

Watson-Aided Non-Linear Problem-Oriented Clinical Visit Preparation on Tablet Computer

Puripant Ruchikachorn, Jennifer J. Liang, Murthy Devarakonda, and Klaus Mueller




Fig. 1. All items that are related to diabetes are highlighted and color-coded according to their categories. Some lists such as *Encounters* are hierarchical. *Lab Tests* is also hierarchical but not fully open. *Procedures* and *Vitals* have no related items so they are folded to save screen space. This provides an uncluttered overview of relevant medical information suitable for tablet computer.

Abstract— Tablet computers have been growing in popularity and finding their way to workspace. Their ubiquity is attributable to their ease of use and portability, also applicable to hospital work environment. We have designed an Electronic Medical Record (EMR) application for clinical visit preparation task on tablet computers. There are many design challenges due to the limited form factor for both input and output. According to our interviews with physicians, our solution utilizes IBM Watson’s machine learning ability to analyze unstructured clinical notes and create a relation graph of medical entities. With the medical knowledge graph, the system enables easy relation retrieval to show only relevant information as needed in simple and easy-to-understand displays. We informally evaluated our system by walking through sample visit preparation steps from a physician.

Index Terms—Interactive systems, machine learning, medical information systems, tablet computer, visualization

1 INTRODUCTION

Tablet computers are popular. Their ease of use and portability are not only for home users but also gaining their place in business. Their flexibility of user interface and work location is desirable in work environments including hospitals and clinics.

To increase the adoption of tablet computers in the workplace, we explore how to bring the power of Watson to enable an intelligent EMR that has a conceptual understanding of medicine. We target the first part of the clinical workflow, visit preparation. Physicians as our target users in this task want to get each patient’s history, ongoing medical problems, and latest data since the last visit.

As one patient may have many medical problems and take numerous medications, there is a lot of information to be shown to physician in charge. Given the limited input and output areas of tablet computers, not all information can be shown at once and the physician should be able to interactively recall any pieces of information and turn them off to give way to other pertinent information.

Our input for a single record is unstructured clinical notes. Ma-

chine learning utilization is twofold: to extract medical entities and to connect them for information organization. The physician can browse through relevant medical entities and quickly look into medical conditions in a non-linear fashion, not only chronologically or hierarchically. Essentially, we combine relation analysis, timelines, and tablet computing paradigm to mirror the way physicians create their mental map of medical knowledge while reviewing patient health records.

2 RELATED WORK

There are numerous visualizations and interactive systems for EMR [1]. Many focus on visual display design, especially for time critical situations, and post-clinical-visit analysis. Pioneered by LifeLines [2], most show medical data chronologically. Without concrete links between timelines, the relations between entities have to be visually deduced. Other approaches include adopting journalistic reporting [3]—the Five W’s: Who, What When, Where, and Why—to connect across data dimensions for higher-level queries. However, the information architecture is often strictly hierarchical. Although this is intentional, it may not support open-ended data access and investigation. Our system emphasizes more on relations between data across categories and hence non-linear exploration; physicians can start exploring from any areas of interest.

There are a few visualization systems designed for tablet computers, specifically touch-input small screens. Two recent notable systems are TouchWave [4] and TouchViz [5] but both primarily propose new interactive techniques for specific visualizations. There is an EMR system on tablet computers called Allscripts Wand [6]. However, unlike ours, it does not employ a machine-learning-assisted information organization that allows non-hierarchical exploration and displays

- Puripant Ruchikachorn is with Stony Brook University and Chulalongkorn University. This work was done during an internship at IBM Thomas J. Watson Research Center. E-mail: pruchikachor@cs.stonybrook.edu.
- Jennifer J. Liang is with IBM Thomas J. Watson Research Center. E-mail: jjliang@us.ibm.com.
- Murthy Devarakonda is with IBM Thomas J. Watson Research Center. E-mail: mdev@us.ibm.com.
- Klaus Mueller is with Stony Brook University. E-mail: mueller@cs.stonybrook.edu.

only related entities as needed.

Machine learning has been used in visual analytics to complement human reasoning. One example of such systems is WatsonPaths that analyzes unstructured data to explore possible outcomes of clinical analyses. We use the same proven question-answering architecture to analyze clinical notes to extract relations for clinical visit preparation.

3 DESIGN REQUIREMENTS

As the target users are physicians and the main task is visit preparation, our system principally aims to give a quick overview with technical terms in a limited time frame before seeing a patient and possibly necessary data drill-down in some cases. Unlike post-visit analysis, this system does not need to support unexpected relation discovery which may require more advanced statistical or visualization tools.

We conducted interviews in 5 sessions with 12 physicians in total to confirm this. They explained their routines, problems with the current EMR systems, and desired functionality. We compiled the comments for visit preparation task and their major problem was that they were inundated with impertinent information.

We want to create a system for visit preparation on tablet computer because physicians can often be busy in other wards and run out of time to do pre-visit routine in examination room. With tablet computer, physicians can quickly access patient's information in other areas even from home in case of emergency.

Given the form factor and task, the design requirements can be listed as follows:

- R1 *Design for small form factor*: Our previous system design [7] was a dashboard web application that needs at least 1,650-pixel-wide screen. The widest logical dimension of Apple iPad, one of the most popular tablet computers, is only 1,024 point. Simply scaling the previous system would not work well as its typeface and other visual elements would be too small to read. The display needs to be not only simple and easy to understand but also compact and readable.
- R2 *Touch-friendliness*: Tablet computers require larger input target size than traditional computers due to physically larger input pointer i.e. finger. Some interactions such as mouse hover are also unavailable for touchscreen.
- R3 *Drill-down information architecture*: There are many medical entities to be explored and they all cannot be shown at once like in our previous dashboard approach, especially in small screen (R1), so the information has to be organized and allows the users to drill down to retrieve specific relations. This can be facilitated by machine learning to show only relevant information.
- R4 *Interactivity*: To support the information architecture (R3), the system needs to be interactive to show and hide information at will. Animations are used as necessary to give the impression of continuity.

4 USER INTERFACE DESIGN

There were a few iterations before the final design. The relations among medical terms suggested a node-link diagram with a force-directed layout which would be useful in finding data clusters and post-visit analysis. We opted for a simpler design, item lists and interactive highlights to show relations.

The final user interface design has four main parts, namely patient profile, item lists, timelines, and clinical notes. All but clinical notes are shown together in figure 3. Besides the patient profile which persistently shows critical information, the overall interaction structure is hierarchical. Satisfying the requirement R3 to show only relevant information as little as possible at a time, the user starts from the item lists, may drill down to timelines, and then clinical notes. The system supports overview+detail as item lists and timelines and details on demand to bring up medical notes to see more written details or confirm machine learning results.

The requirements R1, R2, and R4 are generally applied to the final design. Each user interface component is explained below in detail.




Fig. 2. Given that each circle represents a visualization state, this chart illustrates how the user interacts with the item lists to drill down and select their items of interest.

4.1 Patient Profile

A new patient's medical record can be queried from their identification number by tapping on its field and edit it. All identifying informations—name, age, sex, identification number, date of birth, and date of admission (for inpatient)—are collected in one place and stay on the screen at all time. This can help reducing medical errors due to the confusion of multitasking between multiple EMRs. For safety concern, there is also some critical information such as allergies highlighted in a bright red background. Shown at the bottom (or the left side when rotated) of figure 3, this panel is generally arranged by change frequency from the most to the least immutable.

The latest vital signs are also included to give the physician a quick overview and indicate any urgent needs to be addressed. We currently have only the most recent measures but we can later introduce sparklines [8] to show vital trends.

4.2 Item Lists

All lists are minimized and folded by default. Each list can be opened and closed as needed. Currently there are three list states: closed, half-open, and open, denoted by empty, half-filled, and filled circular markers on the top left corner of each list. They show list title (folded and rotated to save screen estate), only the first-level list items (for hierarchical lists), and all list items, respectively. We plan to add different data representations or aggregations for various list states or sizes, similar to responsive website design.

The order of items in each list is the result of hierarchical clustering of dissimilarity distance among items. This groups similar items in close proximity.

When touched, each item and its related items in the other open lists are highlighted in different hues, which are used throughout the system such as in timeline markers. The user can touch and move their finger around to browse relations.

Tapping on each item slides up its corresponding timeline. More items can be dragged to the timeline section to be compared over time. At this stage, all related lists are open and only related items are shown in order to minimize the number of items on screen. Tapping the related items traverses the relations. For example, from figure 3, if the user is interested in and tap "Glucose" in *Lab Tests*, it will become the focus and its related problems—"Pancreatitis, chronic", "Diabetes mellitus", and "Hyperlipidemia" for this particular patient—will be shown instead. Tapping the selected item again triggers all items to be shown.

An obvious drawback of showing only related items is the inconvenience of adding two unrelated items to the timeline, possibly to investigate an unusual cause and effect. Figure 2 explains how the user can conceptually pick a certain set of items. After filtering irrelevant entities, the user can either go back up in the information hierarchy to the overview and drill down again or open other lists of unrelated items and "jump" to those subtrees, which are subsets or partial views of the relation graph nodes. This flexibility of graph traversal resembles zooming in geographical mapping application.

4.2.1 Machine Learning

Watson as a question answering system can analyze unstructured data and extract relations between terms. With this relation scoring method, medical terms can be linked from clinical notes. To reduce the user's cognitive load, we then use this to create a relation graph and show

only neighboring nodes (relevant items) when a node is selected. Some relations are unconventional; for instance, “Metformin” is for the treatment of “Diabetes mellitus” but also has some effects on “Hyperlipidemia” and “Obesity” so they are connected.

4.2.2 Relation editing

The user can drag between items of different lists to link them together; they will become highlighted when one of the two newly linked items is touched or tapped. The user can edit the relations as they see fit and this will make the lists function as a personal workspace or virtual mental map of how physicians connect medical data.

4.3 Timelines

Each item has its own timeline whose horizontal axis is temporal and vertical axis represents normalized values. The axes and their labels are adjusted to fit the oldest and latest events and stay readable. With alternate background colors, neighboring timelines are distinguishable. Labels on the left side show item names and they are in color if they are linked to clinical notes as described later. Vertically aligned on the right side to the most recent data point of each timeline, the first and second columns show the latest value and date, respectively.

We follow the convention of LifeLines 2 [9] to show all events in triangular markers that are still discernible even when overlapped. The markers are color-coded according to the lists that they belong to, except those whose values are lower or higher than recommended minimums or maximums (currently applicable only to “Lab Tests”). Any timeline with a recommended range has a solid light green background as a visual indicator of the normal range. A ranged data sample is represented as a straight horizontal line connecting two markers.

To comply with the requirement R2, each marker has an invisible larger touch target. Supported is detail-on-demand tooltip of each marker such as value, its reference range, event date (or range) as shown in figure 3. Each event of *Problems* and *Encounters* is directly related to a clinical note so tapping on its marker opens up the corresponding note.

4.4 Clinical Notes

When a certain item label or marker is tapped, related clinical notes, if available, are shown over both the lists and timelines. The user can scroll to read and tap again to close the notes and go back to the list and timeline view.

Note that the clinical notes are actually the input of the system but intentionally structured at the bottom level of the information architecture. For physicians who rely on clinical notes over other forms of data, they can use the user interface to navigate and drill down to specific notes of interest for supplementary or confirmatory purpose.

5 IMPLEMENTATION AND EVALUATION

We implemented this system as a web application with D3.js and wrapped it with Apache Cordova into tablet form factor for later evaluation. The machine learning result has been conducted offline and provided via an internal web service API. The user interface runs at interactive speed on any recent Apple iPad.

To informally evaluate this system, we used a scenario of endocrinologist seeing a new 70-year-old patient, whose profile is shown in figure 3, to walk through our implementation as follows. Note that different physicians may have different review order preferences.

General examination The physician adds “Primary care” and “Internal medicine” to the timeline to see the most recent note, then checks the vital signs on the patient profile which are in their normal ranges. Then the physician opens up *Problems*, *Medications*, and *Lab tests* lists to identify any endocrine problems. “Diabetes” and “Hypothyroidism” are identified.

Diabetes To review the history of the problem, the physician selects “Diabetes mellitus” to identify when the problem was diagnosed. Related lab tests and medications—“Hemoglobin 1c”, “Metformin”, and “Glipizide”—are added to the timeline. These three timelines

clearly show that the problem was severe in the beginning, the patient took the first medication with increased dose before taking the second medication of a constant dose from mid-2007. Since then, his condition has been improving, become stable, although not under optimal control, as demonstrated by red markers across glycated hemoglobin (Hb A1C) timeline. The physician checks the clinical notes of the mentioned significant events for more details.

Finally, the physician goes back to the problem list to identify any diabetic complications. “Diabetic retinopathy” and “Microalbuminuria” are added to the timeline to review further details

Hypothyroidism The physician clears the timeline and adds “Hypothyroidism” instead. Levothyroxine, Synthroid, and thyroid-stimulating hormone (TSH) are the only related items of the problem in *Medications* and *Lab Tests* respectively. Since 2011, after two low lab tests, the patient has switched from the former to the latter medication. The latest value of TSH in 2012 was back to normal.

6 CONCLUSION AND FUTURE WORK

We have developed an EMR application that fulfills our requirements for tablet computers. Its information organization centers around medical entities whose relations are discovered by state-of-the-art machine learning. From the user interface viewpoint, the relations from machine learning provide a form of subtle suggestions and mixed initiative interaction. We plan to add more machine learning features such as semantic searching and reasoning in the future. The data of user interactions can also be fed back to our machine learning algorithm in the manner of supervised learning for better accuracy.

The current application of our system is mainly for clinical visit preparation that includes only the first few steps of the clinical workflow. We plan to support the whole workflow and consider additional features such as security for personal and sensitive medical data.

Also we need to conduct a more formal evaluation in a hospital work environment. Unlike our laboratory setting, there may be some other requirements we could not capture in interviews. We hope this Watson-powered drill-down user interface and tablet form factor become a fruitful combination for medical purposes.

REFERENCES

- [1] A. Rind, T. D. Wang, W. Aigner, S. Miksch, K. Wongsuphasawat, C. Plaisant, and B. Shneiderman, “Interactive Information Visualization to Explore and Query Electronic Health Records,” *Foundations and Trends in Human-Computer Interaction*, vol. 5, no. 3, pp. 207–298, 2013.
- [2] C. Plaisant, B. Milash, A. Rose, S. Widoff, and B. Shneiderman, “LifeLines: Visualizing Personal Histories,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '96)*, 1996, pp. 221–227.
- [3] Z. Zhang, B. Wang, F. Ahmed, I. V. Ramakrishnan, R. Zhao, A. Viccellio, and K. Mueller, “The five Ws for information visualization with application to healthcare informatics,” *IEEE transactions on visualization and computer graphics*, vol. 19, no. 11, pp. 1895–910, Nov. 2013.
- [4] D. Baur, B. Lee, and S. Carpendale, “TouchWave : Kinetic Multi-touch Manipulation for Hierarchical Stacked Graphs,” in *Proceedings of the 2012 ACM international conference on Interactive tabletops and surfaces (ITS '12)*, 2012, pp. 255–264.
- [5] J. M. Rzeszutarski, “TouchViz : (Multi) Touching Multivariate Data,” in *CHI '13 Extended Abstracts on Human Factors in Computing Systems (CHI EA '13)*, 2013, pp. 1779–1784.
- [6] Allscripts Healthcare Solutions Inc, “Allscripts Wand.”
- [7] M. Devarakonda, D. Zhang, C.-h. Tsou, and M. Bornea, “Problem-Oriented Patient Record Summary : An Early Report on A Watson Application,” in *International Conference on E-health Networking, Application & Services (IEEE Healthcom)*, 2014.
- [8] E. Tufte, *Beautiful Evidence*. Graphics Press, 2006.
- [9] T. D. Wang, C. Plaisant, A. J. Quinn, R. Stanchak, B. Shneiderman, and S. Murphy, “Aligning Temporal Data by Sentinel Events : Discovering Patterns in Electronic Health Records,” in *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 2008)*, 2008.




Fig. 3. Our system shows the medical record of a patient (whose name has been changed). Dates of birth and admission are not available and instead show placeholder text. At the present state, only related items to diabetes are in the lists and some of them are dragged into the timelines to investigate their values over time. A tooltip is present on an out-of-range Hb A1C data point. The size of this figure is the same as our system's actual use on Apple iPad.