

# Identifying antipsychotic treatment episodes using natural language processing and change-point detection

**Tao Wang**

`tao.wang@kcl.ac.uk`

King's College London

**Hamilton Morrin**

King's College London

**Ninoslav Majkic**

South London and Maudsley National Health Service (NHS) Foundation Trust

**David Codling**

South London and Maudsley National Health Service (NHS) Foundation Trust

**Risha Govind**

South London and Maudsley National Health Service (NHS) Foundation Trust

**Matthew Broadbent**

South London and Maudsley National Health Service (NHS) Foundation Trust

**Cecilia Casetta**

King's College London

**Giouliana Kadra-Scalzo**

King's College London

**Rashmi Patel**

University of Cambridge

**Angus Roberts**

King's College London

**Richard Dobson**

King's College London

**Robert Stewart**

King's College London

**Sameer Jauhar**

Imperial College London

**James MacCabe**

King's College London

---

## Research Article

**Keywords:** antipsychotic episodes, change-point detection, natural language processing, antipsychotic prescription, electronic health records

**Posted Date:** September 10th, 2026

**DOI:** <https://doi.org/10.21203/rs.3.rs-10580263/v1>

**License:** © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.  
[Read Full License](#)

**Additional Declarations:** Competing interest reported. RP has participated in Scientific Advisory Boards for Boehringer Ingelheim and Teva, has received grant funding from Janssen, and has received consulting fees from Holmusk, Akkrivia Health, Columbia Data Analytics, Clinilabs, Social Finance, Boehringer Ingelheim, Bristol Myers Squibb, Supernus, Teva and Otsuka. GKS has received research funding from Janssen Research & Development LLC and H. Lundbeck A/S. RS declares research support received in the last 36 months from GSK, and royalties from Oxford University Press. The remaining authors declare that they have no competing interests.

---

# Identifying antipsychotic treatment episodes using natural language processing and change-point detection

Tao Wang<sup>1\*</sup>, Hamilton Morrin<sup>1,2</sup>, Ninoslav Majkic<sup>1,2,3</sup>,  
David Codling<sup>2</sup>, Risha Govind<sup>1,2</sup>, Matthew Broadbent<sup>2</sup>,  
Cecilia Casetta<sup>1</sup>, Gioulana Kadra-Scalzo<sup>1</sup>, Rashmi Patel<sup>4,1</sup>,  
Angus Roberts<sup>1</sup>, Richard Dobson<sup>1,2,5,6</sup>, Robert Stewart<sup>1,2</sup>,  
Sameer Jauhar<sup>1,2,7</sup>, James MacCabe<sup>1,2</sup>

<sup>1\*</sup>Institute of Psychiatry, Psychology and Neuroscience, King’s College  
London, Denmark Hill, London, SE5 8AF, United Kingdom.

<sup>2</sup>South London and Maudsley National Health Service (NHS)  
Foundation Trust, Denmark Hill, London, SE5 8AZ, United Kingdom.

<sup>3</sup>Centre for Translational Medicine, King’s Health Partners, Counting  
House, Guy’s Hospital, London, SE1 9RT, United Kingdom.

<sup>4</sup>Department of Psychiatry, University of Cambridge, Herchel Smith  
Building, Robinson Way, Cambridge, CB2 0SZ, United Kingdom.

<sup>5</sup>Institute of Health Informatics, University College London, Euston  
Road, London, NW1 2DA, United Kingdom.

<sup>6</sup>Health Data Research UK London, University College London, Euston  
Road, London, NW1 2DA, United Kingdom.

<sup>7</sup>Department of Brain Sciences, Imperial College London,  
Commonwealth Building, London, W12 0NN, United Kingdom.

\*Corresponding author(s). E-mail(s): [tao.wang@kcl.ac.uk](mailto:tao.wang@kcl.ac.uk);

## Abstract

**Background:** Prescribing information for antipsychotic medications is predominantly recorded in unstructured clinical text, preventing large-scale analysis of treatment trajectories. We developed and validated a framework integrating natural language processing (NLP) and change-point detection (CPD) to

reconstruct longitudinal antipsychotic treatment episodes from electronic health records (EHRs).

**Methods:** We analysed EHRs from 12,530 patients with schizophrenia-spectrum disorders (ICD-10 F20–F29) at the South London and Maudsley NHS Foundation Trust (2007–2017). Antipsychotic mentions were extracted using a validated NLP pipeline and modelled as categorical time series; the Pruned Exact Linear Time (PELT) algorithm detected episode boundaries. Performance was evaluated against 113 manually annotated episodes from 50 patients. Associations between treatment-episode metrics and healthcare utilisation were assessed using generalised linear models adjusted for age at onset, gender, ethnicity, and diagnosis.

**Results:** The algorithm achieved an F1 score of 0.81 against manual annotations. Across 11,731 patients, 24,648 treatment episodes were identified. Combined antipsychotic use was detected in 31.7% of episodes (1,989 transition overlaps; 5,827 concurrent therapy). Clozapine users had more treatment episodes than non-clozapine users (mean 3.3 vs. 1.9), consistent with clinical expectations. Higher combination-prescribing rates and more pre-clozapine episodes were positively associated with hospital admissions and inpatient stays; a higher proportion of long-acting injectable (LAI) episodes was associated with reduced service use.

**Conclusions:** This NLP-CPD framework provides a scalable approach to reconstructing antipsychotic treatment trajectories from unstructured EHRs. Derived episode metrics demonstrate clinical validity and reveal associations between treatment complexity and healthcare utilisation, offering a generalisable architecture for prescribing quality monitoring and automated identification of treatment resistance.

**Keywords:** antipsychotic episodes, change-point detection, natural language processing, antipsychotic prescription, electronic health records

## 1 Background

Antipsychotic medications are the cornerstone of treatment for schizophrenia and related psychotic disorders, yet real-world prescribing frequently deviates from guideline recommendations [1, 2]. Guidelines recommend a stepwise response to inadequate treatment: optimising the dose, switching to an alternative agent, and, once treatment resistance is established after at least two unsuccessful trials, initiating clozapine; concurrent use of multiple antipsychotics (polypharmacy) is generally discouraged [3–5]. Despite these recommendations, polypharmacy occurs in 20% to 50% of cases [6–8], clozapine initiation is frequently delayed [9, 10], and medication switching is particularly common among patients with treatment-resistant schizophrenia (TRS) [11, 12] – deviations associated with poorer patient outcomes and higher service utilisation. Systematically characterising these treatment trajectories is therefore essential for understanding where and why prescribing diverges from best practice, and what consequences follow.

However, clinicians predominantly document prescribing information in unstructured clinical text rather than structured fields, which limits large-scale analysis [13, 14]. This is especially pronounced in mental health services, where narrative notes capture prescribing decisions, clinical rationale, and longitudinal context that are rarely coded in structured form [13]. Even in healthcare systems with comprehensive structured prescription registers, such as those in Nordic countries [15], identifying treatment episodes from individual prescription records remains challenging: structured data capture discrete dispensing events but lack explicit information on treatment continuity, forcing researchers to apply arbitrary gap thresholds to infer episode boundaries. Existing approaches to constructing episodes from structured data include rule-based episode construction [16, 17], waiting-time distribution inference [18], prescription-to-drug-use modelling [19], and clustering algorithms [20]; Segev et al. [21] recently validated a methodology specifically for identifying clozapine treatment periods within EHRs. However, such methods rely on predefined thresholds or assumptions that may not reflect real-world adherence or flexible prescribing practices, and are sensitive to missing or delayed records. Previous studies have therefore relied on manually compiling treatment episodes from free text [2, 4, 22] or clinician surveys [23], approaches that are labour-intensive and difficult to scale.

Recent studies have applied NLP to identify medication mentions from EHRs [12, 24], but these approaches capture discrete mentions rather than reconstructing longitudinal treatment trajectories; aggregating such mentions into episodes introduces noise from early or late references, routine medication reviews, and overlapping polypharmacy, which can obscure the true timing and continuity of prescribing. While state-of-the-art large language models (LLMs) have shown promise in clinical text analysis, with health system-scale models such as NYUTron [25] and trajectory frameworks such as EHR2Path [26] demonstrating potential for predictive tasks, their application to medication trajectory reconstruction faces substantial barriers. High operational costs and stringent privacy requirements, including the risk of data memorisation and hallucination of clinical facts, limit deployment in routine care [27]. Furthermore, LLMs frequently struggle with temporal consistency in clinical narratives, compounded by the scarcity of high-quality longitudinal training data required for complex temporal reasoning [28–30]. Consequently, mention-based extraction alone remains vulnerable to documentation noise, such as routine medication reviews or duplicate references, leaving a methodological gap that more robust temporal modelling can address.

We propose a framework that integrates NLP with change-point detection (CPD) to identify antipsychotic medication start and stop points from unstructured clinical text. NLP extracts medication mentions from EHRs, which are modelled as a time series; CPD then detects distributional shifts in that series to identify points where one treatment replaces another [31]. CPD methods have been widely adopted to identify distributional shifts in sequential data across domains such as signal processing, finance, and bioinformatics [31, 32]; prior applications to medical and behavioural time series include sequence behaviour change detection using nonhomogeneous Markov models [33] and causal discovery-driven CPD for multivariate distributions [34]. Unlike rule-based thresholding, dynamic programming frameworks such as the Pruned Exact

Linear Time (PELT) algorithm [35] enable optimal, computationally efficient  $\mathcal{O}(N)$  change-point detection without requiring the number of change points to be specified in advance. Despite these advances, applying CPD to reconstruct continuous antipsychotic treatment episodes directly from NLP-extracted EHR mentions remains unexplored.

The automated identification of treatment episodes from unstructured clinical text has direct applications in both clinical care and research. In clinical practice, such methods could support medication review systems that flag patients who may benefit from treatment optimisation. For research, they enable cohort identification based on detailed medication history criteria without requiring manual chart review. Applying this framework to EHR data from the South London and Maudsley NHS Foundation Trust (SLaM), one of the UK’s largest mental health trusts, we demonstrate a scalable approach to reconstructing antipsychotic treatment episodes. We validate the method’s performance and examine real-world treatment patterns, including switching behaviours, combined antipsychotic use, and associations with healthcare utilisation.

## 2 Methods

### 2.1 Data collection

We analysed EHRs from the South London and Maudsley NHS Foundation Trust (SLaM), one of the UK’s largest mental health providers, serving a local population of approximately 1.3 million people in South London. We accessed all data via the Clinical Record Interactive Search (CRIS) platform, a de-identified research database and governance framework that enables secure use of SLaM EHR data for research purposes [36].

As defined in our previous work [37], our study cohort comprised all patients diagnosed with schizophrenia, schizotypal, or delusional disorders (ICD-10 codes F20–F29) between 1 January 2007 and 31 December 2017. Eligible patients were those who received at least one antipsychotic prescription during this period and who were either residing within the SLaM catchment area or homeless but receiving care from SLaM services at the time of their first prescription. For each patient, we extracted all records referencing antipsychotic medication between 1 January 2007 and 1 October 2023, encompassing both structured prescription fields and unstructured clinical text, along with all hospital admission records up to the same date.

### 2.2 Medication timeline

To construct a comprehensive longitudinal profile of antipsychotic exposure, we integrated structured prescription data with unstructured medication mentions extracted via NLP. While structured pharmacy records capture dispensing details – including medication name, formulation, strength, quantity dispensed, and dispense date – they do not record medication start or end dates, making it impossible to directly infer treatment episodes from structured data alone. Free-text clinical notes provide richer information, including prescribing rationale, treatment continuity, and treatment responses such as adverse drug events; however, the temporal information derived

from NLP is inherently noisy. Clinical notes frequently reference medications outside the context of current active use: a note may discuss a patient’s historical medication, consider a hypothetical or planned future treatment, or replicate content through copy-and-paste from prior entries. Consequently, the date of a clinical document does not reliably indicate when a medication was started or discontinued, and raw NLP-extracted mentions cannot be used directly to define episode boundaries without further temporal modelling.

The NLP model has been validated in previous studies for extracting antipsychotic information from SLaM EHRs [24, 38]. The pipeline processed raw narrative notes, hospital discharge summaries, and clinical correspondence to identify antipsychotic mentions, using a combination of rule-based pattern matching and lexical analysis to map synonyms, misspellings, and proprietary drug names to a standardised pharmacological vocabulary. Once a mention was detected, the surrounding text was parsed to resolve drug attributes – such as formulation and dose – and align them with structured metadata. For each integrated medication event, the pipeline produced a structured tuple comprising the standardised medication name, formulation (oral or long-acting injectable [LAI]), recorded dose, and the date of the source document or prescription. This document date serves as an approximate temporal anchor but, as described above, does not reliably mark when a medication was started or stopped; episode boundary identification therefore requires further temporal modelling, described in the next section.

### 2.3 Episode identification

We define an antipsychotic treatment episode as a period during which a patient receives a specific antipsychotic medication, from the first recorded mention to either discontinuation or a switch to another agent. Identifying such episodes from EHR data is challenging because no structured fields explicitly record episode start and end dates, and prior approaches, relying either on structured prescription data alone [4] or on manual coding from free text [11, 37], do not scale to large cohorts.

Here, we propose a scalable informatics approach that applies CPD to time-series representations of antipsychotic records to automatically identify episode boundaries. Fig. 1 illustrates the process. For each patient, we formalised their longitudinal medication records as a discrete time series  $S = \{(t_i, m_i)\}_{i=1}^N$ , where  $t_i$  is the timestamp of the  $i$ -th clinical document and  $m_i$  is the categorical medication variable extracted via NLP. By default,  $m_i$  corresponds to the medication name, though the operationalisation of  $S$  can be adapted to other clinical phenotypes; for example,  $m_i$  can be expanded to represent medication-formulation pairs to detect switches between oral and long-acting injectable (LAI) regimens. Since CPD algorithms conventionally operate on numerical sequences, we mapped the categorical states  $m_i$  to distinct numerical values based on their first appearance in the sequence. To ensure sufficient data density for algorithm convergence, we restricted our analysis to patients with  $N \geq 5$  records.

We applied the Pruned Exact Linear Time (PELT) algorithm to detect abrupt shifts in the statistical properties of  $S$ , indicative of genuine treatment transitions. We selected PELT for its computational efficiency: unlike binary segmentation or standard dynamic programming, which scale poorly with sequence length or require the number




**Fig. 1** Data flow in antipsychotic episode identification. Unstructured clinical free text is processed using natural language processing to extract antipsychotic drug mentions. These are then combined with structured prescribing and dispensing records to create a single chronological timeline for each patient. Each point represents an antipsychotic medication record on a given date. A change-point detection algorithm analyses the temporal density of these records to identify dates when prescribing patterns change, indicated by dashed red lines.

of change points to be specified a priori, PELT incorporates a pruning mechanism that achieves  $\mathcal{O}(N)$  expected time complexity [32, 35]. Mathematically, for a sequence of length  $N$ , the algorithm identifies a set of  $K$  change points  $\tau = (\tau_0, \tau_1, \dots, \tau_K)$  with  $0 = \tau_0 < \tau_1 < \dots < \tau_K < \tau_{K+1} = N$ , by minimising the penalised cost function:

$$\min_{\tau} \sum_{i=1}^{K+1} \mathcal{C}(S_{\tau_{i-1}:\tau_i}) + \beta K, \quad (1)$$

where  $\mathcal{C}$  is a segment-specific cost function and  $\beta$  is the penalty parameter. We used a radial basis function (RBF) kernel as  $\mathcal{C}$ , which captures non-parametric distributional shifts within the categorical medication state space. The algorithm determines the optimal partition by iteratively minimising this cost function; central to its efficiency is a pruning condition whereby any candidate change point whose cumulative cost exceeds the running minimum is permanently discarded.

Episode demarcation is governed by two key hyperparameters: (1) the penalty parameter ( $\beta$ ), which controls the trade-off between detecting true medication transitions and overfitting to documentation noise; and (2) the `min_size` parameter, which enforces a minimum segment length between change points, ensuring that detected episodes reflect sustained treatment periods rather than isolated mentions.

## 2.4 Detection of combined antipsychotic use

Clinical guidelines generally recommend antipsychotic monotherapy, but multiple antipsychotics are often prescribed concurrently in practice, a pattern referred to as antipsychotic polypharmacy [8]. To characterise these patterns, we distinguished

between two types of combined use: temporary overlaps during medication switching (transition overlap) and sustained combination therapy (concurrent therapy). We applied complementary rule-based procedures to identify each pattern.

**Transition overlap (cross-titration):** When switching between antipsychotics, clinicians often overlap medications temporarily to ensure therapeutic continuity while minimising withdrawal effects. To detect these transitions, we examined a 28-day window ( $\pm 14$  days) around each detected episode changepoint. We defined a transition overlap when: (i) both the previous and new antipsychotic were mentioned at least twice within this window; and (ii) the previous antipsychotic continued to be mentioned for at least 7 days after the new antipsychotic was initiated, confirming genuine temporal overlap rather than coincidental documentation.

**Concurrent therapy (intentional combination):** Beyond transition periods, we assessed whether patients received multiple antipsychotics as sustained combination therapy within stable treatment episodes. We defined a concurrent therapy period when a second antipsychotic was mentioned at least three times during an episode, or at least once per month of the episode’s duration, a criterion designed to distinguish intentional co-prescription from short-term adjunctive use or documentation artefacts.

## 2.5 Model validation



**Fig. 2** Interface for annotating medication episodes. Top left: number of mentions of each medication over time. Top right: recorded dose over time; bars of different colours indicate different medications and red vertical lines indicate pre-annotated change points. Bottom: source data for each medication mention, including documentation date, source table, and text snippets.

To validate the CPD algorithm and optimise its hyperparameters, we manually annotated 113 treatment episodes from a random sample of 50 patients drawn from the full cohort of individuals with at least 5 antipsychotic records. To streamline the annotation process, we developed a web interface based on Elasticsearch and Kibana

within the CogStack framework [39], a clinical information extraction and retrieval system. We applied the CPD algorithm to pre-annotate medication time series for each patient; these results were ingested into an Elasticsearch database and visualised through a Kibana dashboard alongside raw structured prescribing fields and relevant text snippets (Fig. 2). When annotators identified discrepancies in episode start dates, they recorded the corrected dates in a spreadsheet.

Although we developed the bespoke annotation system for model validation, manual annotation remains labour-intensive and cannot scale to the full cohort; we therefore performed a complementary known-groups validation to assess the clinical validity of the algorithm at scale. We compared treatment episode counts between clozapine and non-clozapine users: since clozapine is indicated for TRS following inadequate response to at least two prior antipsychotic trials, clozapine users should demonstrate more treatment episodes than non-clozapine users. This known-groups comparison provides evidence that the algorithm captures clinically meaningful treatment trajectories without requiring exhaustive manual review.

## 2.6 Data analyses

Having validated the algorithm, we conducted two sets of analyses. First, we descriptively characterised the treatment patterns identified by the pipeline to assess their face validity against known clinical patterns. Second, we examined associations between the derived treatment-episode metrics and healthcare utilisation to demonstrate their downstream clinical utility.

**Descriptive characterisation of treatment patterns.** Patient characteristics were summarised, including age at first psychotic diagnosis (ICD-10 codes F20–F29), gender, ethnicity, primary diagnosis, number of hospital admissions, and total admission days. Treatment episodes were descriptively analysed to examine antipsychotic use trajectories, medication switching patterns, combined antipsychotic use, and episode count distributions.

**Associations between treatment patterns and service use.** We examined associations between antipsychotic treatment episode characteristics and healthcare utilisation, measured as the number of hospital admissions and total inpatient days. Treatment characteristics captured four dimensions: switching speed (mean episode duration per patient); recent treatment stability (duration of the most recent episode, its medication class, and formulation); treatment complexity (combination prescribing rate, defined as the proportion of episodes involving multiple antipsychotics); and treatment history (number of episodes prior to clozapine initiation; for patients who never initiated clozapine, this was their total episode count). We additionally included the number of LAI episodes and the LAI rate (proportion of episodes involving a long-acting injectable [LAI] agent) to assess the role of formulation. Associations were tested using generalised linear models (GLMs), adjusting for age at onset, gender, ethnicity, and diagnosis. All analyses were conducted in Python (version 3.9.7) using the statsmodels package (version 0.13.1).

## 3 Results

### 3.1 Cohort characteristics

The cohort comprised 12,530 patients with schizophrenia-spectrum disorders (Table 1). The mean age at first psychotic diagnosis was 42.9 years (SD = 16.6). Males comprised 57.0% of the cohort. The largest ethnic groups were Black (44.1%) and White (38.4%). Schizophrenia (ICD-10 code F20) was the most common diagnosis (51.2%), followed by other/unspecified psychotic disorders (F28–F29, 16.9%) and acute psychotic disorders (F23, 14.2%). Overall, 61.3% of patients had at least one hospital admission, with a mean of 3.2 admissions (SD = 3.3) and a mean total inpatient duration of 259.0 days (SD = 438.2).

**Table 1** Patient characteristics

| Variable                             | All patients (N = 12,530) |
|--------------------------------------|---------------------------|
| Age of onset (years), mean (SD)      | 42.9 (16.6)               |
| <b>Gender</b>                        |                           |
| Male, n (%)                          | 7,138 (57.0%)             |
| Female, n (%)                        | 5,392 (43.0%)             |
| <b>Ethnicity</b>                     |                           |
| White, n (%)                         | 4,808 (38.4%)             |
| Black, n (%)                         | 5,527 (44.1%)             |
| Asian, n (%)                         | 1,038 (8.3%)              |
| Mixed, n (%)                         | 342 (2.7%)                |
| Other, n (%)                         | 650 (5.2%)                |
| Unknown, n (%)                       | 165 (1.3%)                |
| <b>First psychotic diagnosis</b>     |                           |
| F20 (Schizophrenia), n (%)           | 6,419 (51.2%)             |
| F28–F29 (Other/Unspecified), n (%)   | 2,113 (16.9%)             |
| F23 (Acute psychotic), n (%)         | 1,781 (14.2%)             |
| F25 (Schizoaffective), n (%)         | 1,326 (10.6%)             |
| F22–F24 (Delusional), n (%)          | 813 (6.5%)                |
| F21 (Schizotypal), n (%)             | 78 (0.6%)                 |
| <b>Hospital admissions</b>           |                           |
| Patients with admissions, n (%)      | 7,679 (61.3%)             |
| Number of admissions, mean (SD)      | 3.2 (3.3)                 |
| Duration (days), mean (SD)           | 259.0 (438.2)             |
| <b>Clozapine use</b>                 |                           |
| Patients prescribed clozapine, n (%) | 2,017 (16.1%)             |

SD, standard deviation.

### 3.2 Performance of episode identification

To validate the CPD algorithm, two reviewers (TW and HM) independently annotated 113 treatment episodes from 50 patients, achieving an inter-rater agreement of

93.8%. Discrepancies were resolved through discussion with the broader team. Fig. 3 illustrates F1 scores of the PELT algorithm across parameter settings, benchmarked against manual annotations. We defined true positives as true change points with an estimated counterpart within a 14-day margin. The algorithm achieved its best performance at  $\text{min\_size}=2$  and  $\text{penalty}=7$ , yielding an F1 score of 0.81.




**Fig. 3** Change-point detection performance across kernels and hyperparameter settings. (a) Best F1 score achieved by each candidate kernel; the radial basis function (RBF) kernel was selected for its superior performance in preliminary analyses. (b) F1 scores of the PELT algorithm with the RBF kernel under various penalty and  $\text{min\_size}$  settings. Low penalty values produce numerous false change points; excessively high penalties miss true change points. When  $\text{min\_size}=n$ , all detected change points are separated by at least  $n$  samples. The algorithm was implemented using the ruptures library [32].

Error analysis identified three recurring error types: (i) *Medication restarts* ( $n=6$ ): the algorithm failed to recognise medication restarts, where a medication was stopped and later resumed, as separate episodes. (ii) *Sparse records* ( $n=4$ ): limited availability of medication records hindered episode detection. (iii) *Polypharmacy* ( $n=3$ ): the algorithm was less sensitive to changes involving transient adjunctive prescribing, such as aripiprazole added alongside risperidone to manage antipsychotic-induced hyperprolactinaemia.

### 3.3 Known-groups validation

The CPD algorithm identified 24,648 treatment episodes across 11,731 patients. Fig. 4(a) shows the distribution of treatment episodes by clozapine use. Most patients experienced relatively few episodes, with a mean of 3.3 ( $\text{SD} = 2.0$ ) among the 2,017 clozapine users with at least five medication records and 1.9 ( $\text{SD} = 1.0$ ) among the 9,714 non-clozapine users with at least five records. As clozapine is indicated for TRS following inadequate response to at least two prior antipsychotic trials, the higher episode count among clozapine users is consistent with clinical expectations, providing evidence that the algorithm captures clinically meaningful treatment trajectories.

### 3.4 Episode characteristics




**Fig. 4** Characteristics of the identified antipsychotic treatment episodes. (a) Number of antipsychotic episodes, grouped by patients prescribed with clozapine or not. (b) Mean episode duration for each medication with 95% confidence intervals (CIs). The value of  $n$  in parentheses indicates the number of patients with at least one episode of that medication. (c) Antipsychotic medication choice by episode order.

Fig. 4(b) shows episode durations by medication; oral and LAI formulations of the same medication are distinguished given differences in dosing frequency and duration of action. Clozapine had the longest mean episode duration at 76.8 months (95% CI:

73.9–79.6). Several typical antipsychotics commonly administered as LAIs, including fluphenazine, flupentixol, zuclopenthixol, and pipotiazine, also had prolonged mean episode durations.

In total, 12,917 medication switches were identified across episode sequences. Fig. 4(c) presents a heatmap of medication sequence across treatment trajectories. Olanzapine was the most common first-line antipsychotic, followed by risperidone and aripiprazole. Paliperidone was most commonly used as a second-line agent. Across all treatment stages, clozapine, aripiprazole, and zuclopenthixol appeared in the widest range of treatment orders. The most common switch was from olanzapine oral to aripiprazole oral, followed by risperidone oral to paliperidone LAI and olanzapine oral to risperidone oral (Fig. 5).




**Fig. 5** Common switch patterns between two consecutive antipsychotic episodes.

We identified combined antipsychotic use in 7,816 episodes, representing 31.7% of all episodes and 25.6% of patients. Of these, 1,989 episodes were classified as transition overlaps and 5,827 as concurrent therapy. Fig. 6 illustrates the most frequent medication combinations within each category. The most common transition-overlap patterns involved switches from risperidone to paliperidone and from olanzapine to aripiprazole. Concurrent-therapy combinations most frequently included olanzapine prescribed alongside other agents, particularly aripiprazole and risperidone.




**Fig. 6** Common combinations of concurrently prescribed antipsychotics. Combinations are shown separately for transition overlap (where the first medication was switched to the second) and concurrent therapy (where a second medication was added to the primary medication).




**Fig. 7** Regression results between service use and patient characteristics.

### 3.5 Treatment episodes and service use

Fig. 7 shows regression results for associations between treatment-episode patterns and service use, adjusted for age at first psychotic diagnosis, gender, ethnicity, and primary diagnosis. Both the combination prescribing rate and the number of treatment episodes prior to clozapine initiation were positively associated with higher hospital admission counts and longer cumulative inpatient stays. The absolute number of LAI episodes was positively associated with service use, whereas a higher proportion of LAI episodes relative to total episodes was associated with reduced service use. Quicker switching (shorter mean episode duration) and stable use (longer duration of the final episode) were not significantly associated with service use. Compared with the reference diagnosis of schizophrenia (F20), patients with schizoaffective disorder (F25) had more frequent hospital admissions and longer cumulative admission stays, whereas other diagnostic categories generally showed lower admission rates and shorter durations. Relative to the White reference group, patients of Black ethnicity had higher service utilisation. Female patients had slightly shorter admission durations than males, although admission counts did not differ significantly by gender.

## 4 Discussion

This study demonstrates a scalable biomedical informatics pipeline for reconstructing longitudinal antipsychotic treatment trajectories directly from unstructured clinical notes. By integrating NLP for entity extraction with CPD for temporal sequence modelling, we address a longstanding methodological limitation in EHR research: the absence of structured prescribing continuity data.

The NLP-CPD pipeline achieves an F1 score of 0.81 by separating genuine clinical transitions from documentation artefacts, a task that manual abstraction [2, 22] and structured-data methods reliant on arbitrary refill-gap assumptions [16] cannot perform at scale. To the best of our knowledge, this is the first study to apply CPD to reconstruct antipsychotic treatment episodes from NLP-extracted EHR narratives; prior CPD applications to medical time series have addressed broader behavioural change detection [33, 34] rather than medication episode reconstruction. PELT’s tunable penalty parameter  $\beta$  is central to this: it suppresses false change points arising from isolated medication mentions while remaining sensitive to sustained treatment transitions. This modular approach is generalisable; it requires only categorical time-series inputs and can be adapted to reconstruct treatment trajectories for other chronic conditions beyond schizophrenia. Potential applications include prescribing quality dashboards, automated cohort phenotyping, and clinical decision support systems that flag patients meeting criteria for treatment resistance.

Beyond methodological validation, our findings replicate and extend established clinical knowledge. We confirmed known patterns of antipsychotic prescribing, including the predominance of olanzapine, risperidone, and aripiprazole as first- and second-line agents [40–42]. The pipeline distinguished sustained concurrent therapy from expected cross-titration overlaps. For example, transitions from risperidone to its active metabolite paliperidone [1] – identifying combined antipsychotic use in 31.7%

of episodes. This distinction has direct clinical relevance: cross-titration overlaps represent planned pharmacological transitions, whereas concurrent therapy reflects a deliberate prescribing decision that guidelines generally discourage.

We used generalised linear modelling to demonstrate how algorithmic outputs, derived metrics of treatment complexity, transition frequency, and combination prescribing, can be integrated into downstream statistical analyses. A higher combination prescribing rate and a greater number of pre-clozapine episodes were each associated with higher hospital admission counts and longer inpatient stays. This aligns with prior literature indicating that delayed treatment optimisation burdens healthcare resources [12, 43], and our approach enables these associations to be measured continuously from routine clinical data rather than through retrospective manual cohort abstraction. Future integration of this framework into live clinical environments could operationalise these longitudinal patterns in real time, converting clinical notes into alerts for emerging treatment resistance and elevated healthcare utilisation risk.

Several limitations must be acknowledged. First, the pipeline’s accuracy depends on the upstream NLP entity extraction phase; persistent false negatives or misattributed contexts (e.g., hypothetical or negated mentions) can propagate into the CPD model, occasionally leading to misclassified episode boundaries. Second, applying this algorithm across fragmented healthcare networks, such as between primary and specialised care, may introduce temporal gaps that violate the continuity assumptions of the CPD model. Third, the NLP and CPD models were parameterised on a single mental health trust’s terminology and may require recalibration for different documentation styles. Fourth, as with all observational studies, the data preclude causal interpretation of the associations reported.

Future research should explore integrating LLMs to perform direct temporal grounding, potentially bypassing rule-based entity extraction pipelines, or extending the CPD framework to support online change-point detection as new clinical notes are recorded. Emerging methods in causal inference for EHR data, such as synthetic data generation [44, 45], could use these algorithmic outputs to simulate counterfactual treatment trajectories.

## 5 Conclusions

This study demonstrates a scalable informatics framework that integrates NLP with change-point detection to reconstruct longitudinal antipsychotic treatment trajectories from unstructured EHR text. Applied to 12,530 patients with schizophrenia-spectrum disorders, the pipeline identified 24,648 treatment episodes, characterised switching patterns and combined antipsychotic use, and revealed that higher combination-prescribing rates and a greater number of pre-clozapine episodes are associated with increased hospital admissions and longer inpatient stays. This data-driven approach overcomes the limitations of manual chart review and standalone entity extraction, offering a generalisable architecture that can support downstream clinical decision-making and enable continuous monitoring of population-level prescribing behaviours.

**Abbreviations.** CPD, change-point detection; CRIS, Clinical Record Interactive Search; EHR, electronic health record; GLM, generalised linear model; ICD-10, International Classification of Diseases, 10th Revision; LAI, long-acting injectable; LLM, large language model; NLP, natural language processing; PELT, Pruned Exact Linear Time; RBF, radial basis function; SD, standard deviation; SLaM, South London and Maudsley NHS Foundation Trust; TRS, treatment-resistant schizophrenia.

**Supplementary information.** Not applicable.

## Declarations

**Ethics approval and consent to participate.** This study was conducted using the Clinical Record Interactive Search (CRIS) platform, which holds ethical approval as an anonymised data resource from the NHS Health Research Authority South Central – Oxford C Research Ethics Committee (reference 23/SC/0257). The present secondary analysis of de-identified electronic health record data was reviewed and approved by the CRIS Oversight Committee. All methods were carried out in accordance with the relevant guidelines and regulations. The study used only de-identified data accessed within a secure firewall, did not involve any contact with patients, and individual informed consent was not applicable.

**Consent for publication.** Not applicable.

**Availability of data and materials.** The code developed for this study, including the NLP-CPD pipeline and episode identification framework, is publicly available on GitHub (<https://github.com/wtgme/med-episode-detect>). The clinical data underlying this study are not publicly available because they contain de-identified patient records held within the CRIS platform at the South London and Maudsley NHS Foundation Trust; access is available to researchers via application to the CRIS Oversight Committee.

**Competing interests.** RP has participated in Scientific Advisory Boards for Boehringer Ingelheim and Teva, has received grant funding from Janssen, and has received consulting fees from Holmusk, Akivia Health, Columbia Data Analytics, Clinilabs, Social Finance, Boehringer Ingelheim, Bristol Myers Squibb, Supernus, Teva and Otsuka. GKS has received research funding from Janssen Research & Development LLC and H. Lundbeck A/S. RS declares research support received in the last 36 months from GSK, and royalties from Oxford University Press. The remaining authors declare that they have no competing interests.

**Authors' contributions.** TW: conceptualisation, data curation, formal analysis, investigation, methodology, software, validation, visualisation, writing – original draft, writing – review and editing. HM: conceptualisation, formal analysis, investigation, methodology, validation, writing – review and editing. NM: formal analysis, methodology, validation, writing – review and editing. DC: conceptualisation, investigation, validation, writing – review and editing. RG: data curation, writing – review and editing. MB: conceptualisation, data curation, writing – review and editing. CC: validation, writing – review and editing. GKS: data curation, validation, writing – review

and editing. RP: methodology, validation, writing – review and editing. AR: funding acquisition, supervision, writing – review and editing. RD: funding acquisition, resources, supervision, writing – review and editing. RS: supervision, validation, writing – review and editing. SJ: conceptualisation, investigation, resources, validation, writing – review and editing. JM: conceptualisation, data curation, investigation, methodology, resources, validation, writing – review and editing. All authors read and approved the final manuscript.

**Acknowledgements.** This study was funded by: (a) Maudsley Charity, (b) the National Institute for Health Research (NIHR) Biomedical Research Centre at South London and Maudsley NHS Foundation Trust and King’s College London; (c) Health Data Research UK, which is funded by the UK Medical Research Council, Engineering and Physical Sciences Research Council, Economic and Social Research Council, Department of Health and Social Care (England), Chief Scientist Office of the Scottish Government Health and Social Care Directorates, Health and Social Care Research and Development Division (Welsh Government), Public Health Agency (Northern Ireland), British Heart Foundation and Wellcome Trust; (d) The BigData@Heart Consortium, funded by the Innovative Medicines Initiative-2 Joint Undertaking under grant agreement No. 116074. This Joint Undertaking receives support from the European Union’s Horizon 2020 research and innovation programme and EFPIA; it is chaired, by DE Grobbee and SD Anker, partnering with 20 academic and industry partners and ESC; and (e) The National Institute for Health Research University College London Hospitals Biomedical Research Centre. These funding bodies had no role in the design of the study, collection and analyses. The views expressed are those of the author(s) and not necessarily those of the NHS, the NIHR or the Department of Health.

RS is part-funded by: i) the National Institute for Health Research (NIHR) Biomedical Research Centre at the South London and Maudsley NHS Foundation Trust and King’s College London; ii) the National Institute for Health Research (NIHR) Applied Research Collaboration South London (NIHR ARC South London) at King’s College Hospital NHS Foundation Trust; iii) the DATAMIND HDR UK Mental Health Data Hub (MRC grant MR/W014386); iv) the UK Prevention Research Partnership (Violence, Health and Society; MR-VO49879/1), an initiative funded by UK Research and Innovation Councils, the Department of Health and Social Care (England) and the UK devolved administrations, and leading health research charities; v) the NIHR HealthTech Research Centre in Brain Health. TW is also part-funded by the King’s Together fund and MHaPS Early Career Researcher Award. NM was supported by the King’s Health Partners Centre for Translational Medicine. The views expressed are those of the author(s) and not necessarily those of King’s Health Partners. HM is a National Institute for Health and Care Research Academic Clinical Fellow (ACF-2023-17-008). AR is funded by the UK Medical Research Council (grant DATAMIND - The Data Hub for Mental Health Informatics Research and Development, number MR/Z504816/1, and grant RE-STAR: Regulating Emotions - Strengthening Adolescent Resilience, number MR/W002493/1). RP has received grant funding from the National Institute for Health and Care Research (NIHR301690).

## References

- [1] Akindele T, De Alwis K. Real-World Prescribing Patterns of Long-Acting Injectable Antipsychotics in Australian Psychiatric Inpatients: Trends, Clinical Outcomes, and Substance Use Prevalence. *Drugs-Real World Outcomes*. 2025;12(3):489–501.
- [2] Patel R, Brinn A, Irving J, Chaturvedi J, Gudiseva S, Correll CU, et al. Oral and long-acting injectable antipsychotic discontinuation and relationship to side effects in people with first episode psychosis: a longitudinal analysis of electronic health record data. *Therapeutic Advances in Psychopharmacology*. 2023;13:20451253231211575.
- [3] Howes OD, McCutcheon R, Agid O, De Bartolomeis A, Van Beveren NJ, Birnbaum ML, et al. Treatment-resistant schizophrenia: treatment response and resistance in psychosis (TRRIP) working group consensus guidelines on diagnosis and terminology. *American Journal of Psychiatry*. 2017;174(3):216–229.
- [4] Smidt Cd, Haffmans J, Hoencamp E. Antipsychotics switching strategies in real life: a longitudinal study in clinical practice. *The European Journal of Psychiatry*. 2012;26(1):41–49.
- [5] Shah P, Iwata Y, Brown EE, Kim J, Sanches M, Takeuchi H, et al. Clozapine response trajectories and predictors of non-response in treatment-resistant schizophrenia: a chart review study. *European Archives of Psychiatry and Clinical Neuroscience*. 2020;270:11–22.
- [6] Correll CU, Gallego JA. Antipsychotic polypharmacy: a comprehensive evaluation of relevant correlates of a long-standing clinical practice. *The Psychiatric clinics of North America*. 2012;35(3):661.
- [7] Gallego JA, Bonetti J, Zhang J, Kane JM, Correll CU. Prevalence and correlates of antipsychotic polypharmacy: a systematic review and meta-regression of global and regional trends from the 1970s to 2009. *Schizophrenia research*. 2012;138(1):18–28.
- [8] Højlund M, Köhler-Forsberg O, Gregersen AT, Rohde C, Mellentin AI, Anhøj SJ, et al. Prevalence, correlates, tolerability-related outcomes, and efficacy-related outcomes of antipsychotic polypharmacy: a systematic review and meta-analysis. *The Lancet Psychiatry*. 2024;11(12):975–989.
- [9] Üçok A, Çikrikçili U, Karabulut S, Salaj A, Öztürk M, Tabak Ö, et al. Delayed initiation of clozapine may be related to poor response in treatment-resistant schizophrenia. *International clinical psychopharmacology*. 2015;30(5):290–295.

- [10] Shah P, Iwata Y, Plitman E, Brown EE, Caravaggio F, Kim J, et al. The impact of delay in clozapine initiation on treatment outcomes in patients with treatment-resistant schizophrenia: a systematic review. *Psychiatry research*. 2018;268:114–122.
- [11] Khandker R, Shepherd J, Chekani F, Qureshi Z, Bailey H, Berry M, et al. Atypical treatment switches in schizophrenia patients: drivers and associated outcomes. *Neuropsychiatric Disease and Treatment*. 2022;18:1057.
- [12] Kadra G, Stewart R, Shetty H, MacCabe JH, Chang CK, Kesserwani J, et al. Antipsychotic polypharmacy prescribing and risk of hospital readmission. *Psychopharmacology*. 2018;235(1):281–289.
- [13] Newby D, Taylor N, Joyce DW, Winchester LM. Optimising the use of electronic medical records for large scale research in psychiatry. *Translational psychiatry*. 2024;14(1):232.
- [14] Roe D, Mazor Y, Gelkopf M. Patient-reported outcome measurements (PROMs) and provider assessment in mental health: a systematic review of the context of implementation. *International Journal for Quality in Health Care*. 2022;34(Supplement\_1):ii28–ii39.
- [15] Højlund M, Pottegård A, Johnsen E, Kroken RA, Reutfors J, Munk-Jørgensen P, et al. Trends in utilization and dosing of antipsychotic drugs in Scandinavia: comparison of 2006 and 2016. *British journal of clinical pharmacology*. 2019;85(7):1598–1606.
- [16] Gardarsdottir H, Souverein PC, Egberts TC, Heerdink ER. Construction of drug treatment episodes from drug-dispensing histories is influenced by the gap length. *Journal of clinical epidemiology*. 2010;63(4):422–427.
- [17] Lo CWH, Gillett AC, Iveson MH, Kamp M, Fabbri C, Wong WLE, et al. Antidepressant Switching as a Proxy Phenotype for Drug Nonresponse: Investigating Clinical, Demographic, and Genetic Characteristics. *Biological Psychiatry Global Open Science*. 2025;5(4):100502.
- [18] Hallas J, Gaist D, Bjerrum L. The waiting time distribution as a graphical approach to epidemiologic measures of drug utilization. *Epidemiology*. 1997;8(6):666.
- [19] Tanskanen A, Taipale H, Koponen M, Tolppanen AM, Hartikainen S, Ahonen R, et al. From prescription drug purchases to drug use periods—a second generation method (PRE2DUP). *BMC medical informatics and decision making*. 2015;15(1):21.
- [20] Pazzagli L, Liang D, Andersen M, Linder M, Khan AR, Sessa M. Rationale and performances of a data-driven method for computing the duration of

- pharmacological prescriptions using secondary data sources. *Scientific Reports*. 2022;12(1):6245.
- [21] Segev A, Govind R, Oloyede E, Morrin H, Jewell A, Jones R, et al. Developing a validated methodology for identifying clozapine treatment periods in electronic health records. *BMC psychiatry*. 2024;24(1):584.
  - [22] Keating D, McWilliams S, Boland F, Doyle R, Behan C, Strawbridge J, et al. Prescribing pattern of antipsychotic medication for first-episode psychosis: a retrospective cohort study. *BMJ open*. 2021;11(1):e040387.
  - [23] Mace S, Taylor D. Reducing the rates of prescribing high-dose antipsychotics and polypharmacy on psychiatric inpatient and intensive care units: results of a 6-year quality improvement programme. *Therapeutic advances in psychopharmacology*. 2015;5(1):4–12.
  - [24] Kadra G, Stewart R, Shetty H, Jackson RG, Greenwood MA, Roberts A, et al. Extracting antipsychotic polypharmacy data from electronic health records: developing and evaluating a novel process. *BMC psychiatry*. 2015;15(1):1–7.
  - [25] Yang J, et al. Health system-scale language models are all-purpose clinical predictive engines. *Nature*. 2023;618(7964):332–340.
  - [26] Lee JH, et al. From EHRs to Patient Pathways: Scalable Modeling of Longitudinal Health Trajectories with LLMs. *arXiv preprint arXiv:250604831*. 2025;.
  - [27] Agrawal A, et al. The Privacy Paradox: Scaling LLMs in Healthcare under Strict Regulatory Frameworks. *Nature Medicine*. 2024;30(4):912–925.
  - [28] Kougia V, et al. Challenges and Opportunities in Temporal Information Extraction from Clinical Text. In: *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*; 2024. p. 112–128.
  - [29] Kruse J, et al. Logical Consistency and Temporal Reasoning in Clinical Large Language Models. *Journal of Biomedical Informatics*. 2025;150:104567.
  - [30] Mahajan D, Liang JJ, Tsou CH, Uzuner Ö. Overview of the 2022 n2c2 shared task on contextualized medication event extraction in clinical notes. *Journal of biomedical informatics*. 2023;144:104432.
  - [31] Fan Y, Lu X. An online Bayesian approach to change-point detection for categorical data. *Knowledge-Based Systems*. 2020;196:105792.
  - [32] Truong C, Oudre L, Vayatis N. Selective review of offline change point detection methods. *Signal Processing*. 2020;167:107299.
  - [33] Leyli-Abadi M, Samé A, Oukhellou L, Cheifetz N, Mandel P, Féliers C, et al. Online common change-point detection in a set of nonstationary categorical time

- series. *Neurocomputing*. 2021;439:176–196.
- [34] Gao S, Addanki R, Yu T, Rossi R, Kocaoglu M. Causal Discovery-Driven Change Point Detection in Time Series. In: *International Conference on Artificial Intelligence and Statistics*; 2024. p. 3754–3762.
  - [35] Killick R, Fearnhead P, Eckley IA. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*. 2012;107(500):1590–1598.
  - [36] Stewart R, Soremekun M, Perera G, Broadbent M, Callard F, Denis M, et al. The South London and Maudsley NHS foundation trust biomedical research centre (SLAM BRC) case register: development and descriptive data. *BMC psychiatry*. 2009;9(1):1–12.
  - [37] Kadra-Scalzo G, Fonseca de Freitas D, Agbedjro D, Francis E, Ridler I, Pritchard M, et al. A predictor model of treatment resistance in schizophrenia using data from electronic health records. *PLoS One*. 2022;17(9):e0274864.
  - [38] Perera G, Broadbent M, Callard F, Chang CK, Downs J, Dutta R, et al. Cohort profile of the South London and Maudsley NHS Foundation trust biomedical research centre (SLAM BRC) case register: current status and recent enhancement of an electronic mental health Record-derived data resource. *BMJ open*. 2016;6(3).
  - [39] Jackson R, Kartoglu I, Stringer C, Gorrell G, Roberts A, Song X, et al. CogStack-experiences of deploying integrated information retrieval and extraction services in a large National Health Service Foundation Trust hospital. *BMC medical informatics and decision making*. 2018;18:1–13.
  - [40] Tiihonen J, Mittendorfer-Rutz E, Majak M, Mehtälä J, Hoti F, Jedenius E, et al. Real-world effectiveness of antipsychotic treatments in a nationwide cohort of 29 823 patients with schizophrenia. *JAMA psychiatry*. 2017;74(7):686–693.
  - [41] Correll CU, Rummel-Kluge C, Corves C, Kane JM, Leucht S. Antipsychotic combinations vs monotherapy in schizophrenia: a meta-analysis of randomized controlled trials. *Schizophrenia bulletin*. 2009;35(2):443–457.
  - [42] Vadieli N, El-Ali J, Delaune J, Wild C, Liu YS. Patterns and predictors of oral antipsychotic prescribing in adult patients with schizophrenia. *Exploratory Research in Clinical and Social Pharmacy*. 2022;6:100148.
  - [43] John AP, Ko EKF, Dominic A. Delayed initiation of clozapine continues to be a substantial clinical concern. *The Canadian Journal of Psychiatry*. 2018;63(8):526–531.

- [44] Butzin-Dozier Z, Ji Y, Wang LC, Anzalone AJ, Hurwitz E, Patel RC, et al. Causal Inference via Electronic Health Records in the National Clinical Cohort Collaborative: Challenges and Solutions in Long COVID Research. medRxiv. 2025;.
- [45] Bouvarel B, Lapidus N. CO11.1 - Causal inference by generation of synthetic data from electronic health records as an alternative to inverse probability of treatment weighting: a simulation study. Journal of Epidemiology and Population Health. 2025;.