



King's Research Portal

DOI:

[10.1186/s13326-020-00220-2](https://doi.org/10.1186/s13326-020-00220-2)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Viani, N., Kam, J., Yin, L., Bittar, A., Dutta, R., Patel, R., Stewart, R., & Velupillai, S. (2020). Temporal Information Extraction from Mental Health Records to Identify Duration of Untreated Psychosis. *Journal of biomedical semantics*, 11(2), Article 2. <https://doi.org/10.1186/s13326-020-00220-2>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

1 Temporal Information Extraction from Mental Health Records 2 to Identify Duration of Untreated Psychosis

3 Natalia Viani¹, Joyce Kam¹, Lucia Yin¹, André Bittar¹, Rina Dutta^{1,2}, Rashmi Patel^{1,2}, Robert
4 Stewart^{1,2}, Sumithra Velupillai¹

5

6 ¹Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, UK;

7 ²South London and Maudsley NHS Foundation Trust, London, UK;

8

9 **Email addresses:**

10 Natalia Viani, natalia.viani@kcl.ac.uk

11 Joyce Kam, joyce.kam@kcl.ac.uk

12 Lucia Yin, lucia.yin@kcl.ac.uk

13 André Bittar, andre.bittar@kcl.ac.uk

14 Rina Dutta, rina.dutta@kcl.ac.uk

15 Rashmi Patel, rashmi.patel@kcl.ac.uk

16 Robert Stewart, robert.stewart@kcl.ac.uk

17 Sumithra Velupillai, sumithra.velupillai@kcl.ac.uk

18

19 **Corresponding Author:**

20 Natalia Viani

21 Institute of Psychiatry, Psychology and Neuroscience, King's College London

22 De Crespigny Park, SE5 8AF, London, UK

1 **Abstract**

2 **Background.** Duration of untreated psychosis (DUP) is an important clinical construct in the field of
3 mental health, as longer DUP can be associated with worse intervention outcomes. DUP estimation
4 requires knowledge about *when* psychosis symptoms first started (symptom onset), and when psychosis
5 treatment was initiated. Electronic health records (EHRs) represent a useful resource for retrospective
6 clinical studies on DUP, but the core information underlying this construct is most likely to lie in free
7 text, meaning it is not readily available for clinical research. Natural Language Processing (NLP) is a
8 means to addressing this problem by automatically extracting relevant information in a structured form.
9 As a first step, it is important to identify appropriate documents, i.e., those that are likely to include the
10 information of interest. Next, temporal information extraction methods are needed to identify time
11 references for early psychosis symptoms. This NLP challenge requires solving three different tasks:
12 time expression extraction, symptom extraction, and temporal “linking”. In this study, we focus on the
13 first step, using two relevant EHR datasets.

14 **Results.** We applied a rule-based NLP system for time expression extraction that we had previously
15 adapted to a corpus of mental health EHRs from patients with a diagnosis of schizophrenia (first
16 referrals). We extended this work by applying this NLP system to a larger set of documents and patients,
17 to identify additional texts that would be relevant for our long-term goal, and developed a new corpus
18 from a subset of these new texts (early intervention services). Furthermore, we added normalized value
19 annotations (“2011-05”) to the annotated time expressions (“May 2011”) in both corpora. The finalized
20 corpora were used for further NLP development and evaluation, with promising results (normalization
21 accuracy 71%-86%). To highlight the specificities of our annotation task, we also applied the final
22 adapted NLP system to a different temporally annotated clinical corpus.

23 **Conclusions.** Developing domain-specific methods is crucial to address complex NLP tasks such as
24 symptom onset extraction and retrospective calculation of duration of a preclinical syndrome. To the
25 best of our knowledge, this is the first clinical text resource annotated for temporal entities in the mental
26 health domain.

27 **Keywords**

28 Natural Language Processing; Electronic Health Records; Temporal Information Extraction;
29 Schizophrenia; Mental Health;

1 **1. Background**

2 In the field of mental health, investigating the duration of untreated symptoms in relation to intervention
3 outcomes represents an important research topic [1]. For patients with a diagnosis of schizophrenia, for
4 example, the duration of untreated psychosis (DUP) is a widely used construct in research cohorts,
5 defined as the period of time between first symptom onset and initiation of adequate treatment. A longer
6 DUP has been linked to poorer cognitive function at the time of first presentation [2], and subsequently
7 predicts more severe symptoms and greater social and functional impairment [3]. Therefore, routine
8 identification of DUP across large clinical groups is a crucial step for prognostic monitoring and could
9 form the basis for nested interventions to improve both clinical and functional outcomes at a service or
10 population level. Electronic health records (EHRs) represent a valuable resource for large-scale
11 retrospective clinical studies, as they contain a large amount of routinely collected patient data. In
12 mental health services, however, relevant information on DUP is documented mainly in text fields and
13 cannot therefore be easily analyzed automatically. To make this information available for computational
14 analysis and clinical research, Natural Language Processing (NLP) methods can be used [4,5].

15 EHR databases are typically large and complex, containing data for all patients in a clinical catchment
16 area. Each patient may have varying levels of contact with health services, forming different trajectories
17 and sets of EHRs. As a first step to develop real-world NLP applications using EHR data, it is important
18 to identify appropriate documents for NLP development, i.e., those that are likely to include the
19 information of interest. In the case of DUP extraction, relevant information would typically be
20 documented in EHRs for patients with a psychotic disorder diagnosis: particularly in initial clinical
21 assessment notes or paragraphs describing the patient's previous clinical history or early psychosis
22 symptoms, recorded around the time of first presentation and assessment. Once a relevant set of EHR
23 documents (corpus) is defined, NLP techniques can be used to identify mentions of relevant symptoms
24 as well as the associated temporal details. This represents a temporal information extraction challenge,
25 which typically requires three different steps: (i) the identification of time expressions (*May 1st*), (ii)
26 the identification of relevant concepts, such as symptoms (*hallucinations*) and treatments

1 (*antipsychotic*), and (iii) the identification of temporal relations between entity pairs (*hallucinations*
2 BEFORE *antipsychotic*), also known as temporal “linking”.

3 In recent years, manually annotated corpora and methods for temporal information extraction have been
4 developed mainly based on the TimeML specification language, which was originally created for the
5 general NLP domain (e.g., newspaper text) [6]. In the clinical domain, few gold standard corpora have
6 been created and made available for temporal NLP development. Moreover, they address temporal
7 modelling on a general level, without being driven by a specific clinical use-case such as DUP. Within
8 the Informatics for Integrating Biology and the Bedside (i2b2) project, 310 de-identified discharge
9 summaries from an intensive care unit were manually annotated for events, time expressions, and
10 temporal relations [7]. This corpus was used in the 2012 i2b2 Challenge on temporal relation extraction,
11 which required participants to develop NLP solutions to automatically extract these temporal elements
12 [8]. In the oncology domain, Styler and colleagues created an annotated corpus of 1,254 de-identified
13 EHR notes, including both clinical and temporal information (the THYME corpus) [9]. This corpus
14 consists of two types of EHR notes: clinical notes, which often include clearly delineated sections
15 describing past and present events, and pathology reports, which contain a detailed analysis of
16 specimens (taken at a single moment in time). Subsets of the THYME corpus were reused in different
17 NLP challenges, among which Clinical TempEval 2015 and 2016 focused on temporal information
18 extraction (440 and 591 documents, respectively) [10,11]. In both i2b2 2012 and THYME, four main
19 TimeML types of time expressions are defined: Date (e.g. *2011, yesterday*), Duration (e.g. *three years,*
20 *one week*), Frequency/Set (e.g. *daily, twice a week*), and Time (e.g. *10am, the morning*). The THYME
21 corpus also includes two additional types specific to the oncology domain: PrePostExp (expressions
22 indicating Pre- and Post-operational concepts, e.g. *postoperative day #4*) and Quantifier (e.g. *twice, four*
23 *times*).

24 Time expression extraction involves not only identifying textual spans representing time references, but
25 also assigning a standardized value to them (*normalization*), which is crucial for anchoring clinical
26 concepts on a patient timeline. In general, normalizing time expressions is a challenging task, especially
27 due to the usage of relative (e.g. *two days before*) and underspecified (e.g. *at 9pm*) time expressions

1 [12], as well as imprecise time expressions (e.g. *several weeks*) [13]. In the i2b2 2012 challenge, for
2 example, 37% of all combined Date and Time expressions were relative: when evaluating the top 10
3 performing systems on these, the normalization accuracy dropped from around 0.67 to 0.32 [12].
4 Various strategies have been employed to tackle this difficulty. The SUTime system normalizes all
5 relative expressions by comparison to the document creation time (DCT) [14]. The HeidelTime system
6 uses rules and heuristics based on the document domain (e.g., news, narrative) to select anchor time
7 expressions [15]. When evaluated on the i2b2 2012 dataset, teams using SUTime and HeidelTime
8 achieved a normalized value accuracy of 0.54 and 0.6, respectively on the test set [8]. For this particular
9 task, the highest normalization accuracies were obtained by using regular expressions (0.73) [16] and
10 combining rules with supervised classifiers (0.72) [17]. To specifically deal with relative and
11 incomplete expressions, Sun et al. combined multi-label classifiers for anchor points and anchor
12 relations, achieving an improvement on relative and incomplete time expression value normalization
13 from 0.45 (the top score on these expressions) to 0.54 [12]. As another interesting approach, the
14 TimeNorm system relies on a synchronous context free grammar, with showed promising normalization
15 performances on general domain datasets: on the TempEval 2013 dataset [18], for example, the reported
16 normalization accuracy is 0.82 for TimeNorm, and 0.79 for HeidelTime [19].

17 Compared to other clinical domains, mental health records are characterized by a greater amount and
18 variety of narrative portions, describing clinical histories and health assessments without necessarily
19 relying on pre-defined structured sections containing a temporal anchor point. In this framework,
20 relevant temporal information on DUP (e.g., associated to symptom onset or treatment initiation) is not
21 always well represented by temporal models relying on TimeML. In previous work, to further
22 investigate this aspect, we annotated a corpus of mental health documents for time expression spans
23 and types, with a specific focus on patients with a diagnosis of schizophrenia [20]. Comparing this
24 annotated corpus to two related works (the i2b2 2012 challenge and Clinical TempEval 2016), we found
25 that mental health documents are much longer, with an average of 3,974 tokens per document (vs. 574
26 for i2b2 2012 and 931 for Clinical TempEval 2016), and contain a larger variety of temporal references
27 (including information taken from structured forms that is not relevant to the patient's clinical history).

1 Moreover, while notes in other domains can contain semi-structured date information (e.g. admission,
2 discharge, or section dates), mental health texts often include various paragraphs describing both past
3 and current events related to the patient. As regards information on DUP, from that work, we concluded
4 that age-related expressions like *at age 8* or *in his teens* are important to temporally anchor the first
5 occurrence of psychosis symptoms (Figure 1). To capture these cases, we introduced a new type of time
6 expression, *Age_related*, which represented 8.9% of all annotations in that corpus. We also evaluated,
7 adapted and refined an existing rule-based system, SUTime [14], to extract time expressions in this
8 domain.

9 **Figure 1 here**

10 **Figure 1. Example of clinical text describing the onset of psychosis symptoms.**

11 Our long-term goal is to automatically extract from mental health notes all the elements needed for the
12 generation of DUP data on a large patient cohort. To address this long-term goal, we have previously
13 developed a corpus annotated with time expressions and adapted a time expression extraction system
14 (SUTime) [20] to be used for temporal NLP development in the mental health domain - in particular to
15 support DUP extraction [21]. Here, we extend this work with the following main contributions:

- 16 1. *Large-scale use of an adapted time expression extraction system for automated identification*
17 *of relevant EHR documents:* We applied the adapted SUTime system on a large set of EHR
18 documents from early intervention services for psychosis, to identify additional documents that
19 would be relevant for calculating DUP.
- 20 2. *Time normalization annotation:* We extended our previous annotation effort on first referral
21 EHR documents by adding the normalized values of the time expressions we had annotated in
22 that corpus [20]. Furthermore, we annotated a new subset of documents from early intervention
23 services, manually identifying time expression spans and values (without marking types).
- 24 3. *Automatic normalization:* We used the finalized annotated corpus to further refine our NLP
25 time expression extraction system, with a focus on the normalization task. We also assessed the
26 performance of this system on the i2b2 2012 dataset, to highlight key differences between the
27 annotation tasks.

1 2. Methods

2 2.1 Dataset

3 In this study, we used data from the Clinical Record Interactive Search (CRIS) resource [22], which is
4 derived from the EHR system adopted by the South London and Maudsley National Health Service
5 (NHS) Foundation Trust (SLaM). Within CRIS, de-identified patient EHRs are rendered available for
6 research within a robust governance framework. These include information from both structured fields
7 and free text fields in the source EHR, the latter including case notes and clinical correspondence with
8 automatically blanked-out identifiers [23]. In general, textual documents do not follow a specific
9 structure and contain different types of patient information, e.g., past history, family history,
10 examination results, and drug prescriptions. Moreover, although a document date can be retrieved from
11 CRIS, it does not necessarily correspond to the actual document creation time (DCT), as there might be
12 a temporal gap between the document creation and its upload to the system.

13 Relevant documents for calculating DUP would include a description of the patient’s clinical history or
14 an assessment of early psychosis symptoms. Following advice from domain experts, two CRIS-derived
15 datasets were considered:

- 16 - *First referral documents for schizophrenia patients.* In previous work, we extracted 52 early
17 documents¹ for patients with a diagnosis of schizophrenia [20]. For each patient, we considered
18 the longest document, on the assumption that this first, long referral document would include
19 the richest description of the patient’s clinical history. We call this dataset the *First referral*
20 *corpus*.
- 21 - *Documents from early intervention services for psychosis.* We extended the set of documents
22 to be used for annotation and NLP development, considering attachments from early
23 intervention services for psychosis, i.e., mental health services that support people who are
24 experiencing untreated psychosis for the first time. We call this dataset the *Early intervention*
25 *corpus*.

¹ Documents were written within 3 months of patients’ first presentation to mental health services.

1 As a first step to extend our *First referral* corpus, we extracted all clinical correspondence attachments
2 from early intervention services for psychosis. For each patient, we considered the documents written
3 within a 3-month window from the team’s acceptance date (36.6k documents for 4,166 patients).
4 Because some of these documents were relatively short or contained forms or questionnaires rather than
5 descriptive assessments, we filtered the initial set by only keeping longer documents, which were more
6 likely to include information on symptom onset. More specifically, we calculated the document and
7 line length on the entire document collection, and kept documents with length (in terms of characters)
8 greater than the 50th percentile (2000 characters) and average line length greater than the 25th percentile
9 (30 characters).

10 **2.2 Large-scale use of an adapted time expression extraction system for relevant EHR document** 11 **identification**

12 Through manual review of a small subset from the *Early intervention corpus*², we identified two criteria
13 that could be useful in filtering documents. Firstly, we observed that documents containing relevant
14 symptom keywords (e.g. *hallucinations*, *delusions*) were more likely to include passages on a patient’s
15 clinical history and thus potentially symptom onset information. For this analysis, we used a list of 26
16 keywords that was defined by two psychiatrists³ (the list is publicly available [24]). Secondly, by
17 applying our adapted version of SUTime on the same texts, we found that the number of extracted time
18 expressions was in general higher in the documents mentioning symptom onset. Therefore, we applied
19 this system on all available early intervention documents, using these criteria as additional filtering
20 steps in order to maximize the amount of relevant documents for calculating DUP and minimize the
21 number of irrelevant documents for costly manual review. More specifically, we kept all the documents
22 containing at least one psychosis symptom keyword, and more than 5 time expressions, as we estimated
23 this threshold would allow us to retain only the most relevant documents (based on our manual review).

² We manually reviewed a sample of 70 documents, marking relevant sentences such as “*his mother first expressed concerns about his behaviour when he was 8 years old*”.

³ Starting from a comprehensive list of psychosis symptoms, two psychiatrists reviewed the most frequent terms in the *Early intervention corpus*, thus agreeing on the final list of 26 keywords to be used as a filter.

1 2.3 Manual time expression normalization

2 The *First referral* corpus consisted of documents that we had previously annotated for five types of
3 time expressions (without normalized values): Date, Duration, Time, Frequency, and Age_related [20].

4 One of the findings in that study, through inter-annotator agreement (IAA) analysis on type
5 classification, was that distinguishing between Date and Duration caused the most annotation
6 disagreements (42%). For example, the expression “*last week*” was interpreted as a point in time (Date)
7 by one annotator, and as a period of time (Duration) by another.

8 In the *First referral* corpus, we extended the time expression annotations by adding normalized values
9 (“value” attribute), mostly following the TimeML specification language. To simplify the assignment
10 of normalized values on difficult instances of Date/Duration time expressions, we instructed the
11 annotators to keep prepositions within the marked text spans. In summary, the normalization rules were
12 the following:

- 13 - Dates were normalized as “YYYY-MM-DD”, “YYYY-MM”, or “YYYY”, depending on their
14 granularity. We also considered vague expressions such as *in the past* (value = PAST_REF)
15 and *now* (value = PRESENT_REF).
- 16 - Times were normalized as “YYYY-MM-DDThh:mm”, when the date was available, or
17 “XXXX-XX-XXThh:mm”, otherwise. We also considered times related to a time of the day,
18 e.g. *in the morning* (TMO) and *at night* (TNI).
- 19 - Durations and frequencies were normalized in the form “P (T)? Digit Granularity”, e.g., *for*
20 *four years* (value = P4Y), *for 15 minutes* (value = PT15M).
- 21 - For durations marking a range of time, we added this information in the normalized value,
22 considering two main cases. Expressions preceded by temporal prepositions like “since” and
23 “until” (e.g. *since 2007*, *until a year ago*), which were common in our dataset, were normalized
24 using the strings SIN and UNT (e.g. SIN2007, UNT2018). For explicit ranges denoted by a
25 start point and an end point, e.g. “2009-2012”, we used the Duration type (rather than two

1 separate Dates) with a brackets notation, e.g. value = (2009, 2012). This choice was made to
2 keep the information on both the actual duration and the two endpoints.

3 - Similar to TimeML, vague durations (e.g. *for hours*) were normalized with a placeholder X
4 (value = PXH).

5 - Age_related expressions were normalized in a similar way to standard durations, distinguishing
6 between two different cases: expressions indicating the current age of the patient (e.g. *a 45 year*
7 *old man*, value = P45Y), and those referring to a previous point in time (e.g. *when he was 15*,
8 value = A15Y). In addition, vague references like *when he was a child* were normalized with
9 specific categorical values: CHILD_REF, SCHOOL_REF, TEENS_REF, ADULT_REF,
10 UNI_REF, and OTHER_REF (for other cases).

11 To assess the impact of type classification on IAA values, the *Early intervention* corpus was only
12 annotated with time expression spans and normalized values, without requiring specifying a type.

13 **2.4 Automated time expression extraction system refinement**

14 Once the two corpora were annotated, we used the time expressions that were annotated with the same
15 value by both annotators to further adapt and refine our automated time expression extraction system.

16 We added post-processing rules on top of the default SUTime normalized values, with a focus on
17 Age_related expressions and durations including prepositions. To develop these rules, we manually
18 reviewed annotations from the *First referral* corpus development set used in our previous work (10 first
19 referral documents, Table 1).

20 **2.5 Evaluation**

21 To measure the IAA for time expression extraction, we computed the lenient F1 score, where a true
22 positive (*match*) is defined as a textual span identified by both annotators (allowing overlapping spans).

23 For the normalization task, we calculated accuracy (*acc*) only on matching spans, counting the

1 proportion of expressions normalized with the same value⁴. The system’s normalization performance
2 was evaluated on these expressions, using the same metric (value accuracy for system true positives).
3 To highlight the specificities of our annotation task, we also applied our adapted time expression
4 extraction system to the i2b2 2012 test set, manually analyzing errors and key differences in the two
5 corpora.

6 **3. Results**

7 **3.1 Large-scale use of adapted time expression extraction system**

8 To obtain relevant documents for DUP extraction, the early intervention services dataset was filtered at
9 different levels (Figure 2) [21]. Note that the order of applying these filtering steps is not important.

10 **Figure 2 here**

11 **Figure 2. Filtering steps from EHR documents related to early psychosis intervention services.**

12 We first removed short documents (length < 2000 characters) and texts consisting of short lines (average
13 line length < 30 characters), which resulted in 16,318 documents for 3,819 patients (i.e. around 55% of
14 the initial documents were excluded, while 92% of the patients were retained). From these, we only
15 retained the documents containing at least one relevant psychosis symptom keyword, which resulted in
16 9,901 documents for 3,466 patients (i.e. around 40% of the documents were filtered out, while 91% of
17 the patients remained represented). Figure 3 shows the number of symptom mentions and time
18 expressions found in these documents (normalized counts⁵): texts containing many temporal
19 expressions are more likely to also include relevant psychosis symptom keywords. By also applying the
20 SUTime-based filtering step, we obtained 9,779 documents for 3,433 patients (i.e. only about 2% of
21 documents were further filtered out, retaining 99% of patients). We then randomly selected 20 of these
22 patients for time expression and normalization annotation (49 documents).

23 **Figure 3 here**

24 **Figure 3. Psychosis symptom keyword and time expression counts in the early intervention services dataset.**

⁴ For relative expressions like *one year ago*, we did not consider “granularity” differences as errors (e.g., given a document creation time of 2012-01-01, *one year ago* could be normalized as “2011” or “2011-01-01”).

⁵ All numbers were normalized to the range 0-1.

1

2 3.2 Manual time expression normalization

3 For the manual annotation task, the *First referral* corpus was pre-annotated with time expressions: in
4 our previous study, the IAA on textual spans was 77% (lenient F1 score). The *Early intervention* corpus,
5 on the contrary, was not pre-annotated. The resulting IAA on textual spans was 85%.

6 Table 1 shows the results of manual annotation for normalized values on both corpora. For the *First*
7 *referral* corpus, we use the same data split as in previous work: development (*dev*), validation (*valid*),
8 test (*test*). For the *Early intervention* corpus, we present results on two batches of 10 patients each
9 (*batchA*, *batchB*). For both corpora, we report the number of documents, annotated time expressions
10 (per annotator), overlapping time expressions, time expressions with the same normalized value, and
11 the normalized value accuracy (the IAA measure).

12 **Table 1.** Manual annotation results on the two EHR corpora (*First referral* and *Early intervention*)

Corpus	Batch	Documents (# tokens)	All annotations (A1, A2)	Overlapping annotations	Same value	IAA (acc)
First referral	dev	10 (49K)	932, 972	913	768	0.84
First referral	valid	23 (83K)	1455, 1475	1429	1254	0.88
First referral	test	19 (74K)	1119, 1159	1100	927	0.84
Early intervention	batchA	14 (18K)	435, 391	353	300	0.85
Early intervention	batchB	35 (57K)	867, 822	714	600	0.84

13 Manual annotation results on the two EHR corpora (*First referral* and *Early intervention*) divided into
14 development (*dev*), validation (*valid*) and test sets, and batches (*batchA* and *batchB*), respectively. IAA: Inter-
15 annotator agreement; A1/A2: annotators 1 and 2.

16 In the development set, the most frequent type of disagreement was in the assignment of normalized
17 values for relative expressions, such as *2 yrs back* (61/145), where it was hard to identify the anchor
18 time in the text. Other disagreements involved the confusion between DCT and “PRESENT_REF” (e.g.
19 *at this time*), non-standard dates (e.g., *week 3, over the weekend*), and time-of-the-day expressions (e.g.
20 *at night*).

3.3 Automated time expression value normalization system refinement

In the proposed automated time expression extraction system, time expression spans are first extracted with the adapted version of SUTime. Post-processing rules are then applied in order to improve the normalization step. Table 2 shows the performance of the developed system in normalizing values. The “reference standard” column represents the reference annotations, i.e., matched expressions where both annotators marked the same normalized value. The “TPs” column represents the time expressions that were correctly found by SUTime, which are used to compute the accuracy of the normalized values in the final system. First, we applied the adapted system as developed in previous work (“System 1” column). Then, we created post-processing rules in order to improve the normalization step (“System 2” column).

Table 2. Automated time expression normalization results on the two EHR corpora (*First referral* and *Early intervention*)

Corpus	Batch	Reference standard	TPs	System1 (value acc)	System2 (value acc)
First referral	dev	768	686	0.77	0.86
First referral	valid	1254	1115	0.76	0.80
First referral	test	927	828	0.66	0.71
Early intervention	batchA	300	272	0.76	0.86
Early intervention	batchB	600	556	0.82	0.86

Automated time expression extraction results (normalized values) on the two EHR corpora (*First referral* and *Early intervention*), divided into development (dev), validation (valid) and test sets, and batches (batchA and batchB), respectively. Accuracy values are reported on overlapping annotations (TPs) for both the first system (System1) and its refined version (System2).

The post-processing rules were developed to handle the following normalization values: 1) Age-related expressions, and 2) Duration expressions marking a range of time. In both cases, we defined specific keywords to be searched for in the corresponding string, e.g., “childhood”, “adolescence”. Moreover, we relied on keywords and regular expressions to disambiguate particular cases (e.g., current age of the patient vs. expressions referring to the past).

As shown in Table 2, the refinements performed in System2 were useful to improve normalization results, especially for the development set in the *First referral* corpus (accuracy from 0.77 to 0.86) and

1 batches A and B in the *Early intervention* corpus (accuracy from 0.76-0.82 to 0.86). However, the
2 improvement measured in the validation and test sets was lower (accuracy from 0.76 to 0.80, and from
3 0.66 to 0.71, respectively). To further investigate this, we analyzed IAA values and system accuracy
4 per time expression type. The results of this analysis are shown in Table 3. Column “System2 (acc)”
5 indicates that the major drop in performance was due to incorrect normalization of Time type
6 expressions, for both the validation and the test sets. To determine the number of Time errors related to
7 an underspecified Date part, we recalculated accuracy (acc*) by considering only the “Thh:mm” portion
8 of the values normalized as “YYYY-MM-DDThh:mm”. For example, for the string *6 pm sharp*, the
9 different values “2011-03-14T18:00” and “2011-03-13T18:00” would be considered as a match (value
10 = T18:00). As shown in Column “System2 (acc*)”, this led to much higher accuracies also for the Time
11 type (results marked in bold).

12 **Table 3.** Automated time expression normalization results on the *First referral* corpus, divided per time
13 expression type.

Batch	Type	IAA (matches)	IAA (acc)	System2 (TPs)	System2 (acc)	System2 (acc*)
dev	Date	572	0.84	427	0.93	0.93
	Time	77	0.87	65	0.51	0.88
	Duration	137	0.82	102	0.74	0.74
	Frequency	58	0.95	52	0.92	0.92
	Age_related	69	0.81	40	0.93	0.93
valid	Date	845	0.91	705	0.85	0.85
	Time	128	0.79	100	0.27	0.64
	Duration	209	0.77	147	0.84	0.84
	Frequency	123	0.98	101	0.95	0.95
	Age_related	124	0.81	62	0.73	0.73
test	Date	554	0.92	482	0.82	0.82
	Time	156	0.78	116	0.09	0.78
	Duration	192	0.72	128	0.77	0.77
	Frequency	90	0.72	48	0.83	0.83
	Age_related	108	0.86	54	0.80	0.80

14 Automated time expression extraction results (normalized values) on the *First referral* corpus (dev, valid, test),
15 divided per time expression type. Results are presented in terms of inter-annotator agreement (IAA), system raw
16 accuracy (System2 acc) and system relaxed accuracy (System acc*), where expressions with type Time are
17 evaluated only on the “Thh:mm” portion.

1 To gain more insight into system normalization performance, we also conducted a manual analysis of
2 the 239 normalization value errors found in the *First referral* test set (from the 828 TPs, see Table 2).
3 The majority of errors (110 expressions, 46%) originated from one particular document. This document
4 had an unusual format with a EHR system-enforced structure with several mini-sections from some
5 type of automated form with regular updates from an inpatient stay which included different paragraph
6 dates – the system did not correctly use these as anchor dates, thus resulting in erroneous references for
7 relative times (80/110) and dates (30/110). The second most frequent type of error was due to an
8 incorrect structured DCT data field from the CRIS system (31 expressions, 13%) – in this case,
9 normalized values were correctly extracted, but the provided structured DCT did not match what was
10 written in the document. Other common errors were caused by relative expressions referred to
11 previously mentioned dates (25 expressions, 10%), full dates not correctly recognized, e.g. *Friday 5*
12 *October 2012* (13 expressions, 5%), and periods not recognized, e.g. *since age 3 months* (13
13 expressions, 5%). Other important areas of improvement involve interval/ranges, e.g. *from 10pm to 4am*
14 (9 expressions) and imprecise references, e.g. *for many years* (6 expressions). The remaining errors (32
15 expressions) did not form consistent patterns.

16 **3.4 Application of time expression extraction system to i2b2 2012 corpus**

17 To further investigate the specificity of our corpus and annotation guidelines as compared to previous
18 work in temporal NLP, we also applied our adapted time expression extraction system to the i2b2 2012
19 corpus, to analyze differences in time expression annotations and extraction performance. By applying
20 the system on the 2012 i2b2 test set (120 documents), we obtained 0.71 F1-score and 0.47 value
21 normalization accuracy, which is lower than the best performing systems in the 2012 i2b2 challenge:
22 the top results for time expression extraction/normalization were 0.9/0.73, using regular expression
23 pattern match and systematic reasoning [16], and 0.91/0.72, combining a CRF model and a context-free
24 grammar algorithm [17]. When looking at the false negatives, two major differences were observed.
25 First, the way dates are usually written in U.K. and U.S. clinical institutions is usually different
26 (YYYY/MM/DD vs. YYYY/DD/MM), which impacts the system performance both in terms of
27 extraction (full dates might not be recognized) and normalization. Second, the 2012 i2b2 corpus

1 included annotations for time expressions related to clinical events, e.g. *at the time of discharge*, or
2 *HD#2* (hospital day #2), which our corpus does not contain. Moreover, differently from the i2b2 2012
3 corpus, we annotated and implemented time expression extraction and normalization rules in our
4 adapted SUTime for age-related expressions and imprecise references (e.g. present, past), which
5 account for 15% and 32% of all false positives in the i2b2 2012 corpus.

6 **4. Discussion**

7 In this study, we have made progress in addressing an ongoing challenge of automatically extracting
8 DUP from mental health narratives. As a first step, we extended a previously annotated corpus of 52
9 first referral texts by including documents from early intervention services for psychosis. To keep only
10 the documents that were relevant for our use-case, we used the output of a temporal information
11 extraction system that we had adapted to the mental health domain. We then retained 49 of these
12 documents (for 20 patients) for further annotation and NLP development. Both document subsets (first
13 referrals and early intervention services) were annotated for time expression normalized values, mostly
14 based on the TimeML specification language. As a final step, the manual annotations were used to
15 refine our temporal information extraction system. The developed annotation guidelines are publicly
16 available [24].

17 When applying NLP techniques to real-world clinical use-cases, selecting a suitable corpus for methods
18 development is crucial. This is particularly true for complex problems such as DUP extraction, where
19 the aim is to identify very specific information on a patient-level: this information could be documented
20 only in a small portion of all patient-related texts, and it could be written in a variety of different ways.
21 In our case, utilizing the developed temporal extraction system on a large dataset in combination with
22 specific psychosis symptom keywords was useful to identify the documents containing information on
23 the patient's clinical history. This dataset is currently being analyzed for temporal relations between the
24 identified time expressions and relevant symptom keywords, with the aim of capturing psychosis
25 symptom onset information, which can then be used for DUP calculation.

1 Automatically extracting symptom onset for DUP calculation requires, among other NLP steps, the
2 identification and normalization of temporal expressions. This is a challenging task in the mental health
3 domain, especially due to the abundance of relative expressions, such as *three years ago*, and imprecise
4 age-related references, such as *when he was a child*. Referring relative expressions to the right anchor
5 point is additionally difficult in EHR data, due to underlying procedures for document creation/upload
6 – where the actual DCT is not always available. Moreover, the same document could contain different
7 section/paragraph times that do not necessarily correspond to the stated document date. As for imprecise
8 age-related references, we used categorical values for normalization (e.g. CHILD_REF or
9 TEENS_REF). For subsequent DUP calculation, these would need to be associated to specific temporal
10 ranges - relying on general world knowledge and shared definitions, e.g. (A13Y, A19Y) for
11 TEENS_REF.

12 Despite the identified challenges, the availability of an automated system for capturing time expressions
13 represents an essential step to anchor clinical concepts as accurately as possible. In the case of DUP
14 extraction, the proposed system will be helpful to link the first onset of symptoms to the corresponding
15 time period - even if the exact temporal reference is not exactly identifiable. Within the CRIS resource,
16 the document/section dates - as written in the text - and the corresponding structured fields are not likely
17 to differ by more than one month. On the contrary, early symptom onset is often documented in an
18 imprecise way, which could lead to a more inaccurate estimation (with a year or more error). For this
19 reason, we believe that differences in written vs. structured DCTs will not have a major impact on our
20 approach for DUP calculation. Similarly, using a range notation to identify a patient's period of life will
21 be helpful to give an estimate of onset dates, where the actual value to be used for DUP calculation can
22 be defined depending on the specific study. For example, it might be important to distinguish between
23 current, recent and historical onsets (e.g., longer than 1 year) – keeping each mention explicitly related
24 to a certain level of uncertainty.

25 In this study, we considered two subsets of documents from EHRs: first referral documents for patients
26 with a diagnosis of schizophrenia and documents from early intervention services for psychosis. The
27 first referral dataset was pre-annotated with adjudicated time expressions (including types), while the

1 early intervention services dataset was not pre-annotated (in this case, only normalized values were
2 required). The IAA on textual spans was higher in this second dataset compared to our previous work
3 (85% vs. 77%), while the agreement on normalized values remained comparable (84%-85% vs. 84%-
4 88%). This indicates that the annotation guidelines developed in previous work were useful even when
5 applied to a different set of documents.

6 As regards the system error analysis performed on the i2b2 2012 corpus, we noticed some key
7 differences between both the corpora and the time expression annotations. First, our adapted version
8 of SUTime missed some full or underspecified dates, partly due to the different date formats that are
9 used in U.S. and U.K. clinical notes. Moreover, we annotated and extracted age-related and imprecise
10 temporal references, as these expressions were deemed as useful for contextual analysis of psychosis
11 symptom mentions. Another important observation concerns the types of documents in the two corpora:
12 the 2012 i2b2 corpus consists of discharge summaries only, while our corpus contains a variety of
13 clinical document types.

14 Despite the inherent complexity of our normalization task, the preliminary time expression extraction
15 and normalization system we developed provides reasonable performance. Future directions will
16 concern the normalization of relative Time expressions, where anchor times are available inside the
17 text, as well as further improvement on Duration extraction and normalization. The first issue could be
18 addressed by changing the anchor date for each expression (in a similar way that the HeidelTime system
19 deals with this), while the second issue could be partly improved by adding and refining rules.
20 Moreover, we will investigate whether all mentioned dates/times are actually useful for clinical timeline
21 reconstruction. For example, there are a number of documents (especially in *First referral* validation
22 and test) which include “structured” section times that are not actually related to any clinically relevant
23 information.

24 This study has some limitations. First, as regards the large-scale application of the adapted SUTime
25 system, we did not verify whether the documents that were filtered out could still be useful for our long-
26 term goal of extracting DUP information. Also, for documents that were excluded in the first step
27 (length and average line length filters), we did not check the presence of symptom keywords and

1 temporal expressions –we will investigate the impact of this in a future extension. Furthermore, for this
2 particular use-case, in addition to accurately finding anchor points in time, appropriate psychosis
3 symptom keywords are essential. Here, we have used a predefined set of terms developed by domain
4 experts, which of course might be too restrictive. We are therefore also looking at methods to
5 automatically extend these vocabularies using data-driven methods [25], and will investigate whether
6 this could impact our proposed document filtering approach as well as downstream automated temporal
7 reasoning steps. As observed, using the time expression count on top of the symptom keyword count
8 did not identify many additional entries to be removed. To further assess the utility of the two filters
9 (separately and in combination), we will further review a sample of these texts, to gain more knowledge
10 on how both time expression and psychosis symptom keyword information is documented within the
11 mental health EHR and how this relates to symptom onset information. As another limitation, the
12 annotations in the *Early intervention* corpus were not manually adjudicated. For this reason,
13 normalization results - which were evaluated on overlapping time expressions where the annotators
14 marked the same value - might slightly change when evaluated on a larger set of adjudicated time
15 expressions. Finally, in this study we only considered rule-based approaches for time expression
16 extraction and normalization. In the future, we plan to investigate supervised machine learning methods
17 and more data-driven approaches.

18 **5. Conclusions**

19 Extracting DUP information from free text is an important step to improve large-scale research in
20 mental health using the increasing volumes of EHR data currently accumulating. This NLP challenge
21 requires different steps, for which developing domain-specific resources and methods is essential. In
22 this study, we identified two relevant sets of EHR documents for our use-case, and annotated them for
23 time expression spans and values - which are needed for an accurate representation of a patient's
24 timeline and, by extension, calculating DUP. We also adapted a rule-based system for time expression
25 extraction and normalization in this domain. To the best of our knowledge, this is the first clinical data
26 resource annotated for temporal entities in the mental health domain.

27

1 **List of abbreviations**

2 DUP: duration of untreated psychosis

3 EHR: electronic health record

4 NLP: natural language processing

5 **Declarations**

- 6
 - Ethics approval and consent to participate

7 The database used in this study has ethical approval for secondary analysis (Oxford REC C,
8 reference 18/SC/0372).

- 9
 - Consent for publication

10 Not applicable.

- 11
 - Availability of data and material

12 The EHR-derived datasets generated and/or analyzed during the current study are not publicly
13 available due to governance regulations. However, there are procedures in place to provide
14 researchers with controlled access to CRIS, and thus the source data are available on request
15 and subject to access approval.

- 16
 - Competing interests

17 The authors declare that they have no competing interests.

- 18
 - Funding

19 NV and SV have received support from the Swedish Research Council (2015-00359), Marie
20 Skłodowska Curie Actions, Cofund, Project INCA 600398.

21 RS, RD, RP and SV are part-funded by the National Institute for Health Research (NIHR)
22 Biomedical Research Centre at South London and Maudsley NHS Foundation Trust and

1 King's College London. RS, RP and SV are also part-funded by the Medical Research
2 Council (MRC) Mental Health Data Pathfinder Award to King's College London. RS
3 receives funding from an NIHR Senior Investigator Award.

4 RP has received support from a Medical Research Council (MRC) Health Data Research UK
5 Fellowship (MR/S003118/1) and a Starter Grant for Clinical Lecturers (SGL015/1020)
6 supported by the Academy of Medical Sciences, The Wellcome Trust, MRC, British Heart
7 Foundation, Arthritis Research UK, the Royal College of Physicians and Diabetes UK.

8 ● Authors' contributions

9 NV designed and implemented the study under the supervision of SV. RS and RP provided the
10 clinical knowledge and guided the study from a clinical perspective. JK and LY were the
11 manual annotators. NV and SV prepared the first draft of the manuscript. AB and RD provided
12 feedback and critically revised the manuscript. All authors revised and approved the final
13 version of the manuscript.

14 ● Acknowledgements

15 We would like to acknowledge the organisers of the Ninth International Workshop on Health
16 Text Mining and Information Analysis (LOUHI 2018).

17 **Reference List**

- 18 1. Kisely S, Scott A, Denney J, Simon G. Duration of untreated symptoms in common mental
19 disorders: association with outcomes. *The British Journal of Psychiatry*. 2006;189(1):79–80.
- 20 2. Lappin JM, Morgan KD, Morgan C, Dazzan P, Reichenberg A, Zanelli JW, et al. Duration of
21 untreated psychosis and neuropsychological function in first episode psychosis. *Schizophr Res*.
22 2007 Sep;95(1–3):103–10.
- 23 3. Hill M, Crumlish N, Clarke M, Whitty P, Owens E, Renwick L, et al. Prospective relationship of
24 duration of untreated psychosis to psychopathology and functional outcome over 12 years.
25 *Schizophr Res*. 2012 Nov;141(2–3):215–21.
- 26 4. Meystre SM, Savova GK, Kipper-Schuler KC, Hurdle JF. Extracting information from textual
27 documents in the electronic health record: a review of recent research. *Yearb Med Inform*.
28 2008;128–44.

- 1 5. Wang Y, Wang L, Rastegar-Mojarad M, Moon S, Shen F, Afzal N, et al. Clinical information
2 extraction applications: A literature review. *Journal of Biomedical Informatics*. 2018 Jan 1;77:34–
3 49.
- 4 6. Pustejovsky J, Castano JM, Ingria R, Sauri R, Gaizauskas RJ, Setzer A, et al. TimeML: Robust
5 specification of event and temporal expressions in text. *New directions in question answering*.
6 2003;3:28–34.
- 7 7. Sun W, Rumshisky A, Uzuner O. Annotating Temporal Information in Clinical Narratives. *J*
8 *Biomed Inform*. 2013 Dec;46(0).
- 9 8. Sun W, Rumshisky A, Uzuner O. Evaluating temporal relations in clinical text: 2012 i2b2
10 Challenge. *J Am Med Inform Assoc*. 2013 Sep;20(5):806–13.
- 11 9. Styler IV WF, Bethard S, Finan S, Palmer M, Pradhan S, de Groen PC, et al. Temporal annotation
12 in the clinical domain. *Transactions of the Association for Computational Linguistics*.
13 2014;2:143–154.
- 14 10. Bethard S, Derczynski L, Savova G, Pustejovsky J, Verhagen M. SemEval-2015 Task 6: Clinical
15 TempEval. *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval*
16 *2015)*. 2015;806–814.
- 17 11. Bethard S, Savova G, Chen W-T, Derczynski L, Pustejovsky J, Verhagen M. Semeval-2016 task
18 12: Clinical TempEval. *Proceedings of the 10th International Workshop on Semantic Evaluations*
19 *(SemEval-2016)*. 2016;1052–1062.
- 20 12. Sun W, Rumshisky A, Uzuner O. Normalization of relative and incomplete temporal expressions
21 in clinical narratives. *J Am Med Inform Assoc*. 2015 Sep;22(5):1001–8.
- 22 13. Tissot H, Del Fabro MD, Derczynski L, Roberts A. Normalisation of imprecise temporal
23 expressions extracted from text. *Knowledge and Information Systems*. 2019;1–34.
- 24 14. Chang AX, Manning CD. SUTIME: A Library for Recognizing and Normalizing Time
25 Expressions. *LREC*. 2012;
- 26 15. Strötgen J, Gertz M. Heideltime: High quality rule-based extraction and normalization of
27 temporal expressions. In: *Proceedings of the 5th International Workshop on Semantic Evaluation*.
28 2010. p. 321–324.
- 29 16. Sohn S, Waghlikar KB, Li D, Jonnalagadda SR, Tao C, Komandur Elayavilli R, et al.
30 Comprehensive temporal information detection from clinical text: medical events, time, and
31 TLINK identification. *J Am Med Inform Assoc*. 2013;20(5):836–42.
- 32 17. Xu Y, Wang Y, Liu T, Tsujii J, Chang EI-C. An end-to-end system to identify temporal relation
33 in discharge summaries: 2012 i2b2 challenge. *J Am Med Inform Assoc*. 2013 Oct;20(5):849–58.
- 34 18. UzZaman N, Llorens H, Derczynski L, Allen J, Verhagen M, Pustejovsky J. SemEval-2013 Task
35 1: TempEval-3: Evaluating Time Expressions, Events, and Temporal Relations. In: *Second Joint*
36 *Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the*
37 *Seventh International Workshop on Semantic Evaluation (SemEval 2013)*. p. 1–9.
- 38 19. Bethard S. A Synchronous Context Free Grammar for Time Normalization. In: *Proceedings of*
39 *the 2013 Conference on Empirical Methods in Natural Language Processing*. 2013. p. 821–826.

- 1 20. Viani N, Yin L, Kam J, Alawi A, Bittar A, Dutta R, et al. Time Expressions in Mental Health
2 Records for Symptom Onset Extraction. In: Proceedings of the Ninth International Workshop on
3 Health Text Mining and Information Analysis. 2018. p. 183–92.
- 4 21. Viani N, Kam J, Yin L, Verma S, Stewart R, Patel R, et al. Annotating Temporal Relations to
5 Determine the Onset of Psychosis Symptoms. *Studies in Health Technology and Informatics*.
6 2019;264:418–22.
- 7 22. Perera G, Broadbent M, Callard F, Chang C-K, Downs J, Dutta R, et al. Cohort profile of the
8 South London and Maudsley NHS Foundation Trust Biomedical Research Centre (SLaM BRC)
9 Case Register: current status and recent enhancement of an Electronic Mental Health Record-
10 derived data resource. *BMJ Open*. 2016 Mar 1;6(3):e008721.
- 11 23. Fernandes AC, Cloete D, Broadbent MT, Hayes RD, Chang C-K, Jackson RG, et al.
12 Development and evaluation of a de-identification procedure for a case register sourced from
13 mental health electronic records. *BMC Med Inform Decis Mak*. 2013 Jul 11;13:71.
- 14 24. Viani N, Velupillai S. Project repository [Internet]. Available from:
15 <https://github.com/medesto/temporal-information-extraction-DUP>
- 16 25. Viani N, Patel R, Stewart R, Velupillai S. Generating Positive Psychosis Symptom Keywords
17 from Electronic Health Records. In: Proceedings of the 17th Conference on Artificial Intelligence
18 in Medicine (AIME 2019). 2019. p. 298–303.

19

20 **Figure Legend**

21 **Figure 1**

22 *Title:* Example of clinical text describing the onset of psychosis symptoms.

23 *Legend:* The example includes two structured dates (visit date and birth date) and four time expressions that are
24 written in the text (“when he was 8 years old”, “oct 2009”, “since his teens”, “today”). As shown in the figure,
25 time expressions can be normalized and placed on a timeline in order to reconstruct patient trajectories.

26 **Figure 2**

27 *Title:* Filtering steps from EHR documents related to early psychosis intervention services.

28 *Legend:* First, we retain documents with length and average line length (`avg_line_length`) greater than a certain
29 threshold. Then, we keep documents including at least one psychosis symptom keyword (from a list of predefined
30 keywords). Finally, we retain documents containing more than five time expressions (as automatically extracted
31 by a rule-based system).

32 **Figure 3**

33 *Title:* Psychosis symptom keyword and time expression counts in the early intervention services dataset.

34 *Legend:* The x-axis represents the number of documents obtained after applying length, average line length, and
35 psychosis symptom keyword filters (9,901). The y-axis represents normalized counts for psychosis symptom

- 1 keywords (blue) and automatically extracted time expressions (orange), normalized to the range 0-1. Texts
- 2 containing many temporal expressions are more likely to also include relevant psychosis symptom keywords.