Mining and Modelling Temporal Clinical Data

Henk Harkema, Andrea Setzer, Rob Gaizauskas, Mark Hepple
University of Sheffield, UK
{harkema, andrea, robertg, hepple}@dcs.shef.ac.uk

Richard Power
Open University
R.Power@open.ac.uk

Jeremy Rogers
University of Manchester
jeremy.e.rogers@manchester.ac.uk

Abstract

The Clinical e-Science Framework (CLEF) demonstrator runs Information Extraction technology over textual, narrative patient notes to assemble repositories of clinical patient data for the purposes of biomedical research and clinical care. Since many important medical events in the course of a patient's treatment are mentioned in multiple documents and most documents will only include partial descriptions of these events, constructing a coherent and complete summary of a patient's history – what we call a *patient chronicle* - requires an information integration step over the output of Information Extraction. In this paper we describe and evaluate an approach to information integration which is based on mining narrative patient notes for temporal properties of medically relevant events and combining these with temporal information about events as provided by the structured (i.e., non-narrative) part of a patient's health record.

1 Introduction

The clinical records of patients undergoing hospital care comprise many different kinds of documents, such as lab reports, case notes, clinic letters and discharge summaries. With the impending introduction of *electronic patient records* (EPR) such documents will increasingly be available electronically. This development is not just a matter of convenience, e.g. allowing clinicians easier access to individual documents, but rather has the potential to allow records to be addressed in innovative ways that can make valuable contributions to clinical care and research. For example, a clinical researcher might be able to extract important generalisations about the effectiveness of a particular treatment regime for a particular class of patients by aggregating results across the records of many thousands of patients. We will argue that such analysis requires that the records of individual patients be firstly compiled to provide a coherent overview of the patient's condition and treatment over time -- what we have termed a *patient chronicle*. This task is relatively straightforward for the part of a patient's EPR that consists of structured data, where the information is easily accessible and time-stamped, but there is also much important information regarding patient care contained in narrative texts, such as case notes. Such information can be derived from texts by the use of *information extraction* techniques, but ambiguity is fundamental feature of language, and so the task of linking the medically significant events and entities discovered within narratives to the corresponding elements in the structured record is non-trivial, but this must be achieved if the additional information found about them in the texts is to be integrated into the chronicle. This paper aims to show how temporal information found within textual records can be exploited to facilitate such integration in creating patient chronicles.

The paper is organised as follows. In Section 2, we discuss use of information extraction in the medical domain, and then in Section 3, explain the idea of the patient chronicle and its potential importance. In Section 4, we discuss how temporal information can aid the integration of information from narrative records into the chronicle. In Section 5, we describe our current efforts towards building a system that can automatically extract the required temporal information, and in Section 6, we describe our exploratory work on using temporal information to help in linking relevant events and entities mentioned in narrative records to those listed in the structured data.

2 Medical Information Extraction

Our work on information integration is carried out in the context of the MRC-sponsored Clinical e-Science Framework (CLEF) project. This project aims to establish methodologies and a technical infrastructure for managing repositories of clinical patient data for the purposes of biomedical research and clinical care. An important aspect of building patient data repositories is information capture. Much of the key information regarding patient treatment is contained in textual, narrative patient notes dictated by doctors. Although the final clinical diagnosis of a patient may be represented once within the structured information of an electronic patient record (EPR) – as well as repeatedly in the text of letters written between members of the healthcare team – much valuable clinical information remains locked in the narrative such as:

- earlier provisional diagnoses
- when a relapse of a disease occurred
- what symptoms the patient experienced
- when treatment was changed and why
- why investigations were ordered

The unstructured format of the narratives and their volume make it difficult to survey even a single patient's complete record; it is practically infeasible to aggregate over the records of groups of patients of the size required to carry out clinical research.

To address the information capture barrier, CLEF employs Information Extraction technology, based on Natural Language Processing methodologies, to identify important entities and events referred to in documents and also significant relations that hold between them. By storing the entities, events, and relationships found in a document in a pre-defined format, the information content of an unstructured, textual document can be represented in a structured manner, making it more accessible and amenable to further processing.

CLEF's Information Extraction component contains rules and resources to identify and classify medically relevant classes of entities and events, including drugs, problems (i.e., symptoms and diseases), loci (i.e., that what is affected by a problem, or the target of an intervention or investigation; this class includes anatomical locations, physiological functions, and mental processes), and investigations and interventions, as well as relationships between such entities and events, e.g., that an investigation has revealed a particular problem, which, in turn, has been treated with a particular intervention.

Information of this kind extracted from narrative patient notes can be combined with other information about the patient from the data repository and compiled into a *chronicle*, which is a coherent overview of a patient's condition and treatment over time for clinicians and researchers

3 The Patient Chronicle

We propose that the information available in a patient's clinical record, both from the structured data and narrative texts, should be integrated into a patient chronicle, which is a coherent overview of the significant events in the patient's medical history, i.e. covering their condition, diagnosis and treatment over the period of care. Such chronicles have the potential to be helpful in regard to both clinical care and research. For the former, for example, a patient's chronicle might be used to generate a textual summary of the key aspects of the patient's history to be read by a clinician who is newly involved in the patient's care. For the latter, consider that many of the questions for which a clinical researcher might seek to find answers from a large database of patient records are ones that require not just aggregation over multiple patients, but which are fundamentally stated in terms that relate to the time-course of patients' conditions, treatments and outcomes. For example, a clinical researcher who has a hypothesis in mind about some significant medical effect, might look for initial indications of the correctness of the hypothesis by asking questions such as "How many patients with Stage II adeno carcinoma who were treated with tamoxifen had tumour recurrence within 5 years?" or "For all patients with cancer of the pancreas, compare the percentage alive at five years for those who had a course of gemcitabine with those who didn't".

The structured data component of a patient's clinical record will cover all or most of the noteworthy medical events occurring during a patient's clinical history, such as major diagnoses, the initiation and discontinuation of drug treatments, and investigations such as X-rays, together with associated information, e.g. the body region that was X-rayed. These events will all be clearly time-stamped in the structured data, allowing them to be readily mapped onto the time-line of the patient chronicle, in effect providing a solid "backbone" for the chronicle. However, as discussed in the preceding section, there is additional valuable information to be found in narrative records that will not be found in the structured data. For example, a clinic letter might mention that examination of a previous X-ray has led to a particular diagnosis, or that a new X-ray is being ordered with a view to eliminating a possible diagnosis. To integrate this additional information into the chronicle, we must resolve this mention of an X-ray investigation to one that is listed in the structured data, but if there are several of these, a problem of ambiguity arises. We view this resolution process as one of constraint satisfaction, i.e. one which mobilises constraint information about items in narratives and structured data to rule out incorrect linkages, a process which will be implemented using constraint logic programming techniques, as discussed later in the paper. A valuable source of such constraints is temporal information, as we will illustrate in the next section.

4 Exploiting Temporal Information

When looking at the narrative data for a patient, we are not only faced with the problem of ambiguity as discussed above, but also with the following problem. Information extraction over the set of these documents produces a collection of potentially fragmented and duplicated descriptions of medical entities and events and for inclusion in the chronicle, these various bits and pieces have to be integrated into a coherent whole. Figure 1 shows an example of the same X-RAY event being mentioned in two texts, and we need to be able to map those two events to the same – and correct – X-RAY in the structured record to be able to construct an accurate and complete chronicle.

The narrative texts provide us with temporal clues which can aid us in finding the corresponding structured event. These clues contain temporal expressions of various kinds, including absolute expressions of time, e.g., *She had a mastectomy on 23/5/89*, or expressions to be interpreted relative to the date of the document, e.g., *a chest X-RAY is arranged for next week* (example in figure 1). Temporal information can also be gleaned from the tense and aspect of verbs in combination with the date of the document, e.g. the tense and aspect of *the X-RAY performed* in the example locates the event before the date of the article.

As shown in figure 1, narrative X-RAY event[1] 1 is scheduled for the week after the date of the patient note, and we can see that, based solely on the temporal information, structured events 1, 2 and 3 are potential candidates. In addition to the temporal informatio, we can also employ medical information to narrow down the set of possibilities. In the example, the type of the event, e.g. whether it is an X-RAY or MRI scan, rules out structured event 1 and the locus of the X-RAY (chest) rules out event 3 and so the narrative X-RAY event 1 can be mapped uniquely onto structured event 2. The second narrative event is not located in time by a temporal expression, but we can use the tense information as well as the locus information to map it also to structured event 2.

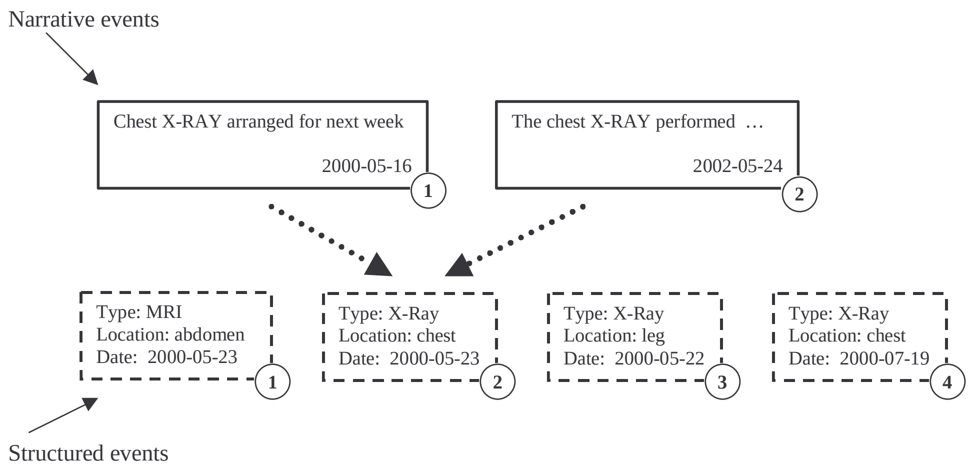

Figure 1: Example of matching narrative and structured events

Accordingly, our system uses the structured data to constrain the set of possible dates for narrative events. First, we use the medical[2] information to select a set of structured events. For example, if the structured information

[1] For conciseness of expression, we write "extracted event" or "narrative event" to mean an event extracted from a textual patient document by the information extraction component, and, similarly, "structured event" to mean an event that is mentioned in the structured information part of a patient's health record.

[2] Note that the only medical information we are using at the moment is the type of event, e.g. X-ray vs. CT scan, etc.

mentions five X-rays that have been performed, then any one X-ray mentioned in a note must be one of these (under the assumption that the structured information is complete).

Then the temporal information about narrative events, gathered from both sources, is formulated as a set of constraints, each constraint having the effect of restricting the time domain or set of possible dates on which an event could have taken place, as described above. Resolving the time domain of an event is treated as a Constraint Satisfaction Problem, and solved using existing Constraint Logic Programming (CLP) tools (see section X). Depending on the number and content of the constraints, the time domain of an event may contain exactly one date, a few dates, many dates, or, in fact, be empty.

In a further step, each extracted event is associated with zero or more structured events according to whether the resolved time domain of the former contains the date of the latter. These associations will be used to evaluate the performance of our method (see section 5.2). This step also facilitates the combination of extracted information and structured information for the chronicle (see section 3).

In the rest of the paper we will discuss the implementation and evaluation of the approach to information integration outlined in the foregoing. In this initial investigation of our approach we are particularly interested in answers to the following two questions: (1) do the narratives contain enough temporal clues to allow non-trivial restrictions of the time domains of extracted events? (2) Is the temporal information in narratives expressed accurately enough to prevent the derived constraints from being contradictory?

5 Annotating Temporal Information

In order to build an automated system which can extract the information from both narrative and structured events necessary for chronicle building, we need to annotate the information we discussed in the section above. The information contained in the structured data is straightforward to extract (as the name suggests), and we will discuss here the mark-up of the narrative texts.

We annotate expressions conveying investigations and time expressions for both absolute and relative times. If the event is subject or object of a “show” or “perform” verb like *The X-RAY shows a shadow on the lung* or *An X-RAY was performed last week*, then we copy the tense and aspect details of the verb onto the event

5.1 Annotation Schema

The annotation scheme used for CLEF is a slimmed down version of the XML-based TimeML annotation scheme. We will give here a brief overview of those parts of the annotation scheme we are using for CLEF; for more detailed information please see Pustejovsky et. al. (2003).

5.2 Annotating Events

The events (investigations) are marked up by annotating key words like X-RAY or CT-scan in the text. The attributes of events are a unique identifier, tense, and aspect. The following fully

```
<TimeML>
<DOC>
<TIMEX3 tid="t6" type="DATE" value="19970707"><DATE>07/07/1997</DATE></TIMEX3>
Seen in Janeway's Outpatients Department.

Mr Sulu came in this morning complaining of pain between his scapula. On examination, he has localised tenderness around the C6-T1 region. Sensation was difficult to ascertain but we have done a cervical thoracic <EVENT aspect="PERFECTIVE" eid="e1" tense="PRESENT">x-ray</EVENT> which shows collapse of the vertebrae. There is collapse of C5/C6/T1 so we have arranged an <EVENT aspect="NONE" eid="e2" tense="NONE">MRI</EVENT> urgently at Kirk Memorial <TIMEX3 tid="t4" type="DATE" value="19970707">today</TIMEX3> to exclude spinal cord compression.

Letter to Dr S. McKoy from Dr Phlox
</DOC>
<MAKEINSTANCE eiid="ei1" eventID="e2"/>
<MAKEINSTANCE eiid="ei2" eventID="e1"/>
<TLINK eventInstanceID="ei1" lid="l2" relType="BEFORE" relatedToTime="t6"/>
<TLINK eventInstanceID="ei1" lid="l3" relType="BEFORE" relatedToEventInstance="ei2"/>
<TLINK eventInstanceID="ei2" lid="l4" relType="IS_INCLUDED" relatedToTime="t4"/>
<TLINK eventInstanceID="ei2" lid="l5" relType="IS_INCLUDED" relatedToTime="t6"/>
</TimeML>
```

Figure 2: An annotatated example

annotated example for an event expression illustrates the event annotation.

An <EVENT eid="1" tense="past" aspect="NONE"> *X-RAY* </EVENT> *was taken*.

5.3 Annotating Times

The annotation of times was designed to be as compatible with the TIMEX2 time expression annotation guidelines (Ferro et al., 2001) as possible. The XML tag for time expressions within TimeML is called TIMEX3 to distinguish it from both the tags in Setzer (2001) and TIMEX2. Within the CLEF scenario, we use the following attributes for TIMEX3:

- A unique identifier: tid,
- A TIMEX3 object is of a certain type, either a date, a time, or a duration.
- The actual ISO value of a temporal expression, for example *1998*, is stored in the value attribute.

Straightforward time expressions like July 1966 are annotated as follows:

<TIMEX3 tid="1" type="DATE" value="1966-07"> *July 1966* </TIMEX3>

The smallest granularity of times is days, we do not take information like *3pm* into account.

5.4 Annotating Temporal Relations

To annotate the temporal relations holding between events and times (e.g. the date of a patient note or a different temporal expression in the note), we use the TLINK tag. A TLINK records the relation between two events, or an event and a time. Thus, the attributes include the IDs of the source and the target entity and the relation type. The temporal relations we use are: before, after, includes, is included, holds, simultaneous, immediately after, immediately before, identity, begins, ends, begun by, ended by (see TimeML specification on www.timeml.org).

For example, the *expression An X-RAY was taken on February 4* states that there is an X-RAY event, a time February 4 and implies a temporal relation IS_INCLUDED between those two entities. The resulting TLINK would look like this:

<TLINK eventInstanceID="2" relatedToTime="4" relType="IS_INCLUDED"/>

Figure 2 shows an simple example of an annotated text.

5.5 The Annotation Pipeline

Our current efforts focus on constructing a semi-automatic pipeline of processing modules, as depicted in figure 3, which, given a patient note and the structured data of a patient, will output a set of events and temporal constraints on these events that can be fed into our CLP tool for resolution. For use in the eventual CLEF prototype, processing must be fully automatic. However, the present pipeline does not only allow us to assess whether our approach to data integration using temporal information is effective, by suppressing the polluting effects of erroneous input to the CLP module due to imperfections in the preceding processing modules[3] - it is also necessary to produce gold standard annotations necessary for evaluating the system's performance.

Let us describe the pipeline step-by-step.

1. In the pre-processing step, a text is tokenised and split into sentences, and the tokens are analysed morphologically and assigned a part-of-speech category. This is implemented as a GATE (see Cunningham et. al. 2002) application.
2. Next, the GUTime tagger (developed at Georgetown University; extends the TempEx tagger, see http://timex2.mitre.org/) is run over the pre-processed text to identify and annotate temporal expressions according to the TIMEX3 standard. The GUTime tagger has been implemented in Perl.
3. The next module in the pipeline recognises and annotates expressions conveying events in TimeML. (As mentioned above, we restrict events to investigations, such as *X-ray* and *PET scan*). Like step 1, this is another GATE application.
4. A text thus marked-up with time and event annotations is then manually inspected using the Callisto annotation tool (a Java appication, developed by MITRE, see http://callisto.mitre.org), providing an opportunity to correct any errors introduced by previous processing steps. Note that this module will not be part of the automated system (hence the dashed outline), but is used for checking the system performance and creating gold standard annotations.
5. In the next step, we use the TANGO annotation tool (developed as part of the TimeML effort, http://www.timeml.org/tango) to manually annotate temporal relations between events and times, again following the TimeML scheme. This module will eventually be automated.

---

[3] Error patterns observed at the manual steps of the pipeline will also be used to improve the automatic system

6. Finally, the input to the CSP module is created. This involves, amongst other things, pairing events extracted from the text and events mentioned in the structured data. These pairs take into account the type of event, to avoid linking, for example, an "extracted" X-ray event to a "structured" PET scan event. Because of the complexity of processing involved, the last two steps of the pipeline are currently manual (step 5) and semi-automatic (step 6).

We are working on converting the pipeline into a fully automated system. This involves solving technical problems such as integrating applications written in different languages into one application, as well as developing and implementing software to automatically establish temporal relations between events and times – the latter being a non-trivial undertaking.

6 Chronicle Integration

As stated above, the process of linking events and entities mentioned in narratives to those that appear in the structured record, as part of the process of information integration needed to construct the chronicle, is performed by a constraint satisfaction component. In this section, we give more detail on how constraint satisfaction is performed, and then go on to report some initial results of an effort to evaluate the potential effectiveness of our approach in aiding chronicle construction.

6.1 Constraint Satisfaction

A Constraint Satisfaction Problem (CSP) is defined by a set of variables, each with a finite domain of possible values, along with a set of constraints over these variables. A solution to a CSP is any assignment of values to variables that respects the domains and constraints. Many CSPs can be solved efficiently by Constraint Logic Programming (CLP): the key idea of this technique is to utilize constraints not only to test potential solutions, but also to restrict domains where possible, thus narrowing the search space.

To solve temporal reasoning problems using CLP it is convenient to represent dates by integers, representing the number of days that have passed since an arbitrary start date. Unknown dates can then be modelled by variables for which the finite domain is an integer range. To determine whether an event in the narratives can be identified, without contradiction, with an event in the structured data, we impose all other constraints, and then test the effect of further stipulating that the dates of the two events are equal. If this further constraint succeeds, the two events might be the same; if it fails, they cannot be the same (assuming that the other constraints are correct).

The structure of the program is as follows. First, it sets up a variable representing the date of each relevant event, and initialises the domains to a safely wide range (e.g., 1..10000). Next, it fixes the values of dates that are already known, and equates the values of variables that are already known to refer to identical events. Next, constraints based on the relations of before and after are imposed, using a greater-than or less-than relation between two domain variables.; these include constraints based on the tenses of verbs. Finally, each potential identification of an event from the narratives with one from the structured data is tested, as described above, and marked as either possible or impossible.

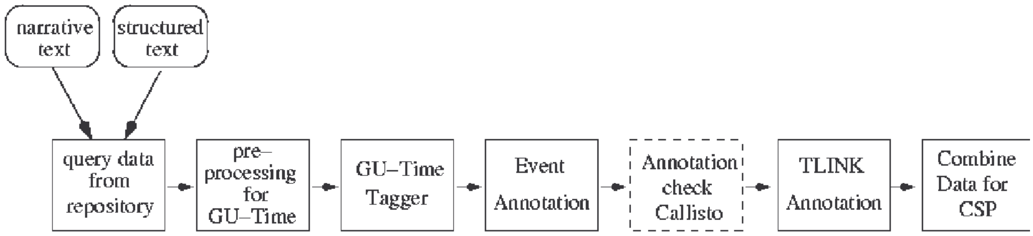

Figure 3: The CLEF pipeline

6.2 Evaluation

The semi-automatic pipeline produces temporally annotated narrative documents, corresponding to a version of what a fully automated temporal analysis component might produce, except that they are `idealised' in containing rather less errors than could be expected of an automated system. Such annotated documents can be provided as input to the satisfaction component, to evaluate the potential effectiveness of the approach for information integration, although the results so produced must be viewed in the light of the fact that the input is idealised.

6.3 The Data and the Gold Standard

Using the semi-automatic pipeline, we have analysed and annotated the patient notes of 5 patients, confining our attention to investigation events. This collection comprises 446 document, of which 94 contain investigation events, which total 152 in number. The corpus is small, due to the large overheads of manual annotation. For the purposes of this evaluation, this temporally annotated data has been augmented with details of the correct linkage between investigation events in the narratives and events from the structured data, i.e. to provide a `gold standard' by which the results of constraint satisfaction can be scored. More specifically, the gold standard annotations indicate, for each narrative event, the full set of possible events within the structured record to which this narrative event could potentially be linked, and also picks which amongst these possible targets are correct. For example, for an X-ray narrative event, this would explicitly list identifiers for all the X-rays investigations appearing in the patient's record, and would indicate for each whether it is the correct target or not. In most cases, a narrative event is mapped to exactly one structured event, but mappings to zero or more than one structured event are also possible. Thus, not every event mentioned in the patient's notes has a structured counterpart. For example, the MRI scan in *If his symptoms worsen, we will perform an MRI scan* does not necessarily happen. Similarly, if the structured data contains more than one MRI scan for one particular date, we cannot know which one is referred to without additional information, in which case the Gold Standard marks all of the indistinguishable alternatives as correct targets.

6.4 Evaluation Metrics

We need metrics to quantify the impact of using the constraint satisfaction process in reducing the ambiguity for mapping narrative events to those that appear in the structured record. Ideally, the process should be such as to greatly reduce ambiguity, by eliminating incorrect candidates from the set of possible targets for each narrative event. At the same time, the process should be such as only rarely to eliminate the true target of a narrative event from the set of possible candidates.

Recall and Precision are two metrics commonly used in NLP which similarly divide in terms of whether they seek to quantify the extent to which truth is preserved and to which ambiguity is eliminated. Here, Recall might be the proportion of the correct targets (i.e. events in the structured data to which narrative events might be resolved) that are recognised by the system as possible targets, whilst Precision would be the proportion of the elements recognised as possible targets that are correct. By applying these metrics both before and after constraint satisfaction is done, and comparing the results so produced, we can verify that truth is preserved and ambiguity eliminated. Recall and Precision scores for the overall data set, i.e. comprising all five patients, both before and after constraint satisfaction, are shown in Table 1. The limitation of these metrics here is that a limited number of events which have a lot of possible targets can serve to give the impression that most events will be incorrectly resolved even when this is not true, e.g. if a few events retain a large number of possible targets this can produce a low precision score even if most events are close to being correctly resolved.

Consequently, we will include also two “accuracy” based scores, which quantify for *each narrative event* the extent to which it is correctly resolved, and then averages performance across all narrative events in the data. In this case, the overall accuracy score depends on whether most events are well resolved, i.e. these scores will not be too badly damaged if a few elements remain highly ambiguous, and hence poorly resolved.

The first of these measures is a `liberal' measure that attempts to assess for each narrative event only whether the true targets have been correctly preserved (i.e. not falsely excluded). Thus, the *liberal score for a single event* is 1 if at least one correct target is preserved amongst the possible targets that remain, and 0 otherwise. The *liberal accuracy* is produced by averaging this score over all narrative events. The *strict score for a single event* is the *proportion* of the possible targets that remain that are correct, and the corresponding *strict accuracy score* is produced by averaging this value over all narrative events. This measure incorporates the extent to which the retention of correct targets for each event are `diluted' by incorrect targets. Liberal and Strict Accuracy scores for our data both before and after constraint satisfaction are also shown in Table 1. A more detailed breakdown of these results for individual patients are shown in Tables 2 and 3.

6.5 Discussion

Comparing the results for Recall of targets before and after CSP indicates that there is some loss of correct targets, but this is limited. The corresponding figures for Liberal Accuracy start of at a value below 1, due to the presence of a reasonable number of events that have no correct targets in the gold standard, but the initial score for this measure is only slightly reduced by CSP. The Precision and Strict Accuracy scores both indicate that there is a substantial amount of ambiguity at the start of the process, which is reduced by CSP, although the Strict Accuracy scores give a better impression of the fact that CSP makes reasonable progress in reducing the ambiguity for a reasonable proportion of the narrative events. The extent to which ambiguity remains after CSP here looks perhaps to be somewhat disappointing, but we should note a number of points regarding these initial experiments. Firstly, the results do show a benefit from exploiting temporal information. Secondly, temporal information has been used somewhat conservatively here. For example, for a patient who has had many X-rays over a long period of care, the fact that a given narrative X-ray is a future event may eliminate some past X-ray investigations, but leave many other candidates as logical possibilities. We could instead here apply heuristic reasoning that the narrative event will probably relate to an investigation within the next few weeks after the document date, allowing perhaps a large number of other possible targets that range several years into the future to be excluded. Thirdly, we have not in these experiments exploited any other possible sources of constraint, e.g. the locus of an investigation, which could serve to eliminate a subset of incorrect targets.

Table 1: Overall results

	Before CSP	After CSP
Recall	1.0	0.94
Precision	0.05	0.09
Liberal Acc.	0.83	0.78
Strict Acc.	0.08	0.27

Table 2: results by patient, pre CSP

	Pat1	Pat2	Pat3	Pat4	Pat5
Recall	1.0	1.0	1.0	1.0	1.0
Prec.	0.1	0.05	0.05	0.08	0.04
Lib.Ac.	0.68	0.73	0.90	0.78	0.90
Str.Ac.	0.10	0.10	0.07	0.07	0.07

Table 3: results by patient, post CSP

	Pat1	Pat2	Pat3	Pat4	Pat5
Recall	0.87	0.93	0.92	1.00	0.96
Prec.	0.27	0.07	0.07	0.16	0.08
Lib.Ac.	0.59	0.68	0.82	0.78	0.86
Str.Ac.	0.28	0.12	0.13	0.37	0.37

7 Conclusion

We have argued that there is an important role for the creation of patient chronicles in the effective exploitation of electronically stored clinical records, in both the context of clinical care and clinical research. We have also argued that temporal information found in narrative textual records can be exploited to achieve the effective integration of information extracted from such records into patient chronicles, and we have described our on-going work aimed at automating this integration process, and evaluating its potential effectiveness.

References

1. Pascal van Hentenryck (1989) Constraint Satisfaction in Logic Programming. MIT Press, Cambridge Mass.
2. Harkema et. al. 2005. Information Extraction from Clinical Documents. In: this volume
3. H. Cunningham, D. Maynard, K. Bontcheva, V. Tablan. GATE: A Framework and Graphical Development Environment for Robust NLP Tools and Applications. *Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistics (ACL'02)*. Philadelphia, July 2002.
4. Ferro L, Mani I, Sundheim B, and Wilson G 2001 *TIDES Temporal Annotation Guidelines, Version 1.0.2* MITRE Technical Report, MTR 01W0000041.
5. Pustejovsky P, Castaño, J, Ingria R, Saurí R, Gaizauskas R, Setzer A, and Katz G *TimeML: Robust Specification of Event and Temporal Expressions in Text.* In: Proceedings of the IWCS-5 Fifth International Workshop on Computational Semantics, 2003.
6. Setzer, A. 2001 *Temporal Information in Newswire Articles: an Annotation Scheme and Corpus Study*, Unpublished PhD thesis, University of Sheffield.