



Enabling scalable clinical interpretation of machine learning (ML)-based phenotypes using real world data

Owen Parsons^{1*}, Nathan E. Barlow^{1*}, Janie Baxter¹, Karen Paraschin², Andrea Derix², Peter Hein², Robert Dürichen¹

¹Arcturis Data Ltd., Oxford, UK; ²Research and Development, Pharmaceuticals, Bayer AG, Wuppertal, Germany

Contributions: (I) Conception and design: All authors; (II) Administrative support: A Derix, R Dürichen; (III) Provision of study materials or patients: R Dürichen; (IV) Collection and assembly of data: O Parsons, NE Barlow; (V) Data analysis and interpretation: O Parsons, NE Barlow, J Baxter, K Paraschin, P Hein, R Dürichen; (VI) Manuscript writing: All authors; (VII) Final approval of manuscript: All authors.

Correspondence to: Robert Dürichen. Arcturis Data Ltd., Oxford, UK. Email: robert.durichen@arcturidata.com; Peter Hein. Research and Development, Pharmaceuticals, Bayer AG, Wuppertal, Germany. Email: peter.hein@bayer.com.

Background: Large and deep electronic health record (EHR) datasets have the potential to increase understanding of real-world patient journeys, and to identify subgroups of patients currently grouped with a common disease label but differing in outcomes and medical need. However, working with EHRs is still relatively new and challenging due to the heterogeneous nature of data. Increasing interest in machine learning (ML)-based EHR aggregation is mostly method-driven, i.e., building on available or newly developed methods. These methods, input requirements, and output are frequently difficult to interpret, especially without data science or statistical training. This endangers the ultimate aim of such analyses: generating actionable and clinically meaningful interpretation.

Methods: We conducted stratification and sub-phenotyping of cardiovascular patients from combination of NHS EHR datasets with encounters from March 2014 to August 2020 from Oxford University Hospital (OUH) and from February 2014 to March 2020 from Chelsea Hospital and Westminster Hospital (ChelWest). The dataset contained diagnoses, laboratory tests, medications, and procedures from 1,480 and 918 patients from OUH and ChelWest respectively. Different clustering and preprocessing resulted in more than 100 clinical reports which would require significant time to be clinically interpreted.

Results: We have developed a new framework facilitating clinical evaluation and interpretation of unsupervised patient stratification. For each new step within this framework we present example methods—pattern screening, meta clustering, surrogate modeling, and curation—which can be used at different stages within the analysis. Compared to a standard approach, we demonstrate the ability to condense results and optimize analysis time. For meta clustering, we show an example where the number of patient clusters are reduced from 72 to 3. With surrogate models, we quickly identified “blood sodium” as a stratification measure in heart failure patients, a likely identification of data bias as this is a routine measurement. By using cohort and feature curation, these and other irrelevant features were removed, increasing clinical meaningfulness.

Conclusions: This study investigates approaches to perform patient stratification analysis at scale using large EHR datasets and multiple clustering methods for clinical research. We show examples on the effectiveness of the methods and hope to encourage further research in this field.

Keywords: Patient stratification; electronic health records (EHRs); clinical evaluation; machine learning (ML); real-world data analysis

* Former member of this organisation.

^ ORCID: Owen Parsons, 0000-0003-0163-5609; Nathan E. Barlow, 0000-0001-7546-6890; Peter Hein, 0000-0002-5899-3979; Robert Dürichen, 0000-0002-4014-3941.

Received: 20 June 2022; Accepted: 04 January 2023; Published online: 15 February 2023.

doi: 10.21037/jmai-22-42

View this article at: <https://dx.doi.org/10.21037/jmai-22-42>

Introduction

Research in the life sciences has come to rely heavily on large scale digital data acquisition and analysis (1). Indeed, digitalization in health care, and specifically documentation of electronic health records (EHRs), is developing into a standard practice for many health care providers. The increased availability of large clinical datasets, combined with recent advances in machine learning (ML) methods, have led to an increasing number of studies [reviewed in ref. (2)]. Although this trend started over a decade ago, the available data is becoming more and more relevant as datasets are (I) more complete, (II) extended more longitudinally, and (III) more horizontally integrated with increased linkage to other relevant datasets from the same patients like laboratory values, imaging, and other diagnostic procedures. These advances have underpinned the growing interest in the application of this large-scale real-world data not only for epidemiological purposes, but also to understand trajectories of patient subgroups within large but heterogenous diseases like heart failure (HF),

chronic kidney disease, or stroke. As the trend towards more complete EHR dataset is accompanied by development of analysis methods (e.g., ML), this approach is becoming more likely to reveal relevant and actionable insights that have the power to benefit patients directly (e.g., through better care and precision medicine) and indirectly (e.g., by facilitating development of tailored drugs) (3-6).

One high impact area of research is patient stratification where, amongst others unsupervised learning methods, patients that share a similar clinical history (e.g., similar comorbidities) are clustered into sub-phenotypes to support disease understanding and facilitate more targeted treatment options (7-13). This has been applied to problems such as identifying subgroups of intensive care patients with common clinical needs (14,15) as well as finding subgroups of patients that have distinct responses to a fixed treatment (16). Generally, there are two core stages to the process (*Figure 1*). The first is the identification of patient subgroups using data-driven methods. The second is the clinical evaluation and interpretation of these subgroups using statistical methods. New studies in this area often focus on the development and application of novel clustering methodologies (11,17,18); however exploration of methods to facilitate and accelerate the clinical evaluation are not considered to the same extent. This development results in an increasing number of methodological choices and often it is impossible to determine initially which approach will lead to the most insightful outcome. Consequently, multiple approaches are applied which results in an increased number of potentially relevant outcomes to be evaluated by clinical experts (see *Figure 1*). As dataset sizes and number of model experiments increase, there is a growing need for novel methods that specifically support interpreting results from complex studies, where many parallel approaches are applied. Note, within the context of this publication, an experiment is defined as applying a specific stratification algorithm on data of a patient cohort including specific preprocessing steps and algorithms parameters which results in a mapping of patients to different clusters.

In this publication we present a new framework how to conduct large-scale, ML driven clinical analyses of unsupervised patient stratification and demonstrate example methods for each new step within the framework. At a high-

Highlight box

Key findings

- This study presents and discusses tools facilitating clinical review of electronic health record (EHR) analyses.

What is known and what is new?

- Data heterogeneity and the amount and complexity of result output by machine learning (ML)-based analyses of EHR require significant efforts by clinicians for review and medical interpretation;
- We report a framework with novel workflows: pattern screening, meta clustering, surrogate modeling, and curation. These methods facilitate scalable review, iterating, and clinical interpretation of such analyses.

What is the implication, and what should change now?

- The framework presented here is an approach helping clinicians who are often not trained in data science to interpret EHR analyses and derive clinically meaningful insights;
- Further research should not only focus on ML method development but also implement approaches facilitating further review and interpretation.

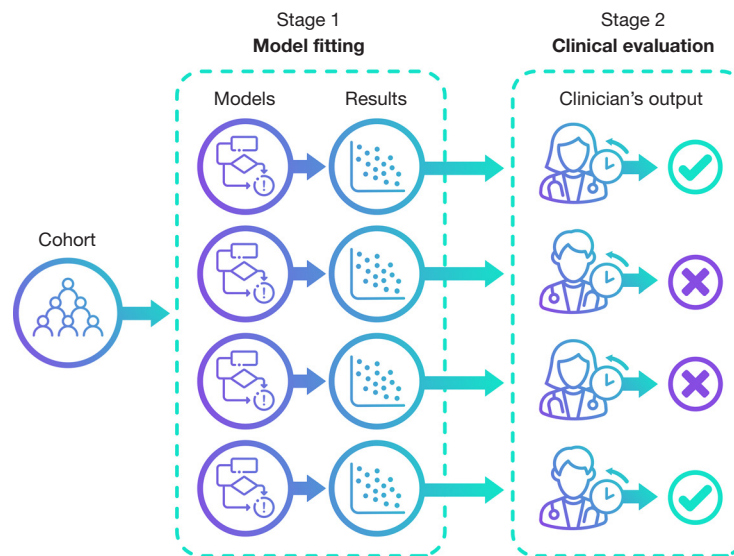


Figure 1 General workflow of patient stratification where multiple models are fit based on a single patient cohort. The outcomes are validated by a domain expert in parallel for each model experiment, which is highly labour-intensive.

level, these challenges break down into managing large volumes of evaluation results that need to be interpreted by clinicians, facilitating the extraction of insights in studies with large number of observations and support fast iterations of results to increase clinical relevance. The solution we propose (*Figure 2*) extends the clinical evaluation process with the addition of a key results identification stage (stage 2), an explainability stage (stage 3), and an optimisation loop (stage 5). Note, even though not every method presented is new, we focus in this publication on how the discovery process using ML methods and large scale EHR datasets can be improved by reducing the burden on clinical domain experts. We hope that this initiates further discussion within the community about the development of more appropriate methods.

Methods

We present the conceptual approach and methodological details of our proposed new scalable clinical evaluation framework including methods for new stages. Additionally, we provide details of a use case study, along with the objectives of the study and the analytical approaches taken.

Proposed scalable clinical interpretation approach

In this publication we propose to extend the current two-step workflow (see *Figure 1*) by three additional stages with

the objectives of (I) reducing the number of results a clinical researcher needs to investigate, (II) increasing explainability of individual results, and (III) to quickly iterate and optimize results. Importantly, we suggest that this extension makes the clinical interpretation scalable. Our proposed workflow is shown in *Figure 2*. Note, the original two stages presented in *Figure 1* remain the same [model fitting (stage 1) and clinical evaluation (now stage 4)]. However, with the introduction of the Identification, Explainability and Optimisation stages the overall time required for a clinical researcher to review results will be reduced.

Stage 2—Identification

The main objective of this stage is to reduce the number of available results generated in stage 1. As motivated in the introduction, for any patient stratification analysis, there exists an extensive number of potential experiments and often it is not possible to determine *a priori* which setting is most appropriate for analysis. Some of the analytical choices to be considered are:

- ❖ Data types: as modern EHR datasets become more complete, multiple data types such as diagnosis codes, laboratory values, oncology medication, or clinical reports become available. Even though ideally as much information as possible should be considered for the analysis, each data type might contain biases which could impact the stratification result;

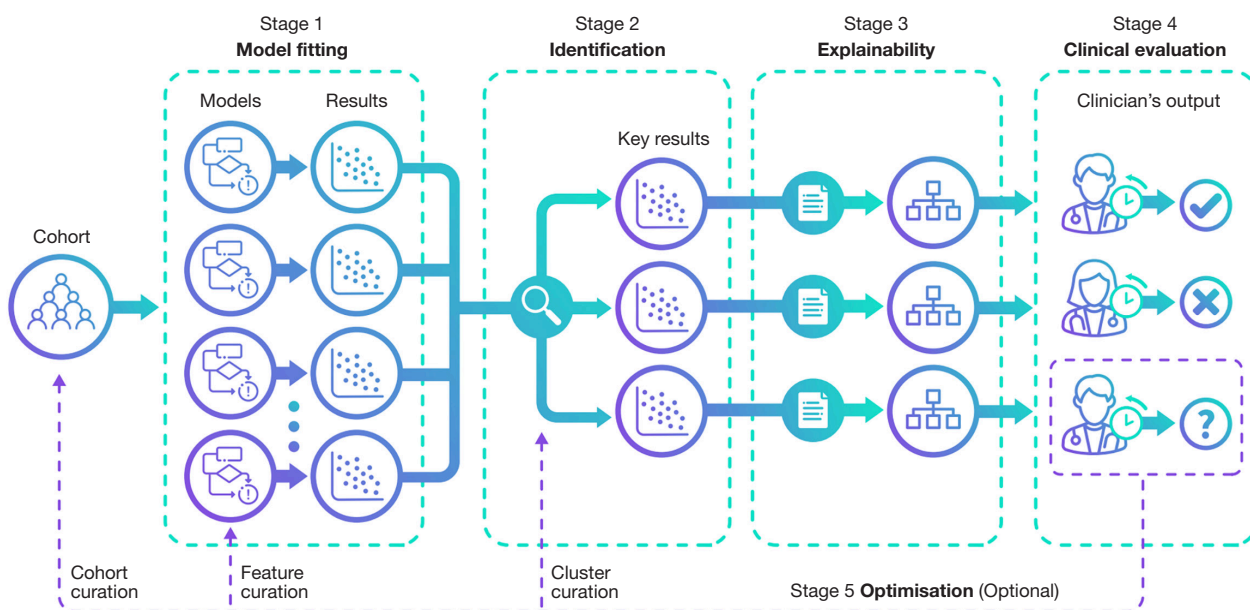


Figure 2 Proposed clinical interpretation approach which allows a faster evaluation and iteration of ML-derived patient clusters to reduce the time burden on clinical researchers. By introducing the identification (stage 2), the explainability (stage 3), and optimisation stage (stage 5) the time required to evaluate a single set of results will be reduced dramatically in stage 4. ML, machine learning.

- ❖ Data preprocessing: information in EHR data can be preprocessed in different forms, some of the choices to be made are: handling of continuous data types (e.g., laboratory values could be coded in a binary format, only indicating presence or absence of a laboratory test, or using the raw values), filtering of data elements (e.g., removing of data elements which are less than $X\%$ present in the dataset to remove noise), or hierarchy adjustment of data elements [e.g., the 10th revision of the International Statistical Classification of Diseases and Related Health Problems (ICD10) based diagnosis codes have a multiple hierarchy levels];
- ❖ Handling of temporal progression: EHR data ranges across multiple years, consequently different types of analysis are possible, e.g., considering only data recorded during a specific hospital admission or data across a specific time range;
- ❖ Cluster methods: as mentioned in the introduction there is a large number of cluster methods available ranging from simple k-means clustering (19) to more complex deep learning based clustering (18), each with its advantages and disadvantages;
- ❖ Expected numbers of clusters: for many clustering methods, the number of expected clusters k in a

dataset needs to be defined. As this is normally not known *a priori*, often multiple numbers of clusters are evaluated. Alternatively some cluster methods have built-in mechanism to find the optimal number of clusters based on a predefined criteria (20).

A reduction in the number of experiments to be analyzed can be achieved through various approaches, such as: (I) automatic screening of all results and identification of commonalities, e.g., certain patients are always grouped together, or (II) automatic ranking with respect to the clinical objectives of the study and identification of the most relevant results. Two examples of such approaches are meta clustering and pattern screening, respectively.

Meta clustering

To easily extract the common trends across multiple experiments, we can turn to approaches that consider the consensus across several different sets of results. Indeed, from a clinical perspective, it is particularly useful to have access to both qualitative and quantitative measures of the extent of agreement between different clustering algorithms. We adopted a procedure for combining results from multiple analyses using meta consensus clustering (21). A previous example of the application of meta clustering for patient sub-phenotyping can be found from Aure *et al.* (22).

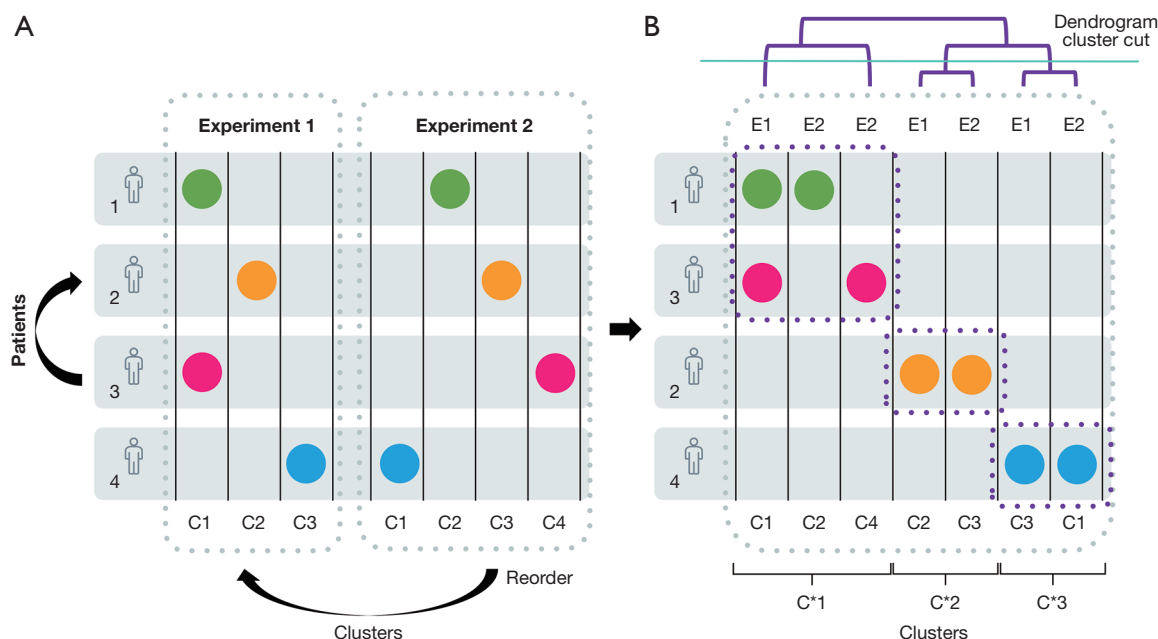


Figure 3 Illustration of meta clustering approach, where two experiments E1 and E2 are combined by reordering the patients and cluster labels so that the similar patients are in proximity. The original set of experiments and clusters (A) are reordered so that the new meta clusters (B), denoted by an asterisk * and bounded by the purple dotted lines, are now found to be C*1: E1-C1, E2-C2, E2-C4; C*2: E1-C2, E2-C3; C*3: E1-C3, E2-C1. A dendrogram in purple above (B) demonstrates how the clusters are split at each level of hierarchy. The height of the blue dendrogram cluster cut line controls the number of clusters. Note that the patients are colour coded for clarity.

In addition, we generated hierarchical clustering heatmaps, or dendrograms, which display the cluster assignments of all individual patients across the different experiments included in the meta consensus clustering analysis. These dendrograms can illustrate the degree of overlap in cluster assignments between distinct experiments.

Assuming n cluster analysis experiments were performed in stage 1 on a dataset with p patients, each individual patient x_i , with $i = \{1, \dots, p\}$, is assigned to a cluster $c_{i,j}$ for experiment j with $j = \{1, \dots, n\}$. The data is converted to a binary dataframe where the individual patients are represented along the rows and clusters for each individual experiment were represented across the columns (Figure 3A). This data is then clustered by agglomerative hierarchical clustering (23) using the average linkage technique which works by iteratively pairing and merging clusters until all clusters have been merged into a set of meta clusters c_k^* . We used the Hamming distance to quantify the degree of dissimilarity between patients in this cluster space. This generates a dendrogram representing a hierarchical structure of patient clusters that can be split at different levels of hierarchy to obtain different numbers of

clusters (Figure 3B).

Note, also for the meta clustering approach, the number of clusters k needs to be defined. To address this, meta clustering was performed for $k = \{n: n \text{ is an integer with } 1 < n < 13\}$ clusters. Each set of meta clustering results was evaluated using the average silhouette index. Higher silhouette index values indicate better cluster separability, i.e., clustering quality (24). Consequently, the numbers of clusters resulting in the first and second (if available) local maxima of silhouette index values are automatically selected to produce the results. Note, the parameter depends on the dataset and specific application and are considered as an example. Here, we focused on smaller numbers of k to ensure that identified sub phenotypes have sufficient number of patients to be practical relevant and focus only on the best 2 results to effectively reduce results which require clinical evaluation. An alternative approach for this step was presented by Rose et al, which has an automatic determination of optimal number of clusters included (25).

The meta clustering process is described succinctly below.

- (I) For n clustering experiments generate cluster label

- $c_{i,j}$, where i is the patient and j is the experiment;
- (II) Encode the cluster labels into a 2D binary matrix with the dimensions patient id and experiment cluster label (*Figure 3A*);
 - (III) Perform an ‘agglomerative hierarchical clustering’ algorithm with the ‘average linkage technique’ and assign new cluster labels c^* for every patient i ;
 - (IV) Rearrange patient ids and cluster labels within binary matrix to match cluster label E assignment (*Figure 3B*);
 - (V) (Optionally) add dendrogram on the x-axis to visualize the patient clusters (see section “Meta clustering HF experiments”);
 - (VI) Add dendrogram on the y-axis to visualize the agglomerative clustering on the experiment cluster labels $c_{i,j}$ (see section “Meta clustering HF experiments”).

Pattern screening

Usually clinicians follow an intuitive, experience-based approach when evaluating clinical cohorts and their properties; naturally this limits throughput. Therefore, we developed a pattern screening method for emulating a clinical review of experiment results. When evaluating results from multiple cluster analysis experiments, typically a reviewer will screen reports by applying a set of—often implicit—concepts derived to determine which are of clinical interest. The pattern screening process takes these concepts and implements simple algorithms, such as identifying clusters with an increased mortality risk, to automatically rank the set of results. Depending on the clinical objective, a wide range of rules can be defined to reflect the different study objectives. Note it is critical to determine rules before analysis to prevent the introduction of selection bias.

As outlined in the “Use-case method” section, the objective of our example study is to identify patient clusters which differ with respect to outcomes such as mortality. Additionally, the identified patient clusters should be easily explainable (see section “Stage 3—Explainability”). These objectives were translated into the following rules:

- (I) For all clusters m across the individual set of results n , compute the hazard ratio $HR_{m,n}$ using the “cluster vs. rest” Cox proportional hazard model with respect to mortality, recurrent stroke, bleeding events and re-hospitalisation (these are the relevant outcomes for the stroke and heart failure cohort—see “Use-case method” section);
- (II) Calculate the log rank P value, p_v , from the Cox

model and compute a result specific and outcome ranking score $R_{m,n} = -\log(p_v)$;

- (III) Repeat step 1 & 2 with patient clusters identified via the surrogate model (see section “Stage 3—Explainability”);
- (IV) Evaluate ranking results by, e.g., visualisation or sorting of results (see section “Pattern screening”) and comparison between base and surrogate model scores.

Note that comparison of base and surrogate model scores provide one means to judge the performance of surrogate models; specifically, when the ranking scores of the surrogate models are worse compared to the base results. This indicates that the surrogate model was not able to define a similar patient cluster using simple criteria.

Stage 3—Explainability

With hundreds of different data points captured in modern EHR systems, it becomes increasingly challenging to capture the key characteristics for a specific patient cluster (see *Figure 4*) especially if classical enrichment analysis are performed. The purpose of this stage is to translate complex patient clusters, which might be generated by black-box deep learning models, into understandable (and ideally explainable) results using surrogate models, such as decision trees, which provide explainability by design (26). This is a very active research field within the data science and ML community and a full presentation of all possible approaches is beyond the scope of the paper. Therefore, the explainability method employed in our analysis is chiefly a decision tree that is trained on the input model features and the cluster labels of the black-box model, which is referred to here going forward as the surrogate model. This simple surrogate model is in line with the clinical question (in particular question III) of our example patient stratification study (see “Use-case method” section), but will depend on the specific application.

Surrogate models

In general, there is a trade-off between model complexity and interpretability (27). Due to the increasing number of data points, there is a tendency to use more complex clustering models which makes rapid model interpretation difficult. Additionally, further methods might be required if consensus clustering approaches, e.g., meta clustering (28), are used. This challenge is addressed by using surrogate models, which are secondary white-box models trained to predict the outputs of the more complex model (29-31). Surrogate models give a more complete picture than enrichment

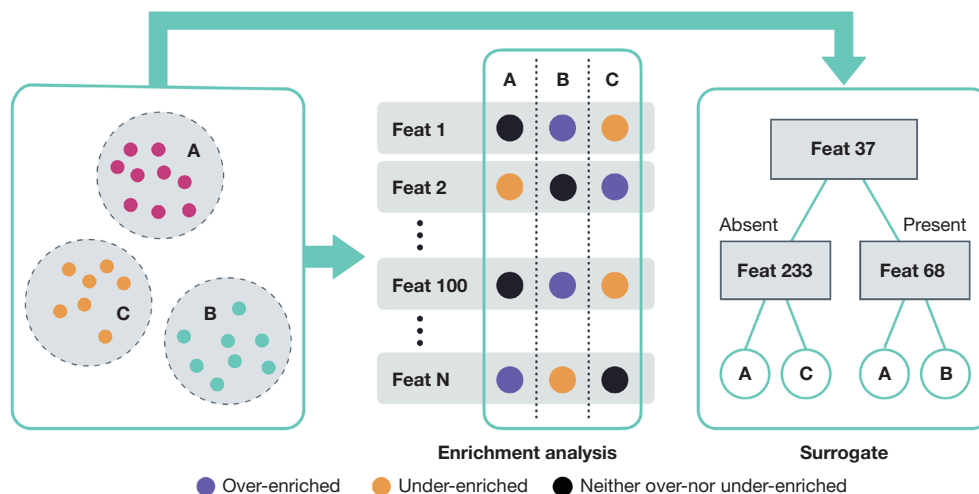


Figure 4 An illustration of two methods for interpreting patient stratification results. In this example, we assume the stratification analysis revealed 3 clusters A, B and C. A simple and direct approach is enrichment, where feature counts (and percent) for each cluster of patients is calculated. The OR between enrichment across clusters indicates the extent of over-enrichment ($OR > 1$) or under-enrichment ($OR < 1$) of a feature. With an increasing number of features in modern EHR datasets, direct enrichment analysis is impractical and alternative approaches, e.g., surrogate modelling, should be applied. A supervised decision tree finds criteria for cluster labels based on the important features in the model. Consequently, only the most important features are used which vastly simplifies the analysis. EHR, electronic health record; OR, odds ratio.

analysis alone (see *Figure 4*). While enrichment is useful for finding features that are significantly more prevalent in each cluster, surrogate models can be used to understand feature interactions and specific feature thresholds that determine patient cluster assignments.

Using a surrogate model, the trained parameters can be used to interpret the extent to which individual features influenced the clustering process. In the case of patient stratification from EHR data, this means we can explain which medical features were most important in determining which cluster a patient should be assigned to. Though there are several methods available, we use primarily supervised decision trees in our surrogate model approach. Also, the ground truth prediction labels are defined as the one-hot-encoded cluster labels from our black-box clustering algorithm. For each of these methods, we trained a model using 5-fold cross-validation. This approach involves randomly splitting the dataset into 5 groups of equal size and then iteratively selecting each of these 5 groups to be the validation set while the remaining 4 groups were used as the training set. Unlike most ML analyses, we did not use an additional test set to evaluate the final model performance as the models were intended for purely explanatory purposes. The performance of the surrogate

model can be evaluated using standard metrics such as accuracy or F1-scores.

Stage 5—Optimisation

The purpose of this stage is to provide methods for iterative optimisation of patient stratification results. For example, a model may find a patient cluster with a very high risk of mortality—in line with the objective of the study—however, the cluster might be defined by non-clinically meaningful features. Alternatively, a cluster might be identified due to biases, e.g., a data recording bias between different hospital systems. Therefore, methods to remove or modify features and/or patients to validate potential insights and safeguard against data bias are crucial. This step relies on close interaction between clinical and data experts as e.g., the removal of a non-clinical feature might introduce bias.

We propose here methods which can be applied at different stages of the stratification analysis, namely to cohort, feature, or cluster curation, to enhance the quality of the cluster analysis (see *Figure 2*). Note that this stage is highly dependent on domain expert knowledge.

The process of curating either features or patient cohorts leads to changes in the model input data which necessitates the clustering process to be re-run prior to clinical




Figure 5 Feature curation examples where diagnoses codes I48 are generalized, procedures codes Y982 are excluded and BNF medication codes are combined where the prescription type, INP versus TTA, is not clinically relevant. BNF, British National Formulary; INP, inpatient; TTA, to take away.

evaluation. Curating the clusters directly—i.e., before review and analysis—means the input data does not change and the clinical evaluation can be performed directly.

Cohort curation

The objective of cohort curation is to remove or add additional patients from the cohort of interest. This might sometimes be required as an initial cluster analysis could reveal patient clusters which are defined by a previously unknown data bias, e.g., a patient cluster contains only patients before a specific year due to a change in standard practices over time, such as the brain natriuretic peptide (BNP) test versus the later introduced N-terminal pro-hormone BNP (NT-proBNP) test for HF. Naturally, such patient groupings typically have poor clinical utility in terms of providing meaningful stratification despite formally meeting the criteria for inclusion in the cohort. In this case it can be beneficial to further exclude some patients before clustering analysis is re-run, thereby avoiding the risk of diluting potentially relevant signals in the data.

Feature curation

As real world EHR data tends to be messy and biased, there will be instances where clusters are defined based on clinically irrelevant or obvious features. The objective of feature curation is to either remove non-informative individual features or combine multiple non-clinically relevant features to a single clinical meaningful feature (see examples in *Figure 5*). This process strongly depends on domain expertise as deciding which feature is useful and not strongly depends on the objective of the analysis. Allowing the clinical researchers to quickly curate features and rerun the analysis will result in a clearer and more clinically relevant cluster definition. Note, that when feature curation is applied, the patient cohorts are unchanged, but the clustering analysis and subsequent clinical analysis are re-evaluated using the curated feature list.

Cluster curation

There are several reasons for identified clusters to be removed or to be combined with other clusters. For example, if a cluster appears to suffer from bias (such as by data source), it can be removed or if two clusters appear to exhibit similar attributes then they can be combined. It is sometimes the case that clinical evaluation of clusters leads to two or more clusters being identified as having very similar survival or enrichment profiles. When this occurs, it can be useful to combine these clusters into a single cluster. Similarly, if there are multiple clusters that are not informative or relevant to a given clinical question then it can be useful to condense these less relevant clusters into a single cluster.

As cluster curation involves manually adjusting the cluster definitions after the cluster analysis, it is not required to re-run the clustering for this type of curation to be applied. Instead, just the clinical evaluation stage is re-run with the expectation that the results from these analyses are easier to interpret due to the reduced number of clusters and therefore a smaller number of comparisons.

Use-case methods (dataset, clinical questions, and analytical analysis)

To illustrate the challenges with large scale patient stratification studies and the advantages of our proposed framework, we use results of a stratification study within the cardiovascular domain with a specific focus on stroke and HF patients. In the following section, the EHR dataset, clinical questions, cohorts, and analytical approaches will be described.

The study was conducted in accordance with the Declaration of Helsinki (as revised in 2013). The data were extracted, anonymized, and supplied by the Trust in accordance with internal information governance

Table 1 Patient data breakdown from OUH and ChelWest for ischaemic stroke, and acute heart failure

Variables	Source cardiovascular dataset		At heart failure event		At stroke event	
	OUH	ChelWest	OUH	ChelWest	OUH	ChelWest
Total patients (n=860,545)	608,759	251,786	1,430	1,062	1,480	916
Outcome mortality	38,039	21,028	516	505	448	389
Demographic						
Age (years), mean [SD]	53.1 [20.0]	59.9 [20.2]	78.1 [13.5]	78.7 [11.6]	77.6 [12.8]	77.1 [12.9]
Female, n	328,760	133,535	688	533	769	469
Male, n	279,885	118,075	742	529	711	447
No. of clinical observations, mean [SD]	39.1 [27.5]	65.1 [40.7]	75.4 [33.7]	90.0 [38.9]	77.2 [19.0]	93.2 [37.5]
Time coverage (days), mean [SD]	675 [637]	762 [1,097]	506 [453]	605 [680]	462 [442]	560 [711]
No. of unique diagnosis codes	10,800	8,907	1,064	1,137	1,186	1,219
No. of unique procedure codes	6,793	4,758	173	229	225	211
No. of unique medication (BNF) codes	14,372	3,619	116	105	116	116
No. of unique laboratory values or vital signs	90	1,599	47	123	45	123

Numbers do not add up due to additional valid entries over male and female (“not known” and “not specified”). OUH, Oxford University Hospital; ChelWest, Chelsea and Westminster Hospital; SD, standard deviation; BNF, British National Formulary.

review, National Health Service (NHS) Trust information governance approval, and the General Data Protection Regulation (GDPR) procedures outlined under the Strategic Research Agreement (SRA) and relative Data Processing Agreements (DPAs) signed by the Trust and Arcturus Data Ltd.

Dataset

The dataset used for our analysis consisted of anonymized EHRs obtained from two NHS trusts, Oxford University Hospital (OUH) and Chelsea and Westminster Hospital (ChelWest). The total number of patients available in the study was 860,545, with records spanning from August 2010 to March 2020.

The dataset contains 6 different types of clinical features: diagnosis codes (ICD-10 codes), procedure codes (OPCS-4 codes), medication codes [British National Formulary (BNF) codes], laboratory values, demographic information, and administrative information. Laboratory values are mainly continuous, while diagnosis, procedure, and medication codes are binary or categorical features (indicating presence or absence). Administrative information contained data such as start and end date of admissions and admission type (e.g., inpatient or outpatient). Diagnosis codes could appear either as a primary diagnosis (in the case that the diagnosis was the primary reason for the hospital admission) or secondary

diagnosis (in the case that the diagnosis was a comorbidity).

Clinical questions

The objectives of this patient stratification study are:

- (I) Can we identify clinical meaningful sub-phenotypes of patients within patients who have a first diagnosis of ischemic stroke or an acute HF episode?
- (II) Do these sub-phenotypes differ with respect to their clinical outcomes such as mortality rates?
- (III) Are these sub-phenotypes practically relevant, meaning that they can be defined by using few inclusion and exclusion criteria with a high degree of clinical meaning to define the population?

Cohort and outcome definitions

Based on the clinical question above, relevant criteria were applied to the starting cohort 608,759 (OUH) and 251,786 (ChelWest), which left 1,430 (OUH) and 1,062 (ChelWest) HF patients, and 1,480 (OUH) and 916 (ChelWest) stroke patients (*Table 1*). To obtain these sub-populations of patients with relevant medical profiles for our analysis, we defined the cohorts by an index date for each patient which is the first diagnosed acute event in their disease course. For the HF cohort, patients were included in the cohort if they had a HF event, which was defined as the occurrence of any of the following ICD-10 codes as a primary diagnosis:

- ❖ I50* Heart failure;
- ❖ I11.0 Hypertensive heart disease with (congestive) heart failure;
- ❖ I13.0 Hypertensive heart and renal disease with (congestive) heart failure;
- ❖ I13.2 Hypertensive heart and renal disease with both (congestive) heart failure and renal failure.

Patients were only included in the cohort if they had one of these events as well as at least 3 months of medical data available prior to the first admission where one of these diagnosis codes was recorded. We also excluded patients from the cohort based on several additional criteria, defined as:

- ❖ Patients whose first admission (as defined above) was under 48 hours and had a HF related procedure code in the 30 days following the first admission (OPCS-4 codes: K59*, K60*, K61*, K72*, K73*, K74*);
- ❖ Patients who had HF related ICD-10 codes recorded as a secondary diagnosis prior to their first admission (ICD-10 codes: I50*, I11.0, I13.0, I13.2);
- ❖ Patients who had been prescribed Eplerenone, Sacubitril with Valsartan or Spironolactone at a dose of either 25 or 50 mg prior to their first admission;
- ❖ Patients with a recorded New York Heart Association classification or patients with a recorded ejection fraction under 40% prior to their first admission;
- ❖ For the stroke cohort, we included patients who met one of the two following criteria:
 - ◆ Patients aged 18 or older who had an admission for ischaemic stroke as a primary ICD-10 code (I63* Cerebral infarction) and had at least 6 months of medical data available prior to the first admission where they had this event;
 - ◆ Patients with a primary or secondary I63* Cerebral infarction ICD-10 code OR I69.3 Sequelae of cerebral infarction ICD-10 code that occurred prior to the first record of an ischaemic stroke admission (as defined above).

The outcomes considered for these patients were mortality and for stroke patients recurrence of stroke and bleeding. Readmission for HF was also considered. The definitions for these end points are available in [Appendix 1](#).

The above cohort criteria quite significantly decrease the number of patients available for analysis, all cohorts are less than 0.5% of their original (*Table 1*). With respect to cohort features, a 1% filter was applied that ensured

all patients had feature coverage greater than 1%, e.g., at least 1% of possible features were present for a given patient. The criteria (unsurprisingly) increase the percent of patients with an outcome of mortality, where the original cohorts had 6–8% mortality, the sub-cohorts ranged from 30–47%. Note that the HF and stroke cohort demographics maintained an even sex balance, however the average patient age increased from 1950s to late 1970s. In these instances, the patient age is defined as the difference between birth date and date of event. The mean patient clinical observations increased in the sub-cohort populations as many patients with a low number of observations were filtered out. The time coverage is highly dispersed and skewed—many patients have single day coverage; hence the standard deviation is larger than the mean. However, we are focusing on “at-event” and data outside the event will be disregarded. The discrepancy between the number of unique laboratory values in OUH *vs.* ChelWest is simply a matter of tests available at given trust. Furthermore, there is only a subset of laboratory values that are relevant for the cardiovascular patient stratification study.

Statistical analysis

Data pre-processing

The data preprocessing step consists of cleaning, quality checks, standardizing, time interval aggregating, and converting to a more useful data representation.

(I) Cleaning and quality checks

There are several sources of errors and irregularities in raw datasets, such as missing values, mismatched data, typos, superfluous formatting, and duplication. Quality checks are designed to catch these types of problems, in particular demographic data is checked that the age of the patient is greater than 18 and that no clinical observations occur after the date of death. Further, demographic and laboratory measurements are checked to be within physiological ranges and that the values are of correct type and not empty.

(II) Standardization

The cleaned data is then standardized across trusts. The standardization includes variations in naming convention, e.g., ‘Amikacin’, ‘AMIKACIN’, ‘AMIK’, ‘ARM’, ‘AMIKCAIN LEVEL’, as well as by units, e.g., ‘g/L’ *vs.* ‘mg/dL’ which must be scaled. Additionally, granular subcategories are standardized, such as the mapping of medications to a parent code, e.g., ‘Timolol’, ‘Pilocarpine Nitrate’, ‘Pilocarpine Hydrochloride’ to ‘Treatment of Glaucoma’. The structure of the raw data was also harmonized, for example, data columns that were spread

across multiple columns were combined to a single column.

(III) Aggregation

As the raw EHR data consists of datapoints across time, it is necessary to aggregate and simplify the time intervals. For continuous variables, the observations within a time window are summarized using medians, median absolute deviation (MAD), count, minimum, maximum, and last observed value. For categorical variables, the observations are condensed to unique values or counts for each feature. Data was filtered by encounters for patients that contain the event of interest where the start and end dates are directly extracted from the EHR.

(IV) Filtering

Patients were filtered for sparsity. If a patient had a feature density of less than 1%, they were excluded from the cohort. Specifically, this was applied for patients with all possible laboratory test, medication, procedure, and diagnoses.

(V) Data representation

Before clustering, the data is transformed into one of the following data representations: (I) one-hot-encoding (OHE), (II) ‘GloVe’ embedding, or (III) quantisation. The simplest data representation method is OHE where the presence or absence of a feature is denoted in a binary fashion as a 1 or 0. The more complex ‘GloVe’ embedding (32) is a common natural language processing (NLP) technique which learns a dense vector representation for words trained on the word-word co-occurrences in a document (or corpus). Quantisation is a method used for handling continuous values, such as found in laboratory tests. In short, quantisation puts the continuous values into bins, quantising the values within a bin to a single value. In particular, we used distribution aware quantisation (33), where the number of data points in each quantisation bin are consistent across all bins. Note, each data representation technique has advantages and disadvantages such how continuous values are considered or with respect to how dense the representation is. Often it is not clear, which data representation results in the best result, therefore multiple approaches are tested.

Clustering methods

The cluster analysis was performed using the 4 cohorts defined in *Table 1* (2 HF and 2 stroke cohorts with approx. 4,000 patients in total) with three different data representations. We ran experiments for two clustering methods described briefly below (6 experiments in total; 3 different data representations and 2 clustering methods) per cohort. The first method is Deep Embedded Clustering (DEC). This technique was based on the work of Xie *et*

al. and consists of an auto-encoder with an additional clustering layer (34). Similarly, a Modified Variational Autoencoder (MVA) was deployed. This method uses a modified Kullback-Leibler (KL)-divergence term in the loss function during training which encourages separation of the individual clusters within the embedding space (19). The optimum number of clusters k was determined using a bootstrapping approach. Briefly, the bootstrapping approach consists of reference models and subset models. The reference models are the trained unsupervised clustering models with all data for each cluster $k=\{2,\dots,11\}$. The subset models are generated from $n_s=10$ subsets where each division represents 75% of the data sampled at random. The optimum cluster k is defined as the cluster number with the highest agreements between subset and reference models. The agreement is defined using an average Jaccard index,

$$J(C_R, C_i)_{avg,k} = \frac{1}{n_s} \sum_{i=1}^{n_s} \frac{|C_R \cap C_i|}{|C_R \cup C_i|} \quad [1]$$

where C_R and C_i is the cluster assignment of the reference model and of subset model i .

Classical clinical evaluation

(I) Enrichment methods

Enrichment analysis is a useful method for inferring feature importance from clustering results (35). In particular, enrichment tables can be generated which summarize the extent to which patients adhere to a subgroup relative to the rest of the cohort. For categorical features, we measure the total and frequency (percent) in each group. For continuous features, we measure the mean and standard deviation as well as median and interquartile range for each group. We further employ statistical analysis by calculating a P value (Fisher’s exact test or Chi-squared if all four counts if the contingency table had values greater than 10), and the odds ratio. Note that the odds ratio is simply calculated from a contingency table and is informative of the direction of the enrichment, i.e., odds ratio values greater than 1 are *over-enriched*, and less than 1 are under-enriched. In the case of continuous feature enrichment, such as laboratory values, numerical features were evaluated using the Kruskal-Wallis test. Note that we defined a P value threshold (adjusted using the Benjamini Hochberg method) of <0.05 to determine significance.

(II) Kaplan-Meier and Cox

The survival analysis for each cluster was performed on the right-censored time to event data. The Kaplan-Meier (KaplanMeierFitter) curves and Cox (CoxPHFitter)

Table 2 Experiments used for demonstration purposes in for meta clustering, pattern screening, surrogate modelling, and curation. Trust: data from either OUH and ChelWest. Investigated disease type either HF or stroke

Trust	Disease	Cluster	Experiment name	Used in
OUH	HF	3	A	Meta clustering
OUH	HF	5	B	Meta clustering
OUH	HF	2	C	Meta clustering + pattern screening
OUH	Stroke	3	D	Pattern screening
ChelWest	Stroke	3	E	Pattern screening
ChelWest	Stroke	6	F	Pattern screening
ChelWest	HF	3	G	Surrogate
OUH	HF	6	H	Feature & cohort curation (original)
OUH	HF	3	I	Feature & cohort curation
OUH	Stroke	5	J	Feature curation (original)
OUH	Stroke	5	K	Feature curation

OUH, Oxford University Hospital; ChelWest, Chelsea and Westminster Hospital; HF, heart failure.

proportional hazard values were calculated along with the 95% confidence intervals using Lifelines (36). The Cox models were adjusted for age and sex, the baseline hazard was calculated using Breslow's method, and ties were handled using Efron's method. No regularization was applied

(III) Software

The models were developed in Python (RRID:SCR_008394) using the standard open-source packages: Pandas (RRID:SCR_018214), Numpy (RRID:SCR_008633), Scikit-Learn (RRID:SCR_002577), Matplotlib (RRID:SCR_008624), and Tensorflow (RRID:SCR_016345). Additionally the statistics and survival analysis was performed using Lifelines (36).

Results

In the following four sections, results will be presented for the meta clustering, pattern screening, surrogate modelling, and curation process. These results are generated using the patient stratification methods outlined above. To demonstrate the power of this process, we illustrate with examples on a clinical study using a large-scale EHR dataset (using the cohorts defined in *Table 1*) which focuses on the identification of novel HF and stroke sub-phenotypes. Note, the focus of this section is not to examine clinical details of the results which have been identified, rather to investigate how the clinical evaluation process was impacted by the proposed

methods. To avoid confusion, the individual experiments generated for the different scenarios were labelled numerically, and experiments reviewed in this section using meta clustering are labelled using letters A–K (*Table 2*).

Meta clustering HF experiments

The benefit of meta clustering is illustrated by the HF use-case with data obtained from OUH as described in the methods section. As outlined in the data pre-processing section and the clustering methods section, 2 clustering algorithms and 3 different possible pre-processing steps were used. For each algorithm and pre-processing step, the cluster number k was iterated between 2 to 11. In theory, this would imply the generation of up to 60 reports (2 cluster methods \times 3 pre-processing steps \times 10 cluster combinations). By using the bootstrapping approach defined in the Clustering Methods section, the number of clusters k was reduced to 1 or 2 per experiment resulting in total in 10 different sets of patient stratification results with 72 clusters (*Table 3*). Applying the meta clustering approach on this set of initial results leads to 2 meta clustering reports automatically determined by the silhouette score with $k=3$ and 5 (see *Figure 6* and experiment A and B in *Table 2*). The different clusters are colour coded by the colour bars to the left and right of the dendrogram. The different cluster algorithms and pre-processing steps resulted in similar clustering results which could be identified by our

Table 3 Overview of initially generated patient stratification results using different cluster algorithms, pre-processing steps, and numbers of clusters k before applying meta clustering. Type of pre-processing described in section “Data pre-processing”, which can be OHE or ‘Glove’ embedding

Algorithm	Type of pre-processing	Number of clusters k	Number of experiments (total $n=10$)
DEC	OHE	5	1
VAE	OHE	5, 8	2
DEC	GloVe	6, 9	2
VAE	GloVe	8	1
DEC	Glove with quantisation	7, 10	2
VAE	Glove with quantisation	5, 9	2

OHE, one-hot encoding; DEC, deep embedded clustering; VAE, variational autoencoder.

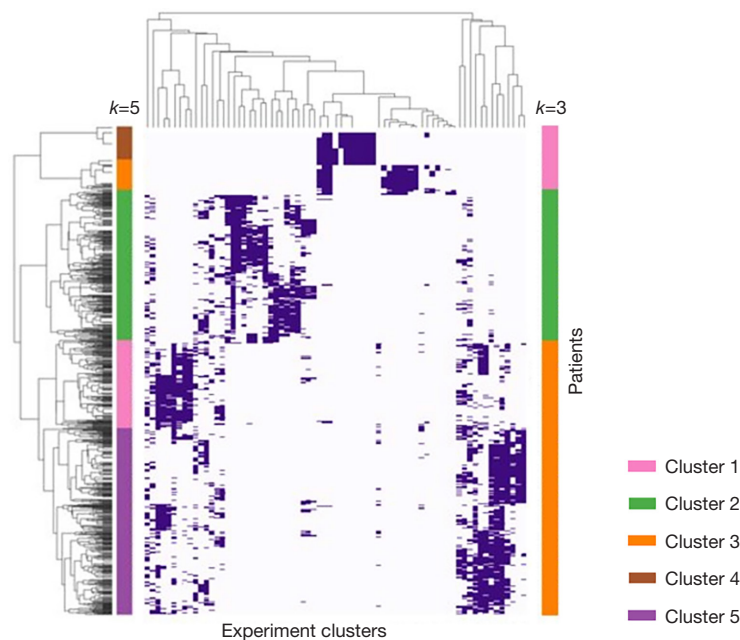


Figure 6 Example of meta clustering results using the heart failure use case and data from OUH. On the x-axis, the experiment and clusters are listed, however the x-axis labels are removed for simplicity and the y-axis shows individual patients within the dataset. The meta-clustering results are illustrated by the different colours next to the dendrogram—experiment A left side ($k=5$); and experiment B right side ($k=3$). The method identified similar cluster assignments reducing the number from 10 experiments (and 72 clusters) to 2 experiments with either 3 or 5 clusters. OUH, Oxford University Hospital.

consensus approach as indicated by the dense purple patient clusters in the dendrograms.

As an aside, another useful aspect of this approach is the ability to tune the number of clusters, whereby increasing the number of clusters, cluster groupings are further separated in a hierarchical manner. This is visible for the meta clustering results with $k=3$ clusters (Figure 6 right colour bar). The clusters 1 and 3 of $k=3$ are split into the clusters 3 and 4, and

1 and 5 for $k=5$ respectively (Figure 6 left colour bar).

Pattern screening

The value of the pattern screening method is illustrated on the stroke stratification example using results from both trusts OUH and ChelWest. Like the HF scenario, multiple encoding and cluster techniques were resulting

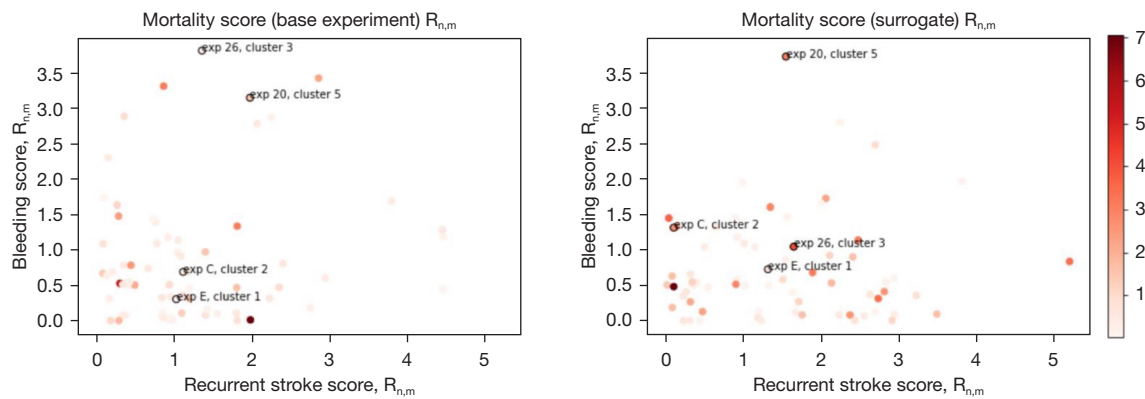


Figure 7 Scatter plot of cluster specific pattern screening scores $R_{n,m}$ of the stroke stratification experiment using data from OUH and ChelWest for the primary outcome mortality (heatmap), and the secondary outcomes bleeding (y-axis) and recurrent stroke (x-axis). Scores range from 0 to infinity which correlates with an increasing significant difference between the cluster and the remaining patients of the cohort. The scores are calculated for the base experiment (left), the surrogate model (right). The experiment number index (exp) and cluster number (k) are annotated for a few selected significant experiments. exp, experiment; OUH, Oxford University Hospital; ChelWest, Chelsea and Westminster Hospital.

in 14 experiments: 7 cluster results for OUH and 7 for ChelWest. Additionally, meta clustering was applied to the results of OUH and ChelWest separately (experiment C–F). In total, there were finally 71 sets of results to be analysed. For each identified cluster, generated either by one of the cluster algorithms or meta clustering, a pattern screening score $R_{n,m}$ was computed for the outcomes: mortality, bleeding, and recurrent stroke. In the following example, we focus on mortality as the main outcome of interest, with bleeding and recurrent stroke scores used to supplement the analysis. To visualize these scores, we provide a heatmap of the mortality score for each experiment, which are placed in a scatterplot with the bleeding and recurrent stroke scores on the y- and x-axis, respectively (Figure 7). The experiment indices are arbitrary and are not significant to the example. As outlined in the method section the pattern screening score was computed for base and surrogate models results to evaluate if a much simpler model with a few inclusion and exclusion criteria can replicate the performance of the base model.

In contrast of analysing all 71 cluster results individually, Figure 7 provides a quick overview which cluster is relevant with respect to the outcomes. In case of experiment 20, cluster 5, the mortality, bleeding and recurrent stroke score are similar high for the base and surrogate model, indicating that the surrogate model approach was able to identify a cluster with similar outcomes using a few inclusion and exclusion criteria which makes this cluster very interesting from a practical point of view. In contrast experiment 26,

cluster 3, there a big performance drop can be observed for the bleeding score, while also the mortality score increases. This indicates that the surrogate model could not identify the same patients as in the base model. The patient cohort of experiment 20, cluster 5 of the base model is most likely defined on a more complex interaction of multiple inclusion criteria.

It is also noteworthy that the meta clustering experiments, such as C and E perform relatively poorly across the board, indicating that consensus clustering, which does not take the outcomes into account, might not always result in the best results.

Surrogate models of HF experiments

In some experiments surrogate models were able to produce simple, clinically interpretable decision trees which provided clear definitions for each cluster. This can be seen in the example from meta clustering experiment G (HF, ChelWest) shown in Figure 8. Whilst surrogate models replace the need to attempt to determine cluster definitions from complex, extensive enrichment tables, the simplified definitions provided by surrogate models can be understood in more depth by subsequent and focused review of the enrichment table. For example, the definition of cluster 1 shown in Figure 8 could be further understood by reviewing the enrichment table which showed a higher prevalence in this cluster of co-morbidities such as respiratory and renal

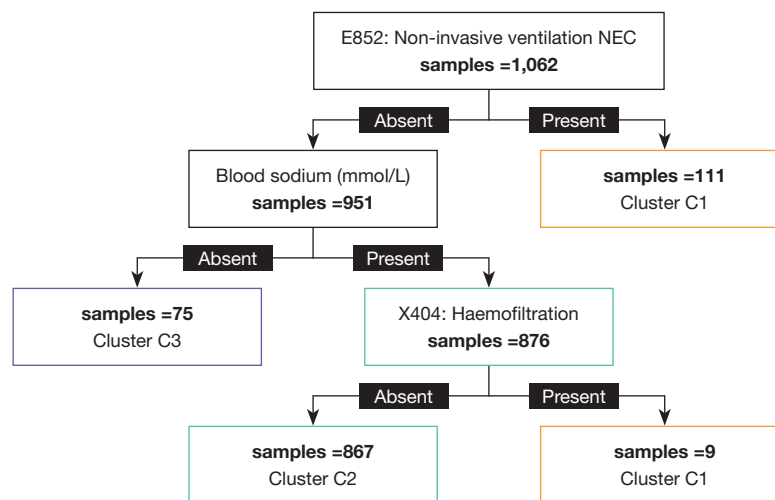


Figure 8 Decision tree from experiment G (HF, data from ChelWest). NEC, not elsewhere classified; HF, heart failure; ChelWest, Chelsea and Westminster Hospital.

diseases which are associated with the procedure codes used in the surrogate models (*Table 4*).

The surrogate model's predictive performance was variable in practice, however in some cases the surrogate model was able to predict the original clusters with a high degree of accuracy, even at a very low tree depth. For example, in meta clustering experiment G (HF, ChelWest) discussed above, the balanced accuracy of the surrogate model was 0.995 demonstrating successful prediction of the original clusters. Furthermore, both the original and surrogate model clusters showed significant differences for mortality between the clusters and therefore the clinical utility of the clusters had been preserved following application of the surrogate model (see *Figure 9*).

Cohort, feature and cluster curation

Cohort curation was used in experiment H (HF, OUH) where the clustering had identified a cluster of patients who lacked any blood tests during the period considered (*Figure 10* left). This cluster was deemed irrelevant by clinical review as it was unlikely these patients were in acute HF; typically, blood tests are required as part of routine care. Further investigation revealed these patients frequently appeared to be admitted for day case procedures such as cardiac magnetic resonance imaging (MRI), providing additional evidence against acute HF. This clustering analysis had therefore identified a previously not considered criteria which could be used to exclude patients

from the cohort on the basis that they were highly unlikely to be in acute HF. A cohort curation was subsequently performed to remove these patients, thereby removing their influence on clustering, and the experiment was re-run (*Figure 10* right). This process can be used to remove the undesirable influence of certain groups of patients identified through clustering as well as potentially adding new groups of patients. Importantly the removal of patients in this case was deemed clinically justifiable. Removal of patients could also occur based on evidence of bias producing irrelevant clusters.

Figure 11 shows a sample of surrogate model without feature curation for experiment J (stroke, OUH). Some of the features used by the model to define clusters, and appearing in the enrichment tables, in this case were discovered to be clinically irrelevant and could be removed completely, such as “Y981: Radiology of one body area (or <20 min)”. Further, several novel features with improved clinical relevance were created, for example: “Novel: Computed tomography angiography of cerebral vessels” was defined as “U212: Computed tomography NEC” AND (“Z342: Aortic arch” OR “Z35: Cerebral artery” OR “Z361: Carotid artery NEC”). *Figure 12* shows a sample of the surrogate model for the same experiment post feature curation (experiment K). All curated features are labeled as “Novel”, which are defined in [Appendix 2](#). As visible in *Figure 12*, majority of features used within the surrogate model are curated features, demonstrating the ability to improve clinical relevance of features used to define clusters.

Table 4 Truncated enrichment table showing some relevant features, such as E852 non-invasive ventilation NEC, which is enriched in cluster C1; X404 haemofiltration procedures, which is under-enriched in cluster C2 and C3, but enriched in cluster C1

Feature	All patients, n (%)	Cluster 1 (120 patients), n (%)	Cluster 2 (874 patients), n (%)	Cluster 3 (68 patients), n (%)
Diagnoses				
E872: Acidosis	53 (5.0)	28 (23.3)*	22 (2.5) [#]	3 (4.4)
J440: Chronic obstructive pulmonary disease with acute lower respiratory infection	53 (5.0)	14 (11.7)*	36 (4.1)	3 (4.4)
J9600: Acute respiratory failure; Type I [hypoxic]	20 (1.9)	9 (7.5)*	8 (0.9) [#]	3 (4.4)
J969: Respiratory failure, unspecified	43 (4.0)	30 (25.0)*	13 (1.5) [#]	0
J9690: Respiratory failure, unspecified; Type I [hypoxic]	37 (3.5)	14 (11.7)*	23 (2.6) [#]	0
J9691: Respiratory failure unspecified; Type II [hypercapnic]	39 (3.7)	29 (24.2)*	10 (1.1) [#]	0
N179: Acute renal failure, unspecified	277 (26.1)	50 (41.7)*	214 (24.5)	13 (19.1)
N185: Chronic kidney disease, stage 5	18 (1.7)	8 (6.7)*	8 (0.9) [#]	2 (2.9)
Procedures				
E852: Non-invasive ventilation NEC	111 (10.5)	111 (92.5)*	0 [#]	0 [#]
X404: Haemofiltration	20 (1.9)	20 (16.7)*	0 [#]	0 [#]

*, positive odds ratio; [#], negative odds ratio. NEC, not elsewhere classified.

Cluster curation can also be demonstrated in this experiment. In the original model for this experiment, there were 7 clusters generated where cluster 7 exhibited an increase in mortality, see *Figure 13A,13B*. However, the surrogate model failed to adequately capture and define this cluster (balanced accuracy 0.509). A cluster curation was therefore undertaken which compared patients within cluster 7 to the remaining cohort in attempt to simplify the surrogate model and obtain a definition for cluster 7, see *Figure 13C,13D*. This increased the ability of the surrogate model to discriminate cluster 7 (balanced accuracy 0.659) from the remaining patients within the cohort. Unfortunately, the surrogate model predicted cluster 7 no longer showed a significant difference in mortality compared to the rest of the population, see *Figure 13E,13F*. However, the cluster curation resulted in a simplified enrichment table when compared to the original experiment which could be used to define some features of cluster 7. These included enrichment for relevant co-morbidities such as atrial fibrillation, markers of stroke severity including gait/mobility issues and under enrichment for thrombolysis or thrombectomy (*Table 5*).

Discussion

The increasing availability of EHRs has the potential to change many aspects of healthcare and drug development. Indeed, ML analysis of EHR data has great potential to produce insights into real-world patient journeys and identify novel patient subgroups. However, we have identified three key barriers to practical application, namely: the identification of clinically relevant results in a sea of analytical outputs, the interpretation of complex black box models, and model parameter and feature optimization. We have demonstrated, based on two (HF and stroke) patient stratification analyses use cases with real-world UK-based EHR data, strategy for handling challenges in identifying patient clusters using meta clustering and pattern screening. Furthermore, we show how surrogate models can be used to explain patient phenotypes (as well as evaluating the performance of surrogates in the first place, see below), and illustrate how feature, cluster, and cohort curation can be applied to optimize model results. The output of the use cases demonstrates how to produce condensed, prioritised, interpretable, and clinically relevant results. Significantly, the proposed framework is independent of the used patient

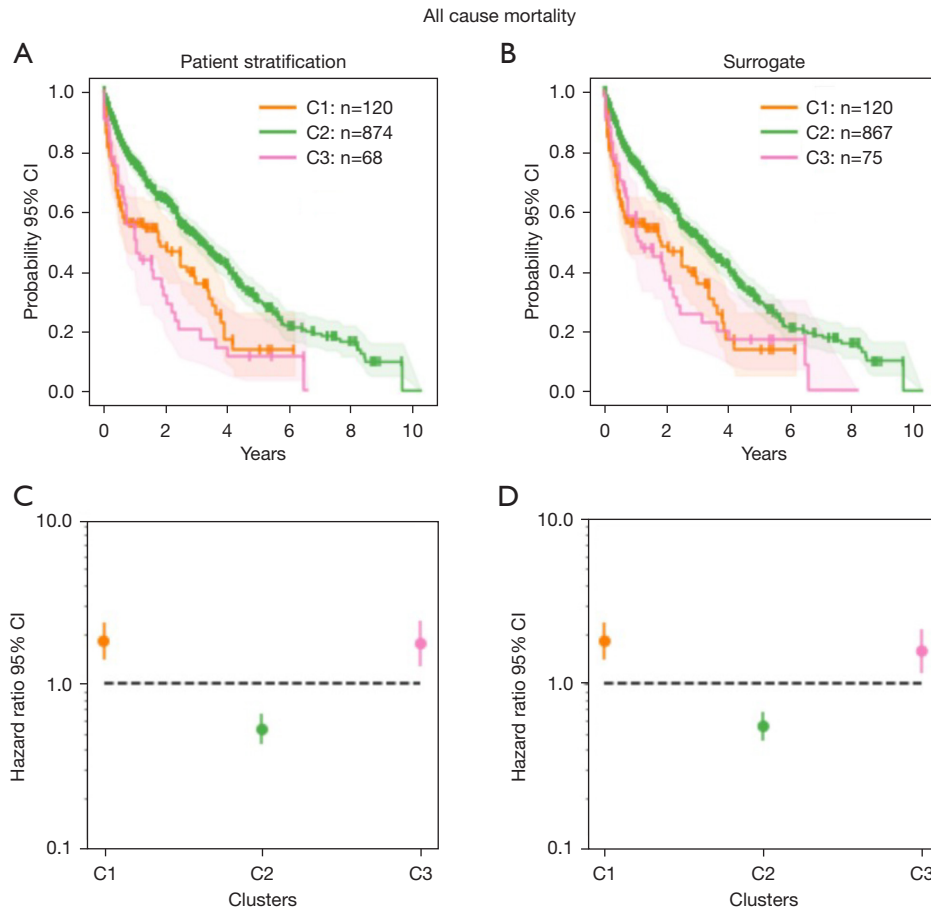


Figure 9 Experiment G, ChelWest HF survival Kaplan-Meier plot (A,B) with HRs of Cox PH models (C,D) for original patient stratification clusters and surrogate model clusters, respectively. Note that Cox PH model is controlled for patient age and sex, and the log-rank P values is less than 0.001 for both cases. CI, confidence interval; ChelWest, Chelsea and Westminster Hospital; HF, heart failure; HR, hazard ratio; PH, proportional hazard.

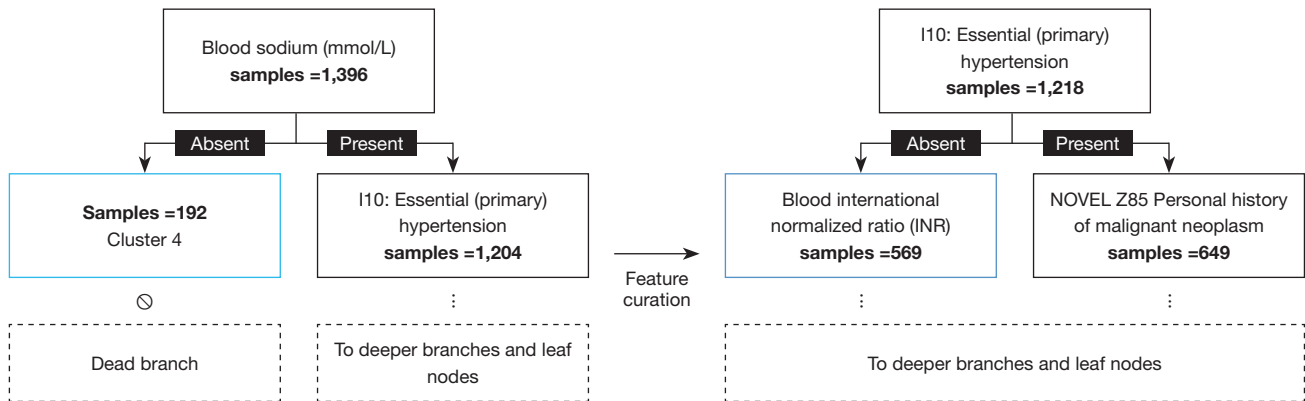


Figure 10 Surrogate model decision tree from original experiment H (left) and curated experiment I (right) (HF, OUH) and after cohort and feature curation. HF, heart failure; OUH, Oxford University Hospital.

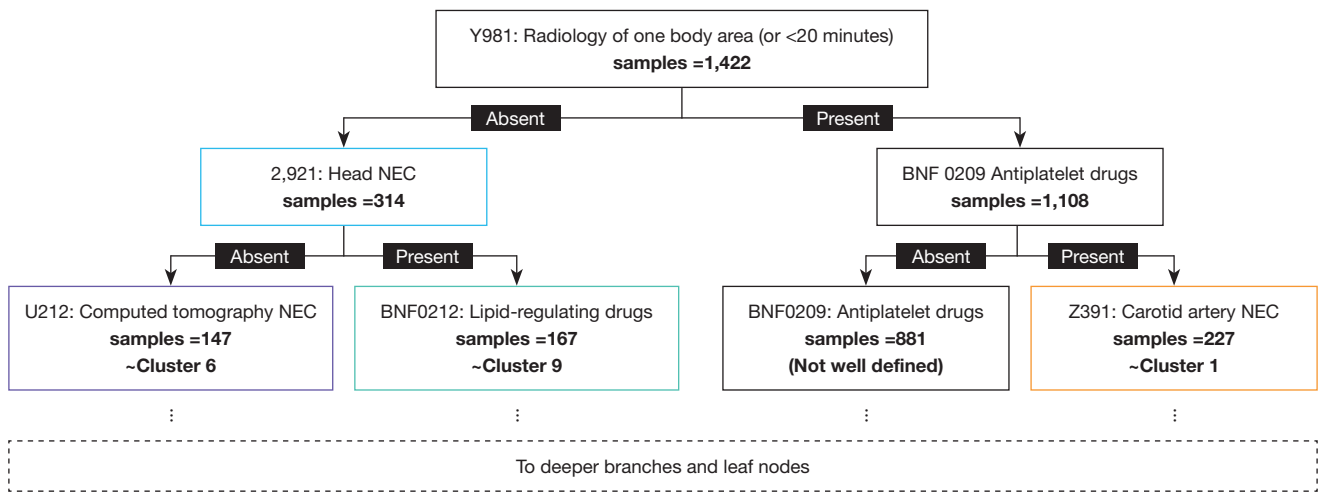


Figure 11 First 3 levels of a surrogate model decision tree from experiment J (stroke, OUH-uncurated). NEC, not elsewhere classified; BNF, British National Formulary; OUH, Oxford University Hospital.

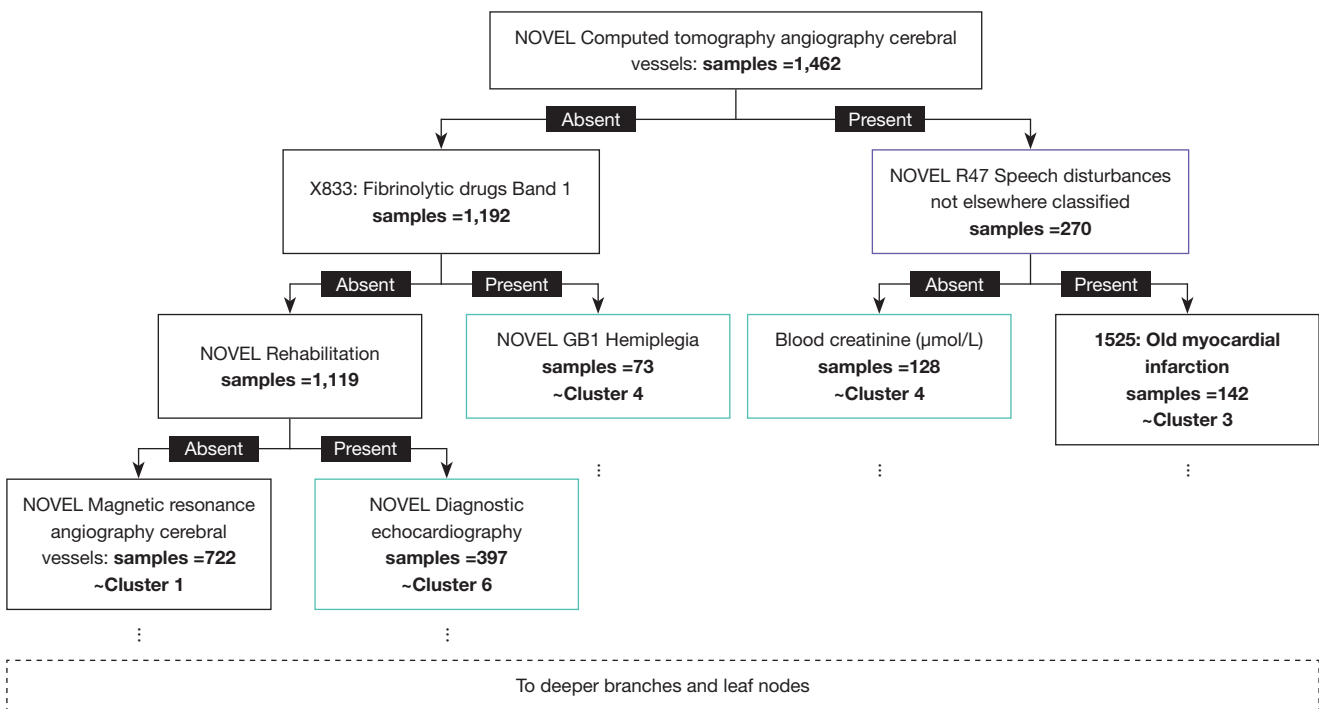


Figure 12 Decision from experiment K (stroke, OUH-curated). OUH, Oxford University Hospital.

stratification methods and does not put any constraints on the complexity of the stratification methods.

As we have shown, meta clustering allows for the rapid and simultaneous use of a wide range of models of varying complexity. There is a wealth of clustering algorithms at

the disposal of a technically savvy researcher, and we do not suggest the specific algorithms we applied above are fit for all situations. It is often best practice to start with simpler models and build complexity to suit the application. For example, though not used in this study, a common baseline

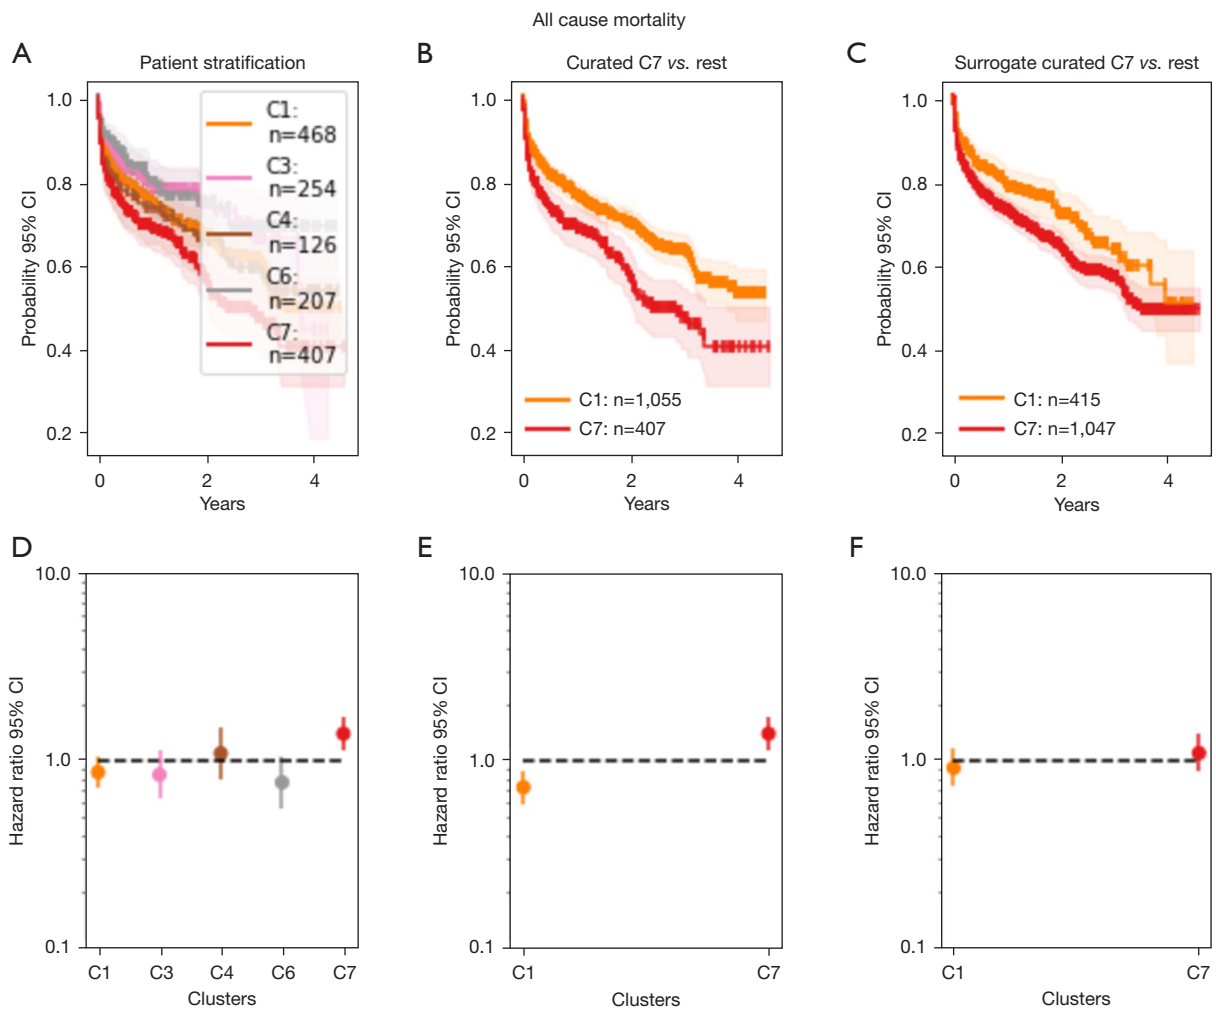


Figure 13 Experiment J and K (OUH, Stroke) survival Kaplan-Meier (A) and Cox PH model HRs (B) for the patient stratification clusters from the original model clustering. (C,D) The Kaplan-Meier curves and HRs for the curated cluster C7 compared to the rest of the cohort. (E,F) The Kaplan-Meier curves and HRs for the curated cluster C7 compared to the rest of the cohort using the surrogate model. The subplot (D) shows the HRs of the curated cluster C7 is significant, however the surrogate model is unable to capture the same behaviour (F). CI, confidence interval; OUH, Oxford University Hospital; PH, proportional hazard; HR, hazard ratio.

clustering method is k-means clustering (37). Additionally, this method can be supplemented by dimensionality reduction by converting the data into a compressed space via classical principal component analysis (PCA) for fitting k-means. To build complexity, the models should suit the nature of the data; for example, sequential or time series data is well suited to modern transformers (38) or recurrent neural networks (RNNs)/long short-term memories (LSTMs) (39); image data have been classically handled using CNNs (40). In our case we have limited the number of models for demonstrative purposes, but we suggest that

for further applications a wider selection of models can be used (not limited to deep learning).

There are several limitations to meta clustering. As shown in the meta clustering example applied to the HF cohort, we were able to reduce the number of initial reports by ~80%. However, it was not possible in all instances to find clinical meaningful meta clusters, such as with the stroke cohort using OUH data (see Appendix 3). In this example, meta clustering was successfully applied to only a subset of results, namely from the DEC model. A further challenge of the proposed meta-clustering approach is the

Table 5 Experiment K, C7 vs. Rest (OUH, stroke) selected enrichment table entries highlighting novel features

Feature	All patients, n (%)	Cluster 1 (1,055 patients), n (%)	Cluster 7 (407 patients), n (%)
Diagnoses			
NOVEL F05 Delirium not induc alcohol or psychoact	73 (4.8)	40 (3.8) [#]	30 (7.4)
NOVEL I48 Atrial fibrillation flutter	522 (35.0)	327 (31.0) [#]	187 (45.9) [*]
NOVEL I67 Other cerebrovascular diseases	183 (12.0)	109 (10.3) [#]	73 (17.9) [*]
NOVEL J18 Pneumonia organism unspecified	100 (6.8)	55 (5.2) [#]	42 (10.3) [*]
NOVEL M15 19 Arthrosis	115 (7.8)	67 (6.4) [#]	46 (11.3) [*]
NOVEL M81 Osteoporosis without path fract	61 (4.1)	28 (2.7) [#]	32 (7.9) [*]
NOVEL R26 Abnormalities of gait and mobility	159 (11.0)	84 (8.0) [#]	75 (18.4) [*]
R296: Tendency to fall, not elsewhere classified	184 (12.0)	112 (10.6) [#]	71 (17.4) [*]
Procedures			
NOVEL Computed tomography angiography cereb vessels	277 (19.0)	269 (25.5) [*]	1 (0.2) [#]
NOVEL Magnetic resonance angiography cereb vessels	27 (1.8)	25 (2.4) [*]	0 [#]
NOVEL Magnetic resonance imaging head	115 (7.8)	90 (8.5)	21 (5.2) [#]

Numbers do not add up since small clusters were removed to safeguard participants' privacy. *, positive odds ratio; #, negative odds ratio. OUH, Oxford University Hospital.

selection of the optimal number of meta-clusters. Further approaches to automate this should be investigated such as proposed in (25). In general, meta clustering evaluation requires that all individual results use the same patient cohorts. Performing meta clustering analysis across multiple datasets is not possible. A further limitation of this approach is that only similarities across cluster assignments are considered and not the clinical outcomes. Both limitations can be overcome by using pattern screening.

Pattern screening is a simple yet effective way to handle large sets of analytical outputs and model reports. Not only can this approach handle multiple results across different models, but it can also compare different datasets and multiple outcomes. In our use case, we have shown it is effective in rapidly pinpointing patient clusters with relevant outcome characteristics and simple inclusion and exclusion criteria (see clinical questions).

Pattern screening is also highly flexible, where any number of metrics or scoring functions can be applied. Note, that the pattern screening metrics we chose for the illustration are simplistic, and going forward, other metrics should be explored for instance metrics which consider the within cluster similarity. The flexibility of pattern screening is underscored by the fact it can be used independently or in conjunction with meta clustering. Furthermore,

in contrast to meta clustering where multiple results are combined and single unique results might be lost, pattern screening focuses on each individual cluster result. This is an important feature as there can be cases where a perceived outlier may in truth be a novel insight. This can also be seen in *Figure 7*, where though both individual as well as meta clustering results were ranked. Clusters of meta clustering results seem to have lower pattern screening scores and therefore have less relevant clinical outcomes. This is sensible, as we are finding the consensus among a set of analyses, which independently of the clinical outcome. Additionally, meta clustering performed poorly in case of experiment C (see [Appendix 3](#)).

Which specific method is used as surrogate model, depends strongly on the objectives of the study. We have used a simple decision tree, which matches the requirements of the third clinical question, meaning that a relevant cluster can be defined by a few clinical meaningful inclusion and exclusion criteria. This is particularly useful in the context of clinical trials, where each additional patient selection criteria can have a big impact on patient recruitment. Additionally, decision trees provide a very intuitive way for clinical experts to understand how specific clusters are defined and to evaluate if the used medical concepts are relevant or if part of the analysis should be repeated (see

optimisation stage—feature curation).

However, there are several drawbacks in using tree-based methods, chiefly among them are the inability to handle temporal data. In our example, all temporal feature data was aggregated before it was applied to the surrogate model. In future work, there is scope to develop surrogate models that can accommodate patient trajectories. This resulted partly in low accuracy values of the trained surrogate models and big difference between the pattern screening scores of the base and surrogate model for some cluster results (see *Figure 7*, experiment 26 $k=3$). However, if the study objective focuses, for instance, only on understanding which clinical parameters are relevant, there are several well-developed more advanced surrogate models, such as Ripper (41), Trepan (42), or RuleFit (43)—the details of which are beyond the scope of this study—which could result in a better overlap between base and surrogate model.

It is worth noting that there is an inherent ‘cost of explainability’ when using surrogate models (27); a black box model may achieve a higher level of accuracy than a white box or surrogate white box model, but in the field of medicine and clinical research, explainability is valued highly, and it is practical to trade accuracy for explainability. This is not necessarily a problem, but it must be understood and considered throughout the model development and interpretation process.

Of the optimisation steps discussed in this study, feature, cluster, and cohort curation undoubtedly requires the most input from a domain expert. The creation of some novel features and the removal of others features can be intuitive in some cases for a clinical researcher and in other cases it is important to consider why irrelevant features may have been selected and what signal they represent from the data. An iterative approach to feature curation was used with different thresholds for combining and removing features to find a balance between searching for unbiased data driven insights versus ensuring clinical relevance of results. The choice of features to use is not trivial. On the one hand, it is critical to avoid confounding variables, which can cause an association to appear that doesn’t exist. Many times this happens when an important variable is not controlled for—sometimes called a forking confounder—which distorts the association (44). This can be achieved by adding all the features and “regressing out confounding effects from each input variable” before model training, or they can be controlled *post hoc* (45). Yet there is a danger of blindly using every covariate available, which comes in the form of colliding confounders (46,47). A famous example

of a colliding confounder is the obesity paradox in HF; in short, it has been found that patients stratified by mild to moderate obesity were associated with a decreased mortality risk (48). Unsurprisingly, there is little evidence to suggest obesity has any real protective properties with respect to cardiovascular outcomes, and in fact has a negative impact on patient survival. To handle these types of biases requires domain knowledge, and these examples highlight the need for careful feature selection.

Curation is a somewhat subjective process. Based on training and experience, clinical researchers can identify meaning—or lack thereof—in EHR signals. This, for example, applies to the relevance of features for individual patients (e.g., discarding a laboratory assessment as not relevant for a specific patient), but also for definitions of a patient cohort. Such definitions should ideally be simple to apply and concentrate on medically meaningful concepts that relate to the condition or outcome of interest. These concepts are difficult to implement in an unbiased way and usually requires expert input. Our approach to curation at different levels of analysis (cohort, feature, and cluster) therefore aimed to make it easy to obtain this input from clinical experts by reducing definition criteria and make their meaning explicit.

Cohort curation has some overlap with data quality analysis and cleaning. Plainly, any spurious patient data should be omitted, e.g., male patients in a gestational diabetes cohort or female patients in a prostate cancer cohort. However, cohort curation should also ensure that the patients represent the target population. This implies that there should ideally be a balance between sexes (if not a sex dependent disease), as well as a distribution of patient ages that correspond to the target population. It is also important that underrepresented ethnicities are not lost. And in some cases, these populations should be enriched, especially if they represent an outsized burden for a disease area. In the context of clinical trials, this has been clearly outlined in a 2022 USA Food and Drug Administration (FDA) draft guidance for industry (49).

Conclusions

The primary objective of presented framework presented is to facilitate the identification of clinical meaningful subpopulations within a larger sample based on patient characteristics in their EHRs. The utility of defining patient clusters with clinically relevant features is underscored by the need to find suitable data driven clinical trial criteria.

A natural potential application of the methods outlined in this study would be to use an understanding of patient subgroups to guide selection for clinical trials, allowing for patients of different profiles to be recruited based on prior diagnoses, blood test results or medical procedures. Moreover, the subgroups produced by the clustering algorithm can support targeted recruitment of patients for clinical trials who are likely to be enriched with the desired outcome such as mortality, recurrence of disease or secondary conditions.

In addition to clinical trial optimisation, this scalable clinical interpretation approach can be applied to fundamental research of a disease area. By pin-pointing sub-phenotypes, we can better understand disease progression for a richer variety of patients. And in the future, with more widespread genomic and transcriptomic testing matched to full patient EHRs, we can drastically increase our understanding of disease and improve therapy selection (50,51).

We have shown in depth the strength of the framework through a use case and have highlighted the key techniques developed. However, we have found that there are several limitations inherent to the process, which include limits in the data itself, the use and misuse of the process. Despite recent advances, it is still common to find EHR datasets which are incomplete, with missing documentation for patients from general practitioners, imaging diagnostic, or surgeries. Furthermore, limitations may arise impacting longitudinal aspects, e.g., data only being available for a certain period making it unfeasible to observe both health, start, and advanced stages for many chronic diseases in the same patient. Moreover, the breadth and quality of the data will naturally affect the feasibility of the proposed scalable clinical interpretation models.

Additionally, though more of a caveat than a limitation, we do not propose a replacement for human analysis, rather a method for optimised human-in-the-loop analysis. A corollary of this is the necessity of a domain expert to define the models and metrics. For example, our current approach uses a silhouette score that automatically determines the optimal number of meta cluster k . However, this formulation may not be well suited to answer any given clinical questions as it only considers the average width of a cluster and its distance to the nearest neighbouring cluster. Metrics such as these still require both clinical and technical expertise for this approach to be successful.

In summary, we have developed a framework to analysis patient stratification results at scale and presented example methods with to support the clinical evaluation process.

Some of them emulate parts of an ‘intuitive’ approach a clinical researcher may choose to take when reviewing EHR analyses. These methods not only increase clinical meaning and facilitate throughput of the analyses, but they also provide a common language for data scientists and clinicians to support collaboration in an interactive approach. With this work, we hope to encourage the community to further investigate methods to facilitate clinical evaluation of EHR analysis and we expect the methods presented here, and future improvement based on this, to substantially increase learnings and help unlock the potential of real-world data in improving clinical practice and focusing drug development efforts.

Acknowledgments

We thank Oxford University Hospitals and Chelsea and Westminster Hospitals trust for access to anonymized data, as well as Sam Allen for designing our figures and charts. Additionally, we thank specifically Christian Diedrich, Eren Elci, Steffen Schaper, Basel Abu-Jamous, Niklas Kokkola, Fernando Andreotti and all Arcturis Data and Bayer members of this research project. This work uses anonymized data provided by patients collected by Chelsea and Westminster Hospital NHS Foundation Trust and Oxford University Hospitals NHS Foundation Trust as part of their care and support. We believe using the patient data is vital to improve health and care for everyone and would, thus, like to thank all those involved for their contribution. The data were extracted, anonymized, and supplied by the Trust in accordance with internal information governance review, NHS Trust information governance approval, and the GDPR procedures outlined under the Strategic Research Agreement (SRA) and relative Data Processing Agreements (DPAs) signed by the Trust and Arcturis Data Ltd. This research has been conducted using the Oxford University Hospitals NHS Foundation Trust Clinical Data Warehouse, which is supported by the NIHR Oxford Biomedical Research Centre and Oxford University Hospitals NHS Foundation Trust. Special thanks to Kerrie Woods, Kinga Varnai, Oliver Freeman, Hizni Salih, Steve Harris and Professor Jim Davies.

Funding: This work was funded by Arcturis Data Ltd. & Bayer AG.

Footnote

Data Sharing Statement: Available at <https://jmai.amegroups.com/article/view/10.21037/jmai-22-42/dss>

Peer Review File: Available at <https://jmai.amegroups.com/article/view/10.21037/jmai-22-42/prf>

Conflicts of Interest: All authors have completed the ICMJE uniform disclosure form (available at <https://jmai.amegroups.com/article/view/10.21037/jmai-22-42/coif>). KP, AD and PH are employees of Bayer AG who funded this project. The other authors have no conflicts of interest to declare.

Ethical Statement: The authors are accountable for all aspects of the work in ensuring that questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved. The study was conducted in accordance with the Declaration of Helsinki (as revised in 2013). The data were extracted, anonymized, and supplied by the Trust in accordance with internal information governance review, NHS Trust information governance approval, and the General Data Protection Regulation (GDPR) procedures outlined under the Strategic Research Agreement (SRA) and relative Data Processing Agreements (DPAs) signed by the Trust and Arcturis Data Ltd.

Open Access Statement: This is an Open Access article distributed in accordance with the Creative Commons Attribution-NonCommercial-NoDerivs 4.0 International License (CC BY-NC-ND 4.0), which permits the non-commercial replication and distribution of the article with the strict proviso that no changes or edits are made and the original work is properly cited (including links to both the formal publication through the relevant DOI and the license). See: <https://creativecommons.org/licenses/by-nc-nd/4.0/>.

References

- Leonelli S. Data-centric biology. University of Chicago Press, 2016.
- Wiens J, Shenoy ES. Machine Learning for Healthcare: On the Verge of a Major Shift in Healthcare Epidemiology. *Clin Infect Dis* 2018;66:149-53.
- Beckmann JS, Lew D. Reconciling evidence-based medicine and precision medicine in the era of big data: challenges and opportunities. *Genome Med* 2016;8:134.
- Collins FS, Varmus H. A new initiative on precision medicine. *N Engl J Med* 2015;372:793-5.
- Parimbelli E, Marini S, Sacchi L, et al. Patient similarity for precision medicine: A systematic review. *J Biomed Inform* 2018;83:87-96.
- Javer A, Parsons O, Carr O, et al. Compensating trajectory bias for unsupervised patient stratification using adversarial recurrent neural networks. *arXiv Preprint arXiv:2112.07239*. 2021.
- Chen R, Sun J, Dittus RS, et al. Patient Stratification Using Electronic Health Records from a Chronic Disease Management Program. *IEEE J Biomed Health Inform* 2016. [Epub ahead of print]. doi: 10.1109/JBHI.2016.2514264.
- Hedman ÅK, Hage C, Sharma A, et al. Identification of novel pheno-groups in heart failure with preserved ejection fraction using machine learning. *Heart* 2020;106:342-9.
- Kobayashi M, Huttin O, Magnusson M, et al. Machine Learning-Derived Echocardiographic Phenotypes Predict Heart Failure Incidence in Asymptomatic Individuals. *JACC Cardiovasc Imaging* 2022;15:193-208.
- Parker JS, Mullins M, Cheang MC, et al. Supervised risk predictor of breast cancer based on intrinsic subtypes. *J Clin Oncol* 2009;27:1160-7.
- Segar MW, Patel KV, Ayers C, et al. Phenomapping of patients with heart failure with preserved ejection fraction using machine learning-based unsupervised cluster analysis. *Eur J Heart Fail* 2020;22:148-58.
- Uijl A, Savarese G, Vaartjes I, et al. Identification of distinct phenotypic clusters in heart failure with preserved ejection fraction. *Eur J Heart Fail* 2021;23:973-82.
- Vazquez Guillaumet R, Ursu O, Iwamoto G, et al. Chronic obstructive pulmonary disease phenotypes using cluster analysis of electronic medical records. *Health Informatics J* 2018;24:394-409.
- Vranas KC, Jopling JK, Sweeney TE, et al. Identifying Distinct Subgroups of ICU Patients: A Machine Learning Approach. *Crit Care Med* 2017;45:1607-15.
- Williams JB, Ghosh D, Wetzel RC. Applying Machine Learning to Pediatric Critical Care Data. *Pediatr Crit Care Med* 2018;19:599-608.
- Karwath A, Bunting KV, Gill SK, et al. Redefining β -blocker response in heart failure patients with sinus rhythm and atrial fibrillation: a machine learning cluster analysis. *Lancet* 2021;398:1427-35.
- Wang Y, Zhao Y, Therneau TM, et al. Unsupervised machine learning for the discovery of latent disease clusters and patient subgroups using electronic health records. *J Biomed Inform* 2020;102:103364.
- Carr O, Javer A, Rockenschaub P, et al. Longitudinal patient stratification of electronic health records with flexible adjustment for clinical outcomes. *Proceedings of*

- Machine Learning Research 2021;158:220-38.
19. MacQueen J. Some methods for classification and analysis of multivariate observations. In: Proceedings of the fifth Berkeley Symposium on Mathematical Statistics and Probability, 1967:281-97.
 20. Hothorn T, Hornik K, Zeileis A. ctree: Conditional inference trees. *Compr R Arch Netw*. 2015;8. Available online: <https://cran.r-project.org/web/packages/partykit/vignettes/ctree.pdf>
 21. Monti S, Tamayo P, Mesirov J, et al. Consensus clustering: a resampling-based method for class discovery and visualization of gene expression microarray data. *Machine Learning* 2003;52:91-118.
 22. Aure MR, Vitelli V, Jernström S, et al. Integrative clustering reveals a novel split in the luminal A subtype of breast cancer with impact on outcome. *Breast Cancer Res* 2017;19:44.
 23. Müllner D. Modern hierarchical, agglomerative clustering algorithms. *arXiv Prepr arXiv11092378*. 2011.
 24. Rousseeuw PJ. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *J Comput Appl Math* 1987;20:53-65.
 25. Rose TD, Bechtler T, Ciora OA, et al. MoSBI: Automated signature mining for molecular stratification and subtyping. *Proc Natl Acad Sci U S A* 2022;119:e2118210119.
 26. Loyola-Gonzalez O. Black-box vs. white-box: Understanding their advantages and weaknesses from a practical point of view. *IEEE Access* 2019;7:154096-113.
 27. Goethals S, Martens D, Evgeniou T. The non-linear nature of the cost of comprehensibility. *Journal of Big Data* 2022;9:30.
 28. Caruana R, Elhawary M, Nguyen N, et al. Meta clustering. In: Sixth International Conference on Data Mining (ICDM'06). 2006:107-18.
 29. Srivastava M. A Surrogate data-based approach for validating deep learning model used in healthcare. In: Applications of Deep Learning and Big IoT on Personalized Healthcare Services. IGI Global, 2020:132-46.
 30. Barmada S, Fontana N, Formisano A, et al. A deep learning surrogate model for topology optimization. *IEEE Trans Magn* 2021;57:7200504.
 31. Thakur A, Chakraborty S. A deep learning based surrogate model for stochastic simulators. *Probabilistic Eng Mech* 2022;68:103248.
 32. Pennington J, Socher R, Manning CD. Glove: Global vectors for word representation. In: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP). 2014:1532-43.
 33. Hong C, Kim H, Oh J, Lee KM. DAQ: Distribution-Aware Quantization for Deep Image Super-Resolution Networks. *arXiv Prepr arXiv201211230*. 2020.
 34. Xie J, Girshick R, Farhadi A. Unsupervised deep embedding for clustering analysis. In: International Conference on Machine Learning. 2016:478-87.
 35. Freudenberg JM, Joshi VK, Hu Z, et al. CLEAN: CLustering Enrichment ANalysis. *BMC Bioinformatics* 2009;10:234.
 36. Davidson-Pilon C. lifelines: survival analysis in Python. *J Open Source Softw* 2019;4:1317.
 37. Horiuchi Y, Tanimoto S, Latif AHMM, et al. Identifying novel phenotypes of acute heart failure using cluster analysis of clinical variables. *Int J Cardiol* 2018;262:57-63.
 38. Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. *NIPS'17: Proceedings of the 31st International Conference on Neural Information Processing Systems*. December 2017:6000-10.
 39. Graves A. Long Short-Term Memory. In: Supervised Sequence Labelling with Recurrent Neural Networks. Studies in Computational Intelligence. vol 385. Springer, Berlin, Heidelberg. 2012.
 40. Sermanet P, Chintala S, LeCun Y. Convolutional neural networks applied to house numbers digit classification. In: Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012). 2012:3288-91.
 41. Cohen WW. Fast effective rule induction. In: Machine learning proceedings 1995. Elsevier, 1995:115-23.
 42. Craven M, Shavlik J. Extracting tree-structured representations of trained networks. *NIPS'95: Proceedings of the 8th International Conference on Neural Information Processing Systems*. November 1995:24-30.
 43. Friedman JH, Popescu BE. Predictive learning via rule ensembles. *Ann Appl Stat* 2008;2:916-54.
 44. Aronson JK, Bankhead C, Nunan D. Catalogue of Bias Collaboration: Confounding. *Cat Bias [Internet]*. 2018. Available online: www.catalogueofbiases.org/biases/confounding
 45. Dinga R, Schmaal L, Penninx BWJH, et al. Controlling for effects of confounding variables on machine learning predictions. *BioRxiv*. 2020. doi: <https://doi.org/10.1101/2020.08.17.255034>.
 46. Lee H, Aronson JK, Nunan D. Catalogue of Bias Collaboration: Collider. *Cat Bias [Internet]*. 2019. Available online: <https://catalogofbias.org/biases/collider-bias/>
 47. Sackett DL. Bias in analytic research. *J Chronic Dis* 1979;32:51-63.

48. Clark AL, Fonarow GC, Horwich TB. Obesity and the obesity paradox in heart failure. *Prog Cardiovasc Dis* 2014;56:409-14.
49. Diversity Plans to Improve Enrollment of Participants From Underrepresented Racial and Ethnic Populations in Clinical Trials. 2022. Available online: <https://www.fda.gov/media/157635/download>
50. Docking TR, Parker JDK, Jädersten M, et al. A clinical transcriptome approach to patient stratification and therapy selection in acute myeloid leukemia. *Nat Commun* 2021;12:2474.
51. Fujio K, Takeshima Y, Nakano M, et al. Review: transcriptome and trans-omics analysis of systemic lupus erythematosus. *Inflamm Regen* 2020;40:11.

doi: 10.21037/jmai-22-42

Cite this article as: Parsons O, Barlow NE, Baxter J, Paraschin K, Derix A, Hein P, Dürichen R. Enabling scalable clinical interpretation of machine learning (ML)-based phenotypes using real world data. *J Med Artif Intell* 2023;6:2.

Appendix 1: Outcome definition

In a similar manner to the definition of mortality, the specific definitions of the outcomes: recurrence of stroke, bleeding events, and readmission for heart failure are provided.

Recurrence of stroke

Presence of repeat hospitalisation for ischaemic stroke (as defined below) following index hospitalisation for ischaemic stroke. Admission to an inpatient unit or emergency department with any of the following primary ICD10 codes:

- ❖ I63* Cerebral infarction.

Excluding:

Any admissions with the admission code:

- ❖ 81-Transfer of any admitted PATIENT from other Hospital Provider other than in an emergency (NHS Data Dictionary).

Bleeding events

Any of the following primary ICD10 codes:

- ❖ K25.0 Acute gastric ulcer with haemorrhage;
- ❖ K25.2 Acute gastric ulcer with both haemorrhage and perforation;
- ❖ K25.4 Chronic or unspecified gastric ulcer with hemorrhage;
- ❖ K25.6 Chronic or unspecified gastric ulcer with both hemorrhage and perforation;
- ❖ K26.0 Acute duodenal ulcer with hemorrhage;
- ❖ K26.2 Acute duodenal ulcer with both hemorrhage and perforation;
- ❖ K26.4 Chronic or unspecified duodenal ulcer with hemorrhage;
- ❖ K26.6 Chronic or unspecified duodenal ulcer with both hemorrhage and perforation;
- ❖ K27.0 Acute peptic ulcer, site unspecified, with hemorrhage;
- ❖ K27.2 Acute peptic ulcer, site unspecified, with both hemorrhage and perforation;
- ❖ K27.4 Chronic or unspecified peptic ulcer, site unspecified, with hemorrhage;
- ❖ K27.6 Chronic or unspecified peptic ulcer, site unspecified, with both hemorrhage and perforation;
- ❖ K28.0 Acute gastrojejunal ulcer with hemorrhage;
- ❖ K28.2 Acute gastrojejunal ulcer with both hemorrhage and perforation;
- ❖ K28.4 Chronic or unspecified gastrojejunal ulcer with hemorrhage;
- ❖ K28.6 Chronic or unspecified gastrojejunal ulcer with both hemorrhage and perforation;
- ❖ K92.0 Hematemesis;
- ❖ I60* Subarachnoid hemorrhage;
- ❖ S06.6 Traumatic subarachnoid hemorrhage;
- ❖ I62.0 Nontraumatic subdural hemorrhage;
- ❖ S06.5 Traumatic subdural hemorrhage;
- ❖ I61* Intracerebral hemorrhage;
- ❖ I62.1 Nontraumatic extradural hemorrhage;
- ❖ I62.9 Nontraumatic intracranial hemorrhage, unspecified;
- ❖ K92.1 Melena;
- ❖ I85.0 Oesophageal varices with bleeding;
- ❖ I98.3 Oesophageal varices with bleeding in diseases classified elsewhere.

OR

Administration of any of the following medications:

- ❖ Dried prothrombin complex;
- ❖ Fresh frozen plasma;
- ❖ Idarucizumab;

OR

Administration of any of the following procedures (OPCS4):

- ❖ G20.1 Fiberoptic endoscopic coagulation of bleeding lesion of oesophagus;
- ❖ G46.2 Fiberoptic endoscopic coagulation of bleeding lesion of upper gastrointestinal tract;
- ❖ X33.2 Intravenous blood transfusion of packed cells.

Readmission for heart failure

Inpatient admission or emergency department admissions 2+ days in length which meet the following criteria:

- ❖ A primary ICD10 code for HF defined as:
 - ◆ I50* Heart failure;
 - ◆ I11.0 Hypertensive heart disease with heart failure;
 - ◆ I13.0 Hypertensive heart and renal disease with (congestive) heart failure;
 - ◆ I13.2 Hypertensive heart and renal disease with both (congestive) heart failure and renal failure.

OR

- ❖ A secondary ICD10 code for HF as (defined above) AND IV Furosemide administered within admission or within 1 day following admission

Excluding

- ❖ Any admission that is less than 48 hours during which one of the following procedures is performed:
 - ◆ Cardioverter defibrillator introduced through the vein (OPCS-4: K59*);
 - ◆ Other cardiac defibrillator (OPCS-4: K72*);
 - ◆ Cardiac pacemaker system introduced through vein (OPCS-4: K60*);
 - ◆ Other cardiac pacemaker system (OPCS-4: K61*);
 - ◆ Other cardiac pacemaker system introduced through vein (OPCS-4: K73*);
 - ◆ Cardiac pacemaker system (OPCS-4: K74*).

Appendix 2: Novel feature definitions

The specific novel feature definitions that are discussed in this paper are provided here. Each line is connected by an ‘AND’ or ‘OR’ statement, which can be further connected by an indentation of the same logic.

Computed tomography angiography of cerebral vessels

U212: Computed tomography NEC
AND
(
Z342: Aortic arch
OR
Z35: Cerebral artery OR Z361: Carotid artery NEC
)

Speech disturbances not elsewhere specified

All R47 subcodes combined

Rehabilitation

Z501: Other physical therapy

OR

Z505: Speech therapy

OR

Z507: Occupational therapy and vocational rehabilitation, not elsewhere classified

Hemiplegia

All G81 subcodes combined.

Magnetic resonance angiography of cerebral vessels

U211: Magnetic resonance imaging NEC
AND
(
Z342: Aortic arch
OR
Z35: Cerebral artery OR Z361: Carotid artery NEC
)

Appendix 3: Example of poor performance of meta clustering (Figure S1)

Experiment C, Oxford University Hospital (OUH) Stroke with $k=5$ is an example of a meta clustering analysis that did not perform well. In contrast to the example in Figure 6 where we see good separation with distinct blocks of patients in the dendrogram, there is no clear consensus of which clustering the patients should belong. Note that the Kaplan Meier and Cox plots are provided purely for illustrative purposes, meta clustering does not optimise the separation of outcomes.

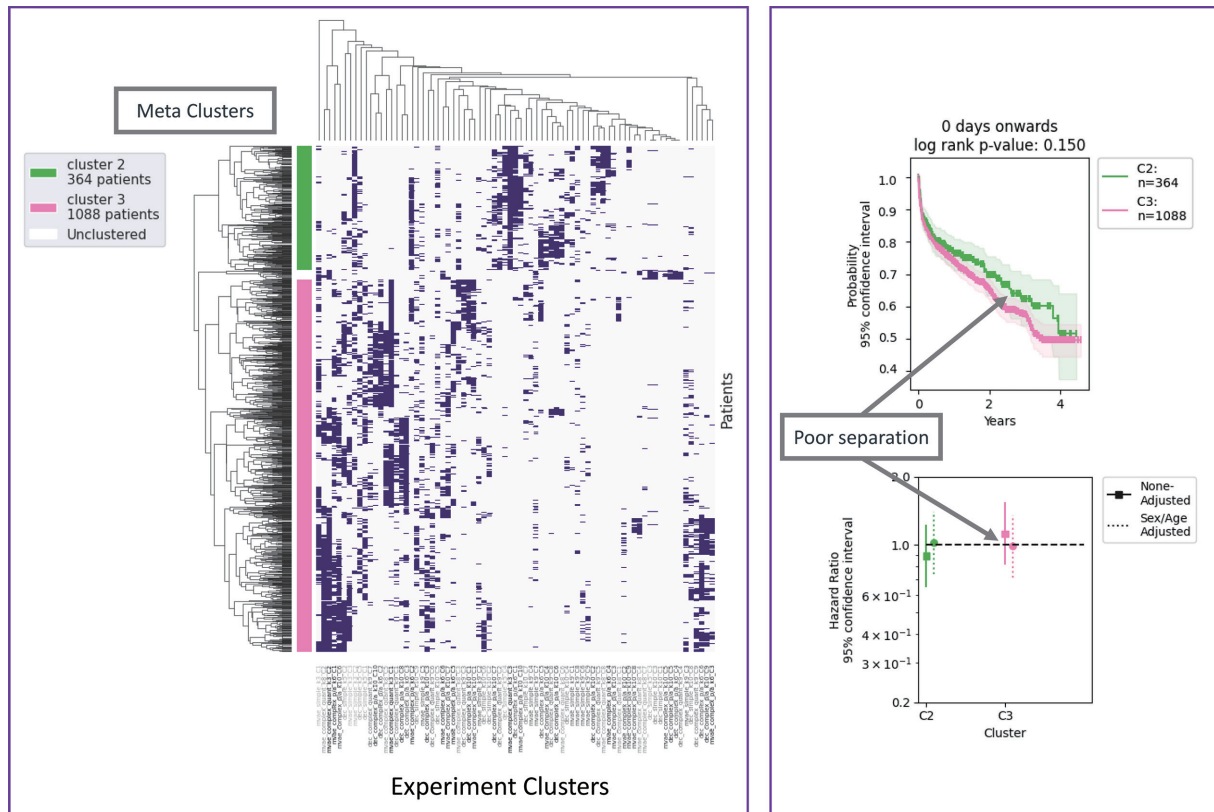


Figure S1 Experiment C, an example of poor performance of meta clustering (left) with a dendrogram heatmap and (right) survival analysis. Here the patient groups do not show significantly different outcomes.