

Multi-modal presentation of medical histories

Catalina Hallett
The Open University
Walton Hall, Milton Keynes, UK
c.hallett@open.ac.uk

ABSTRACT

This paper describes a visualisation architecture that integrates graphical devices and natural language in a cooperative system for navigating through complex images of medical histories. We show how the addition of automatically generated natural language can be used to improve the usability of a graphical user interface and conversely how the graphical user interface can be used to specify the content of user customizable medical reports.

Author Keywords

Natural Language Generation, electronic patient records, navigation tool, visualisation

INTRODUCTION

Records of cancer patients are very rich: in addition to a thousand or more numeric data points arising from successive laboratory tests and a chronology of five or six hundred significant events — such as the dates tests were requested or performed, clinics attended or drugs dispensed — our typical patient files will also contain between fifty and a hundred and fifty narrative clinic letters, together with a similar number of reports interpreting a variety of investigations (e.g., Xray, body scan, radiosintigraphy etc.).

One of the aims of the Clinical e-Science Framework (CLEF) project [1], under which the research reported here is being conducted, is to establish a technical infrastructure for managing research repositories of aggregated patient data arising from routine medical care across potential multiple sites and institutions, in support of biomedical research. Information is extracted from medical narratives¹ and aggregated with structured data in order to build complex images of a patient's medical history which model the story of how patient illnesses and treatments unfolded through time: *what happened, when, what was done, when it was done and why*. The resulting complex semantic network, termed *chronicle*, allows the construction of targeted summarized reports which do more than present individual events in a medical history: they present, in coherent text, events that are semantically and temporally linked to each other.

This paper discusses the problem of presenting the information held in chronicles back to the users (medical practitioners) in a format that could benefit clinical care and medical research. Of particular interest to us here is the problem of integrating generated targetted and comprehensible textual

¹Using Natural Language Processing techniques, see [2].

reports with visual navigation tools that allow time-efficient access to particular events or sets of events in a patient's medical history.

In presenting medical histories we are trying to circumvent the shortcomings of textual reports by combining them with visual navigation tools. In this way, we take advantage of the better accessibility and interactivity offered by visual timelines as well as of the ability of natural language to convey complex temporal information and to aggregate numerical data.

APPLICATION

Input

A CLEF chronicle [3, 4] is a rich semantic network that logically and temporally organises events in a patient's medical history. The chronicle models what happened to the patient, why and how. The chronicle is backed by a logical model of a cancer history and a series of description logic based ontologies which act as knowledge sources for the domain of cancer. As an example, a chronicle can encode the following history (italicised text represents events, the type of event follows in brackets): a patient presents to the clinic with *pain in the right breast* (Problem). A *clinical examination* (Investigation) is performed which suggests *cancer* (Problem). As a result, a *biopsy* (Investigation) is recommended, which returns a histology diagnosis of *invasive ductal adenocarcinoma* (Problem). The patient is treated for cancer with *radiotherapy* (Treatment), which causes an *ulceration in the right breast* (Problem) which is treated with medication (Treatment) over several weeks. This simplified scenario displays several events which occur at various moments in time and are linked through semantic relations such as causality, reason, finding and consequence. In practice, each individual event will consist of much more detail. For example, a diagnosis of cancer will include the size of the tumour, the staging, whether the tumour is hormone receptive or not, genetic information. Similar clusters of events may happen several times during a patient's history and semantic links may exist between events situated at some temporal distance. Whilst it is clear that such views of medical histories present a much richer source of information than traditional log-like patient records (that record events but no dependencies) it is also clear that retrieving information from chronicles will require different visualisation techniques. It is beyond the scope of this paper to describe the methodology involved in transforming an EPR into a chronicle - the chronologicalisation process is complex and involves Infor-

mation Extraction from narratives, solving multi-document coreference, temporal abstraction and inferencing over both structured and information extraction data [2].

User requirements

The intended end-user of the generated reports is a general practitioner or clinician who uses electronic patient records at the point of care to familiarise themselves with a patient's medical history and current situation. They are familiar with finding information in text documents but have little or no exposure to graphical interfaces to Electronic Patient Records. Since this is a novel application we did not have the possibility of observing modalities of interaction with the input data. An initial requirements elicitation process resulted in a set of general requirements which are summarised below:

- Each individual event in a patient's history should be available for inspection
- Events that deviate from the norm are more important than normal events (e.g., an examination of the lymphnodes that reveals lymphadenopathy is more important than an examination that doesn't). However, normal events should also be available on demand.
- Some events are more important than others and they should not only be included in the summary but also highlighted (through linguistic means, colour coding, graphical timelines or similar display features).
- Having different views of the same data is a useful feature, because it allows the clinician to spot correlation between events that they may have missed otherwise.
- The interface should allow focusing on time intervals as well as clusters of related events (i.e. events of the same type as well as events of different types linked through causality relations)
- Summaries that provide a 30-second overview of the patient's history are often desirable; ideally, these should fit entirely on a computer screen. However, users should be able to obtain more detailed information about specific events by expanding their description.

The interface should allow users to search and retrieve information such as:

- Why has a certain event happened? (*Why was the patient's medication changed?*)
- When has a certain event or series of events happened? (*When did the patient undergo surgery?*)
- How many events of a certain type occurred in a certain time interval? (*How many courses of chemotherapy were cancelled in the course of a full chemotherapy cycle?*)
- What was the outcome of a certain event? (*What were the findings of the fine needle aspirate biopsy?*)
- How did a certain event progress in a certain time interval? (*How did the platelet count change between the start and the end of a chemotherapy treatment?*)

Following the user requirements and the type of input we were dealing with, it became clear that we were faced with two relatively distinct tasks: providing navigational capabilities to an Electronic Patient Record and generating short textual summaries describing snapshots of the patient's medical history.

Textual reports have the advantage of offering a snapshot view of a patient's history at any point in time, they can be amended and printed, used in communication between clinicians or clinicians and patients. Text is a good way of describing temporal information (events that happened at a certain position in time with respect to another event), of summarising numerical data (for example, specifying that *liver tests were normal* instead of listing individual measurements for bilirubin concentration, Alanine aminotransferase, Alkaline phosphatase, Aspartate aminotransferase, albumin and total protein). However, pure text is not always the best medium of presenting large amount of information, part of which is numerical and most of which is highly interconnected. Since the generated reports are likely to be very long (thus violating the 30 seconds report requirement), accessing particular pieces of information is much more difficult in textual reports. Linear text loses the ability of navigating through information, of expanding some events and of highlighting important dependencies between pieces of information. A textual report alone cannot effectively combine the time sequence element with the semantic dependencies - both of which are essential in representing patient records. Depending on the type of report chosen, either one or the other of these elements will necessarily be emphasised at the expense of the other.

We envisage therefore that, depending on circumstances, users may want to have fully textual reports (for example, for producing printed summaries of a patient's history) or combined graphical and textual reports (for interactive visualisation). The solution we propose in this project is an integrated visualisation tool where users can interact with a graphical interface coupled with a text generation engine in order to navigate through patient records. The visualisation technique is designed to fulfil three main desiderata: interaction (better achieved through visual navigation), clarity (improved readability of the visual display is achieved through text generation which shifts the burden from graphical representations to textual descriptions), and completeness (if the information is available, it should be available to the user without impeding the clarity of the interface, which is achieved through the combination of visual navigation and text generation).

In the following, we will describe the two types of report generated in either of the two scenarios. Section will describe in more detail the natural language generation techniques employed in generating both independent textual reports and report snippets that support the graphical interface.

VISUALISATION MODALITIES

Visual navigator

he visual navigator presents on one screen a high level overview of a patient's medical history. The visual display is built around a timeline paradigm enhanced with visual techniques for highlighting relationships between events on the timeline.

The information in the chronicle is organised along three parallel timelines, corresponding to *Diagnoses*, *Treatments* and *Investigations*. Events in a patient's medical history are represented as graphical objects visually differentiated by colour and icons. Both snap events (occurring at a certain moment in time) and span events (occurring over a period of time) are represented on the timelines. Furthermore, each event can be either simple or composite. Composite events are events that are from a semantic viewpoint a single entity (they have a common cause and a common aim) and combine under the same label several sub-events. For example, a chemotherapy treatment is a composite event consisting of a number of chemotherapy courses (usually six) administered over a period of time. However there is no direct relationship between span events and composite events. Some composite events are snap events (for example, a liver test is a span event consisting of several separate tests) whilst some span events are simple events (for example, medication can be administered continuously over several weeks). Composite events can have several levels of compositionality. For example, *primary treatment packages* are composed of various instances of radiotherapy, surgery, medication, chemotherapy and each of these events can be either simple or composite events. The user can drill into composite events to reveal details about component events and, conversely they can collapse composite events to obtain a higher level view. By zooming in and out of a time span the user can get more or less detailed views of the events in that time span.

Timeline labels also act as interactive graphical objects that correspond to classes of events of their corresponding category (i.e. *diagnosis*, *treatment* and *investigation*). The same applies to time labels on the X-axis, which correspond to sets of events that occur in the time span represented by the time label. Selection of the labels has as effect all events subsumed by the category being brought into focus.

In a high level view, all events currently in focus are represented in a uniform manner, as graphical objects on a timeline (see Figure 1). In the initial, default view all events are in focus, which results in the timelines being split into spans of equal length corresponding to years in the medical history. The graphical objects corresponding to events are then placed in the time span they correspond to.

Unselecting events or bringing other events into focus results in a reorganisation of the currently focused graphical objects on the timeline. Time spans that contain no focused events are collapsed and the remaining timespans are redistributed equally. Since the chronicle contains both factual and numerical data, we have to provide different types of low level view according to the type of the selected event. For example, in the case of procedures the fact that they were performed is their main defining property, therefore placing a

graphical object corresponding to an instance of a procedure on a timeline is sufficient. However in the case of blood tests, which have numerical values associated with them it is much more likely to assume that the user is interested in the actual value of the measurement, rather than in the fact that such a measurement occurred, therefore the view changes to a line chart with graphical objects representing individual measurements being placed at the appropriate position in the chart according to the measurement value, instead of linearly along a timeline. (see Figure2).

Apart from visualising events and time spans, one of the useful features of the navigator is the ability to visualise relationships between events. Since events in a chronicle are highly interconnected, it is not desirable that all dependencies should be available at all times; instead, the user can hide or show types of dependencies as needed. This can be done either at event level (e.g. showing only the cause and the outcome of a surgical procedure) or at global level (show all causality and consequence relations between multiple selected events). Right click on a graphical object brings up a check-box menu where the user can select from the available relations associated to the selected event, the ones they want to visualise. The result of this selection is three-fold: the visual navigator brings into focus (if not already there) the events that are linked to the current event through the selected types of relationship, it displays the relationships as colour-coded arrows and it also hides from the timeline unconnected events for reasons of improving clarity. This may also result in time spans that have no visible events being collapsed or reduced in size for a better display of the required information on the screen.

The visual navigator allows three main scenarios of data retrieval and navigation. In the first scenario, the user identifies an event on the timeline and proceeds to find more information about that particular event and visualises relationships with other events. This is the type of interaction which occurs for example when the doctor wants to know what was the cause and outcome of a surgical procedure. They can achieve this by selecting a focus event, expanding it if possible and selecting the relationships they are interested in. The second scenario is that of the medical practitioner trying to get an overview of what happened in a patient's history in




Figure 1. Visual history snapshot: high level view




Figure 2. View of the trend of the billirubin concentration

a certain time span, e.g. in the last year or following treatment. This is achieved by selecting time spans and zooming into the selected portion of the medical history, which reveals more detail about individual events.

In the third scenario, the user may want to discover events on the timeline. For example, a doctor may notice abnormal values of liver functions which are not explainable by current events in the patient's history and may suspect a previous history of liver problems. In this case they may want to find previous recordings of diagnoses related to problems of the liver. In this scenario, the user does not interact with events on the timeline, but with classes of events instead.

Textual reports

In the following, the term *Report Generator* will be used to designate the software that performs text generation, as a result of either a direct request from the user for a specific type of report or a selection of events in the graphical timeline. The output of the report generator may be either a full report or a report snippet, but practically, the type of selection employed in choosing the focus of the report does not influence the technique used in generating it.

The chronicle described in whichs represent the input to the Report Generator is a highly structured representation of an Electronic Patient Record, in the form of a semantic network. For the purpose of this paper, we consider the input correct and complete. Each node in the network describes an event in the patient's medical history

The chronicle relations can be categorised into three types according to their role in the generation process. *Rhetorical relations* are relations of causality and consequence between two facts (such as, *Problem CAUSED-BY Intervention* or *Intervention INDICATED-BY Problem*) and are used in the document planning stage for automatically inferring the rhetorical structure of the text, as it will be described in . *Ontological relations* such as *Intervention IS-SUBPART-OF Intervention* bear no significance in text planning and realisation, but can be used in content selection. *Attributive relations* such as *Problem HAS-LOCUS Locus* or *Investigation HAS-INDICATION Problem* are used in grouping messages in a coherent way, prior to the construction of the rhetorical structure tree. The former are in fact not rep-

resented as relations in the chronicle semantic network but rather as fields in an object description, however for reasons of consistency in the design of our system we will consider them as semantic relations.

The system design of the Report Generator follows a classical NLG pipeline architecture [5], with a Content Selector, Content Planner and Syntactic Realiser. The Content Planner is tightly coupled with the Content Selector, since part of the document structure is already decided in the event selection phase. Aggregation is mostly conceptual rather than syntactic, therefore it is performed in the content planning stage as well.

Content selection

The process of content selection is driven by two parameters: the type of summary and the length of summary. We define the concept of *summary spine* to represent a list of concepts that are essential to the building of the summary. For example, in a summary of the diagnoses, all events of type *Problem* will be part of the spine (Figure 3). Events linked to the spine through some kind of relation may or may not be included in the summary, depending on the specified type and length of the summary. The design of the system does not restrict the spine to containing only events of the same type: a spine may contain, for example, *Problems* of type *cancer*, *Investigations* of type *x-ray* and *Interventions* of type *surgery*.




Figure 3. Spine events for a summary of diagnoses

The relations stored in the chronicle help in the construction of clusters of related events. A typical cluster may represent, for example, that a patient diagnosed with cancer following a clinical examination, a mastectomy was performed to remove the tumour, a histopathological investigation on the removed tumour confirmed the cancer, radiotherapy was given to treat the cancer, which caused an ulcer that was then treated with some drug. Smaller clusters are generally not related to the main thread of events, therefore the first step in the summarisation process is to remove small clusters². The next step is the selection of important events, as defined by the type of summary. Each cluster of events is a strongly connected graph, with some nodes representing spine events. For each cluster, the spine events are selected,

²In the current implementation these are defined as clusters containing at most three events. This threshold was set following empirical evidence.

together with all nodes that are at a distance of less than n from spine events, where n is a user-defined parameter used to adjust the size of the summary. For example, in the cluster presented in figure 4, assuming a depth value of 1, the content selector will choose *cancer*, *left breast* and *radiotherapy* but not *radiotherapy cycle*, nor *ulcer*.




Figure 4. Example of cluster

The result of the content selection phase is a list of messages, each describing an event with some of its attributes specified. The number of attributes specified depends on the depth level of a message (i.e., how far from the spine the event is). For example, a *Problem* event has a large number of attributes, consisting of *name*, *status*, *existence*, *number of nodes counted*, *number of nodes involved*, *clinical course*, *tumour size*, *genotype*, *grade*, *tumour marker* and *histology*, along with the usual time stamp. If the *Problem* is a spine event, all these attributes will be specified, whilst if the *Problem* is two levels away from the spine, only the name and the existence will be specified.

Document planning

The document planner component is concerned with the construction of complete document plans, according to the type of summary and cohesive relations identified in the previous stage. The construction of document plans is, however, initiated in the content selection phase: content is selected according to the relations between events, which in turn informs the structure of the target text.

The document planner uses a combination of schemas and bottom-up approach. A full report (generated for the purpose of printing, for example) is typically formed of three parts:

- a schematic description of the patient's demographic information (e.g., name, age, gender);
- a two-sentence summary of the patient's record (presenting the time span of the illness, the number of consults the patient attended and the number of investigations and interventions performed);
- the actual summary of the record produced from the events selected to be part of the content.

The visualisation tool only makes use of the third component, therefore in what follows, we will concentrate on this latter part.

The first stage in structuring the summary is to combine messages linked through attributive relations. This results in in-

stances such as that shown in example (3), where a *Problem* message is combined with a *Locus* message to give rise to the composite message *Problem-Locus*.

In the second stage, messages are grouped according to specific rules, depending on the type of summary. For longitudinal summaries, the grouping rules will, for example, stipulate that events occurring within the same week should be grouped together, and further grouped into years. In event-specific summaries, patterns of similar events are first identified and then grouped according to the week(s) they occur in; for example, if in week 1 the patient was examined for enlargement of the liver and of the spleen with negative results and in week 2 the patient was again examined with the same results and underwent a mastectomy, two groups of events will be constructed, leading to output such as:

- (1) In weeks 1 and 2, examination of the abdomen revealed no enlargement of the liver or of the spleen.
In week 2, the patient underwent a mastectomy.

Within groups, messages are structured according to discourse relations that are either retrieved from the input database or automatically deduced by applying domain specific rules. At the moment, the input provides three types of rhetorical relation: Cause, Result and Sequence. The domain specific rules specify the ordering of messages, and always introduce a Sequence relation. An example of such a rule is that a histopathology event has to follow a biopsy event, if both of them are present and they start and end at the same time. These rules help building a partial rhetorical structure tree. Messages that are not connected in the tree are by default assumed to be in a List relation to other messages in the group, and their position is set arbitrarily. Such events are grouped together according to their type; for example all unconnected *Intervention* events, followed by all *Investigations*.

In producing multiple reports on the same patient from different perspectives, or of different types, we operate under the strong assumption that event-focussed reports should be organised in a way that emphasises the importance of the event in focus. From a document structure viewpoint, this equates to building rhetorical structures where the focus event (i.e., the spine event) is expressed in a nuclear unit, and skeleton events are preferably in satellite units.

At the sentence level, spine events are assigned salient syntactic roles that allows them to be kept in focus. For example, a relation such as *Problem CAUSED-BY Intervention* is more likely to be expressed as:

The patient developed a *Problem* as a result of an *Intervention*.

when the focus is on *Problem* events, but as:

An *Intervention* caused a *Problem*.

when the focus is on *Interventions*.

This kind of variation reflects the different emphasis that is placed on spine events, although the wording in the actual report may be different. Rhetorical relations holding between simple event descriptions are most often realised as a single sentence (as in the examples above). Complex individual events are realised in individual clauses or sentences which are connected to other accompanying events through the appropriate rhetorical relation. Additionally, the number of attributes included in the description of a *Problem* is a decisive factor in realising an event as a phrase, a sentence or a group of sentences. In the following two examples, there are two *Problem* events (*cancer* and *lymphnode count*) linked through an *Investigation* event (*excision biopsy*), which is indicated by the first problem and has as a finding the second problem. In Example 4, the problems are first-mentioned spine events, while in Example 5, the problems are skeleton events (the cancer is a subsequent mention and the lymphnode count is a first mention), with the *Investigation* being the spine event.

- (2) A 10mm, EGFR +ve, HER-2/neu +ve, oestrogen receptor positive cancer was found in the left breast (histology: invasive tubular adenocarcinoma). Consequently, an excision biopsy was performed which revealed no metastatic involvement in the five nodes sampled.
- (3) An excision biopsy on the left breast was performed because of cancer. It revealed no metastatic involvement in the five nodes sampled.

As these examples show, the same basic rhetorical structure consisting of three leaf-nodes and two relations (CAUSE and CONSEQUENCE) is realised differently in a *Problem*-focussed report compared to an *Investigation*-based report. The conceptual reformulation is guided by the type of report, which in turn has consequences at the syntactic level.

Aggregation

The fluency of the generated text is enhanced by conceptual aggregation, performed on messages that share common properties. Simple aggregation rules state, for example, that two investigations with the same name and two different target loci can be collapsed into one investigation with two target loci. Consider, for example, a case where each clinical examination consists of examinations of the abdomen for enlargement of internal organs (liver and spleen) and examination of the lymphnodes. Thus, each clinical examination will typically consist of three independent *Investigation* events. If fully expanded, a description of the clinical examination may look like:

- (4) • examination of the abdomen revealed no enlargement of the spleen
- examination of the abdomen revealed no enlargement of the liver
- examination of the axillary lymphnodes revealed no lymphadenopathy of the axillary nodes

With a first level of aggregation, this is reduced to:

- (5) • Examination of the abdomen revealed no enlargement of the spleen or of the liver.
- Examination of the axillary nodes revealed no lymphadenopathy.

Further aggregation of the two examination events transforms the output into:

- (6) Examination revealed no enlargement of the spleen or of the liver and no lymphadenopathy of the axillary nodes.

However, even this last level of aggregation may be not enough, since clinical examinations are performed repeatedly and consist of the same types of investigation.

The system makes use of two solutions to this problem, both of which make use of domain specific rules. The first is to report only those events that deviate from the norm - for example, abnormal test results. The second, which leads to larger summaries, is to produce synthesised descriptions of events. In the case of clinical examinations for example, it can describe a sequence of investigations such as the one in Example 7 as "*The results of a clinical examination were normal*". If the examination result deviates from the norm on a restricted numbers of parameters, this can be described as "*The results of clinical examination were normal, apart from an enlargement of the spleen*".

INTEGRATING TEXT AND GRAPHICS

The addition of natural language support to the data visualiser serves two main purposes which are achieved by two different types of user interaction.

First, the user may directly require textual reports for the purpose of producing documents suitable for printing or for exchanging information. In this case, the text generation process is triggered by the user through direct selection of the report generation option in the interface. The user can further specify the type of information they would like to be included in the report either by choosing a certain predefined type of report (such as for example, a longitudinal report or a summary of interventions) or by specifying in more detail the type of information they want included in the report (for example, surgical interventions, completed chemotherapy courses, medication). This information can be specified by selecting graphical objects on the timeline (including objects corresponding to individual events, classes of events, years or time spans). Selection of events will produce event-focused summaries, whilst selection of time spans will produce longitudinal summaries for that particular span. Another way of including information in a summary is by dragging and dropping graphical objects from the timeline onto the text. This results into adding the specific event or class of events represented by the graphical object to the content of the report and regenerating the report to reflect the new

content.

The second and more interesting role of natural language is as a support tool for the visual navigator, which facilitates the understanding of complex visual devices and better describes relationships between events. The paradigm that underpins the integration of text and graphics is the fact that the information made available to the user in text format directly reflects the user interaction with the graphical interface in a uniform manner. All generated text is produced by re-organising the information available to the text generation module, therefore the visual front-end interacts directly with the content selection component of the report generator. This involves not only the addition and removal of nodes in the content graph, but also the inclusion of relations in the graph, which has implications in the construction and realisation of the discourse structure tree. In this scenario we can differentiate between three types of graphics-text interaction. The simplest type of text generation is the generation of short tooltips describing graphical objects. This occurs when the user hovers the mouse over a graphical object corresponding to an event on the timeline; this triggers the generation of a tooltip with a short description of the event. A higher level of text generation is performed as a result of the user selecting a graphical object on the timeline. This action is interpreted as a request for a complete textual description of the particular event represented by the selected graphical object. A graphical icon representing a chemotherapy event, for example, “hides” information about the particular drug regimen used, dates of the chemotherapy cycles and reasons for deferring a particular cycle. Since this information is better expressed as text than graphically, each selection of an event will trigger the production of a report snippet that describes in more detail that particular event. This procedure is equivalent with the production of a focused summary with only one central event, therefore the report snippet is more than a simple description of the event; it also contains the dependencies that particular event has with other events in the medical history. For example, a description of a chemotherapy event will include not only features of the treatment - duration, type of drug, status, but also the reason for administering (or deferring/cancelling) chemotherapy, side effects and outcome. If the graphical object selected corresponds to a composite event, the user can drill down into the component sub-events. This action is reflected in text by unpacking aggregate descriptions of events into their components. For example, a composite event such as *chemotherapy package* will be expanded into several sub-events corresponding to the chemotherapy cycles within the package and the textual description will include information such as the dates when chemotherapy cycles, possible side-effects or complaints following individual chemotherapy cycles and reasons for cancelling or delaying them. This additional information is integrated seamlessly with the higher level description of the chemotherapy package.

A different type of text generation is performed for supporting charts and diagrams describing the trend of numerical values. Captions are generated using template-based techniques, which are more appropriate for producing simple

text with little variability than the kind of deep generation used in producing full reports.

RELATED WORK

Natural language generation has been used in the medical domain for various applications. For example: to generate drug leaflets (i.e., pill inserts) in multiple languages and styles (PILLS [6]), letters to patients to help them stop smoking (STOP [7]), individualised patient-education brochures (MIGRANE [8]); HealthDoc [9]); Piglit [10]). There is also a body of work on the generation of summaries of patient records (e.g., [11], [12]). Most computer-based patient record management systems have simple generation facilities built-in, which produce text normally consisting of unconnected sentences and thus lacking fluency. Natural language generation techniques have been applied in various reporting systems for generating telegraphic textual progress notes [13], reports on radiographs [14], and bone scans [15] or post-operative briefings [16]. However, this work differs from ours in that they concentrate on the summarization of *textual* records, as opposed to *data encoded* records.

Outside the medical domain, several NLG systems have been developed for generating reports from data, most notably the SummTime project which deals with the generation of numerical time-series data in several domains, including, of most interest to us, neonatal intensive care[17]. Our project however deals with a combination of numerical and non-numerical data, therefore describing numerical trends is only one aspect of our research.

Unlike previous projects, which used generated text to provide explanations for graphical devices ([18], [19]) or as an alternative means of describing data ([20], [21]), our emphasis is on text generated interactively as a response to the user interaction with the visual display.

The timeline paradigm has been used extensively in visualising patient histories. The Lifelines project [22] provides a method for visualising and accessing personal histories by means of a graphical interface, and has been used for both patient records and legal case histories. TeleMed [23] gathers patient information from distributed databases and presents it in a Web interface as icons on a timeline; interaction with the icons provides access to more detailed descriptions of the individual pieces of information. Various authors describe the advantages of the timeline approach to visualising temporal data of the kind present in patient histories [24, 25]. Our particular application however differs significantly from other timeline methods in the fact that capturing and visualising relationships between events is equally important to anchoring events in time. The second major difference is of course the ability of supplementing visual navigation with textual explanations, which improves the clarity of the interface.

CONCLUSIONS AND FUTURE WORK

We presented in this paper a visualisation tool that integrates natural language and visual devices for presenting and navigating through complex views of medical histories termed

chronicles. The graphical interface is used to specify the content of medical reports, whilst automatically generated textual descriptions are used to support the navigation tool by improving clarity and moving the presentation of details away from the graphical interface. Whilst formative evaluation has been performed, which helped us refine and improve the system with the help of comments from potential users, a final end user evaluation is still outstanding.

The design and implementation of the current visualisation tool is based on the assumption that the input is correct and complete. In practice however the medical chronicle is the outcome of a complex process which involves information extraction from natural language text and heuristics-based inferencing, which can result in missing or incorrect data. Some data may also be missing simply because it was not available to the chronological process. An interesting topic for future research is taking into account modalities of presenting potentially incomplete data to the user without introducing misleading information. This involves using different visualisation modalities for data that is marked as uncertain in the chronicle, and also reflecting this uncertainty in the generated text.

Acknowledgment

The work described in this paper is part of the Clinical E-Science Framework (CLEF) project, funded by the Medical Research Council grant G0100852 under the E-Science Initiative. We gratefully acknowledge the contribution of our clinical collaborators at the Royal Marsden and Royal Free hospitals, colleagues at the National Cancer Research Institute (NCRI) and NTRAC and to the CLEF industrial collaborators.

REFERENCES

1. Rector A, Rogers J, Tawee A, Ingram D, Kalra D, Milan J, Gaizauskas R, Hepple M, Scott D, Power R: **CLEF - Joining up Healthcare with Clinical and Post-Genomic Research**. In *Second UK E-Science "All Hands Meeting"*, Nottingham, UK 2003.
2. Harkema H, Roberts I, Gaizauskas R, Hepple M: **Information Extraction from Clinical Records**. In *Proceedings of the 4th UK e-Science All Hands Meeting*, Nottingham, UK 2005.
3. Rogers J, Puleston C, Rector A: **The CLEF Chronicle: Transforming Patient Records into an E-Science Resource**. In *Proceedings of the UK E-Science All-Hands Meeting* 2006.
4. Rogers J, Puleston C, Rector A: **The CLEF Chronicle: Patient Histories Derived from Electronic Health Records**. In *ICDE Workshops 2006* 2006.
5. Reiter E, Dale R: *Building natural language generation systems*. New York, NY, USA: Cambridge University Press 2000.
6. Bouayad-Agha N, Power R, Scott D, Belz A: **PILLS: Multilingual generation of medical information documents with overlapping content**. In *Proceedings of LREC 2002* 2002:2111–2114.
7. Reiter E, Robertson R, Osman L: **Lessons from a Failure: Generating Tailored Smoking Cessation Letters**. *Artificial Intelligence* 2003, **144**:41–58.
8. Buchanan B, Moore J, Forsythe D, Carenini G, Ohlsson S: **Involving patients in health care: explanation in the clinical setting**. In *Proceedings of the Sixteenth Annual Symposium on Computer Applications in Medical Care (SCAMC'92)* 1992:510–512.
9. Hirst G, DiMarco C, Hovy EH, Parsons K: **Authoring and Generating Health-Education Documents That are Tailored to the Needs of the Individual Patient**. In *Proceedings of the 6th International Conference on User Modeling*, Italy 1997.
10. Binsted K, Cawsey A, Jones R: **Generating personalised patient information using the medical record**. In *Proceedings of Artificial Intelligence in Medicine Europe*, Pavia, Italy 1995.
11. Afantinos SD, Karkaletsis V, Stamatopoulos P: **Summarization from Medical Documents: A Survey**. *Artificial Intelligence in Medicine* 2005, **33**(2):157–177.
12. Elhadad N, McKeown K: **Towards Generating Patient Specific Summaries of Medical Articles**. In *Proceedings of the Workshop on Automatic Summarization, NAACL 2001*, Pittsburgh, USA 2001.
13. Campbell K, Wieckert K, Fagan L, MMusen: **A Computer-based Tool for Generation of Progress Notes**. In *Proceedings of the Symposium on Computer Applications in Medical Care (SCAMC)* 1993:284–288.
14. A Abella JS J Kender: **Description Generation of Abnormal Densities Found in Radiographs**. In *Proceedings of the Symposium on Computer Applications in Medical Care (SCAMC)* 1995:542–546.
15. Bernauer J, Gumrich K, Kutz S, Linder P, Pretschner D: **An Interactive Report Generator for Bone Scan Studies**. In *Proceedings of the Symposium on Computer Applications in Medical Care (SCAMC)* 1991:858–860.
16. M Dalal K S Feiner: **MAGIC: An Experimental System for Generating Multimedia Briefings about Post-Bypass Patient Status**. *Journal of the American Medical Informatics Association* 1996, :684–688.
17. Portet F, Reiter E, Hunter J, Sripada S: **Automatic Generation of Textual Summaries from Neonatal Intensive Care Data**. In *Proceedings of AIME 2007* 2007:227–236.
18. Mittal VO, Carenini G, Moore JD, Roth S: **Describing complex charts in natural language: a caption generation system**. *Computational Linguistics* 1998, **24**(3):431–467.
19. Fasciano M, Lapalme G: **PostGraphe: A system for the generation of statistical graphics and text**. In *Proceedings of the Eighth International Workshop on Natural Language Generation*, Brighton, U.K 1996.

20. Green NL, Carenini G, Kerpeljiev S, Mattis J, Moore JD, Roth SF: **AutoBrief: an experimental system for the automatic generation of briefings in integrated text and information graphics.** *Int. J. Hum.-Comput. Stud.* 2004, **61**:32–70.
21. Kerpeljiev S: **Automatic generation of multimodal weather reports from datasets.** In *Proceedings of the 3rd International Conference on Applied Natural Language Processing*, Trento, Italy 1992:48–54.
22. Plaisant C, Mushlin R, Snyder A, Li J, Heller D, Shneiderman B: **Lifelines: Using visualization to enhance navigation and analysis of patient records.** In *Proceedings of the 1998 American Medical Informatic Association Annual Fall Symposium* 1998:76–80.
23. Kilman DG, Forslund DW: **An International Collaboratory Based on Virtual Patient Records.** *Communications of the ACM* 1997, **40**(8):110–117.
24. Tufte E, Kahn M: *The Visual Display of Quantitative Information.* Cheshire, Connecticut: Graphics Press 1983.
25. Cousins, Kahn M: **The visual display of temporal information.** In *Artificial Intelligence in Medicine* 1991:341–357.