes through the NHSX/NIHR Artificial Intelligence programme. A £1.5M proposal has been submitted building on the results of the BigMedilytics work on Pilot-4. Currently (April 2021) the proposal has been successfully reviewed and called for stage 2 presentation to the NIHR for final decision.

Our work on Pilot-4 demonstrated that machine learning can predict acute COPD exacerbation events several days in advance. Such predictions have huge potential to improve clinical models and drive behavioural change necessary to avoid acute exacerbation. The NHSX/NIHR proposal builds on the fact that we can potentially predict COPD exacerbation events for NHS patients (i.e. UK wide) and use them to trigger interventions through scaling up the current BigMedilytics outcomes in the areas of 1) safe and reliable AI, 2) acceptance and engagement with patients and clinicians, 3) establishment of clinical efficacy, safety, regulatory data, 4) establishment of a competitive AI-enhanced digital health ecosystem.

4. Output

4.1 Papers

“Trust, But Verify: Informed Consent, AI Technologies, and Public Health Emergencies” submitted to MDPI *Future Internet* Special Issue "Human–Computer Interaction Models and Experiences for Internet of Things Systems and Edge Computing" (https://www.mdpi.com/journal/futureinternet/special_issues/HCL_ME_IOT_EdgeComputing)

4.2 Open Source & Resources (refer to ELG)

4.3 Demos