Understanding Care Plans of Community Acquired Pneumonia Based on Sankey Diagram

Shunan Guo[1,2], Ph.D Student, Chaoguang Lin[2], Bachelor,
David Gotz[3], Ph.D, Bo Jin[1], Ph.D, Hongyuan Zha[1], Ph.D
Linhua Shu[4], M.D, Nan Cao[2], Ph.D
[1]ECNU, ShangHai, China; [2]NYU, ShangHai, China;
[3]UNC, NC, USA ; [4]SHChildren, ShangHai, China

Abstract

A care plan is a sequence of medical interventions formulated for curing a specified disease. Doctors usually craft different care plans for different patients based on their knowledge and experience as well as following the clinical guidelines. An intuitive summarization of previous successful care plans not only provides doctors a reference of successful treatment guidelines, but also helps with the detection of anomalous treatments and the improvement of existing care plans. However, producing such kind of summarization is challenging due to the complexity of the data, i.e., large, temporal oriented, and multidimensional. In this paper, we propose a Sankey diagram-based visualization design to visually summarize care plans based on our medical collaborators' requirements. We apply our tool to a medical dataset of pneumonia patients collected from a Children's hospital in Shanghai, China. Based on the visualization results, doctors detected many interesting findings, which will be discussed in the paper.

1. Introduction

Medical care planning is a critical step in the care delivery process, contributing directly to a patient's treatment pathway and associated outcomes. A care plan is a sequence of medical interventions formulated to help patients manage their diseases. It is often crafted by doctors after a patient is diagnosed according to their intuition and experience, and is often based on recommendations in clinical guidelines. However, the correct care plan is not identical for everyone with a given medical condition due to the individual differences between patients. Thus, in order to help patients obtain better treatment, it is often necessary for doctors to customize the treatment plans recommended in clinical guidelines.

The care plan customization process can be aided, in part, by examining the efficacy of alternative care pathways as experienced by previously treated patients. However, there are several complex issues which make this form of care plan analysis hard to accomplish. First, the large numbers of patients as well as the vast variety in drug types and other variables in medical data increase the complexity of datasets. These issues make the comparison and aggregation over large numbers of patients difficult if not impossible with traditional tools. Second, care pathway data contains temporal events, such as all medications and procedures, which make such data even more difficult to analyze. Third, it is important to not only consider the pairwise correlation of variables across patients, but also the combination of these variables (e.g., multiple medications or diagnoses, often overlapping in time).

Although there are some existing solutions [6-8] for displaying time-oriented medical data, many fail to handle the case of simultaneous or overlapping medical treatment. The typical serial time-oriented data visualization techniques often used for medical data are therefore not quite suitable for such situations. There are also methods that have focused specifically on representing medical care plans [1-4], which inspired our work. However, unlike our work, most of these methods formulate care plans based on existing treatment guideline instead of medicines, thus missing details such as when did what medicine has been used for curing a disease.

In this paper, we present a visual encoding schema based on the Sankey diagram to enable the analysis of treatments for a cohort of patients to examine differences in outcomes for variations in care plans. In this visualization, we summarize a patient's care plan based on a layered graph model in which each node indicates a medicine (i.e., a treatment) and different graph layers indicates different time and links are used to connect the same treatment at the different time. Similar path plans of different patients are aggregated together to form patient cohorts that facilitate the comparison between pathways for finding better or worse medical care plan in a single clinical stage or the entire course of disease. Rich interactions are also designed to support an efficient data exploration and filtering. Throughout the paper we describe our design using a motivated problem related to Community Acquired Pneumonia (CAP). We include five sample analysis after discussed with CAP experts to explain the insights found through the visualization. The major contributions of this paper are as follows:

1. Apply Sankey Diagram to real-world medical care plan data and verified its applicability after the evaluation of domain experts.
2. A novel visual encoding method based on Sankey Diagram to display electronic medical care plan records.

The rest of the paper is organized as follows. We review the related work in section 2 and describe our motivation problem in section 3. We introduce our visual encoding design in section 4 and demonstrate preliminary analysis sample and evaluation results from domain experts in section 5. And finally conclude in section 6.

2. Related Work

In this section we review the papers related to work, which basically includes techniques developed for visualizing temporal event sequences and medical care planning.

2.1 Visualizing temporal event sequences

Generally, this topic lies in the direction of visualizing time-oriented data, which is comprehensively discussed in [5]. We focus on the techniques developed for representing event sequences in the healthcare domain. Sankey Diagram is one of the most intuitive and commonly used methods for representing the event sequence, which is also adopted in our visualization design. Besides this approach, many designs such as LifeLines[6], Timelines[7], Eventflow[8] and many other timeline-based representations[9-12] align event sequence horizontally along a timeline in which one patient record is split into different event categories, thus making them inefficient for cohort analysis. LifeFlow[13] aggregates health records of multiple individuals based on a Treemap but each event type are divided into several pieces and is hard to analyze in general. There are also visualization tools like[14-20] aggregates multiple records based on data transformation and mining techniques, DecisionFlow[21] supports an in-depth analysis of a heterogeneous multidimensional event sequence at different phrases via rich interactions and flow based visualization design. Although powerful, none of these techniques are specifically designed for revealing a care plan in the electronic health records.

2.2 Visualizing medical care planning

Among many visualization techniques developed for representing electronic health records [22-25], visually representing care plans attracts more and more research interests in recent years due to its usefulness for supporting clinic process. Most projects dealing with representing care planning are based on flow-chart algorithms, which is widely known by physicians and requires minimal learning efforts. For instance, CareVis[4] and AsbruView[3] gives visualization solution for large and complex flowchart. GapFlow[1] shows the derivation during different medical treatment plans. Most of these techniques take the existing care plans as the input, which are not always available. Our work is largely inspired by CareFlow[2], which is designed to assist doctors in finding better care plans based on the treatment records. We adopt its visual design to help analyze a group of patients. Different from CareFlow, which shows cohorts' physical outcomes after applying a sequence of treatments, our work improves this design by strictly aligning the treatments along a timeline, which help illustrate the combination or correlations among different medicines used at the same and different time.

3. Motivation Problem and Dataset

Community-acquired pneumonia (CAP) is one of the most common infectious diseases and has been recognized as a potentially lethal condition for nearly two centuries[26]. It is also a serious infection that afflicts children throughout the world. The average annual incidence of pneumonia in children younger than 5 years of age is 34-40 cases per 1000, and is increasing every year, generally becoming the largest killer of children[27]. Symptoms suggestive of pneumonia basically include 80 percent of fever combined with respiratory symptoms such as cough, sputum production, pleurisy, and dyspnea. The pathogens responsible for community-acquired pneumonia in children mainly includes mycoplasma, influenza virus, and bacteria. However, there have been few attempts to devise treatment guidelines in China. Guidelines from North America and Europe are not practically useful enough for doctors in China due to different infection environment and etiologic process. And the treatment guideline for CAP in China is lack of support of population statistics, or slightly out-of-date because of pathogenic variation. Besides, the antibiotic abuse is widespread is the medical treatment of CAP in China, which results in greater potential hazards.

With the above issues in mind, our medical collaborators are interested in analyzing the past clinical cases, in order to find care pathways that are mostly used or those anomalous or irrational ones, and their corresponding outcomes. Furthermore, doctors also wish to find the proportion of patients for different pathogens and their correlated physical signs, as well as physical reactions to different drugs. To address these problems, we are given access to a dataset of nearly 3,000 children patients with pneumonia and with various medical treatments based on nearly 100 types of medications. All these patients were hospitalized. The patients' anonymized information, prognosis results (i.e., pathogen types), the full treatment history (e.g., medicines) during the hospitalization, and the monitoring of the patients' body temperatures was also given.

4. Visualization Design

In this section, we first present certain design tasks stemmed from our discussion with two doctors. We then provide a detailed description of our data aggregation and visual encoding methodology. Finally, we introduce a couple of user interactions for data exploration through graph manipulation.

4.1 Design Tasks

Clinical CAP physicians are often faced with the difficulty of making precise care plans for different patients. Currently, an efficient care plan can only be made by these doctors manually by using their own experiences and domain knowledge to bridge the outcomes to the corresponding medical treatments. This procedure is usually extremely inefficient and time consuming. Therefore, the doctors expected a tool or even a system which could provide references from the existing and similar care plans to support their decision making process. This is especially important in China as most of the doctors need to handle hundreds of patients every day. The doctors are particular need a system that can help them automatically integrate the treatment with the outcomes and show existing care plans that are related to a focal disease for their reference. They hope this tool can also help to summarize the existing care plans so that they can easily identify which one is more efficient in terms of curing the disease. To meet their requirements, we compile a list of visualization design tasks as follows.

T1 Integrating the medical treatments and outcomes of each patient. Both medical treatments and the outcomes such as the monitoring results of the body temperatures are recorded independently over time. Therefore, for each patient, we need to align different data records based on their timestamps to build a medical care path so as to identify the order of medicine being used and corresponding physical sign afterwards.

T2 Clustering patients with similar medical care paths. In order to proceed care plan comparison and cohort analysis, our visualization needs to gather similar paths and show average physical outcomes. In this case, different care plans are automatically distinguished, saving the doctors from looking through thousands of individual pathways.

T3 Revealing the statistical details of each treatment. The statistical details of treatment include the number of involved patients and the frequency of a medicine being taken. This helps doctors to quickly differentiate the critical care plan for groups of patients and anomalous plans for different patients.

T4 Associating care plans with pathogens. Finding correlations between the treatments and different pathogens are considered to be very important as it will help doctors to make a correct decision at the early stage. Therefore, the proposed system should be able to differentiate different pathogens of a focal disease and associate it with different care plans. More precisely, the system should be able to illustrate the proportions of patients for each pathogen type and their corresponding reaction after taking each treatment.

T5 Facilitating visual data filtration. Considering the actual care pathways can be chaotic due to the complicated situations the doctors may face, the system should enable doctors to filter exceptional cases and explore data through interactions in order to reveal the main patterns.

T6 Easy browsing of raw data. The raw data, such as the name of a medicine, a pathogen type, and specific amount of the cohort's size can help doctors to a better understanding a care plan. Some other features of the patient, such as the lab test results, the exact date of hospital admission and discharge can be of great value for doctors to perceive the essence beyond the visual representations. Thus the visualization should enable analysts to explore raw data easily.

4.2 Data Model

With the above dataset and design tasks, our goal is to find and differentiate care plans that are used among groups of patients and compare their outcomes. We approach this goal by aggregating all the patient records based on how the treatments are performed. Specifically, patients take the same medicines at the same stage (defined by the number of days in hospital) are grouped together. Thus the grouping results at different stages forming a summary of the treatment history as shown in Figure 1. In this data model, nodes indicate medicines and links indicate groups of patients, which have two primary

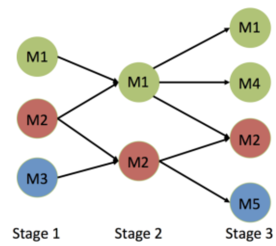

Figure 1. The data model for summarizing the care plans

attributes: the size of the patient group and the outcome (e.g., body temperature) after taking the previous treatment.

As a result, the final directed-acyclic-graph captures every existing medical care pathway as well as the number of patients flowing through the same sub-paths, and the average body temperature during each time interval. This graph serves as the prototype and preparation of our Sankey Diagram implementation.

4.3 Visual Encoding

The above data model can be intuitively visualized by a Sankey diagram and the aforementioned tasks guide our design of the visual encoding schema. Figure 2 illustrates our visual encoding design. It illustrates (a) the proportion of patients taking each medicine at the same time interval, (b) the proportion of patients for each pathogen type, and (c) the average body temperature for each time interval after or before a treatment was taken. In particular, this visualization consists of several components, which are described as follows:

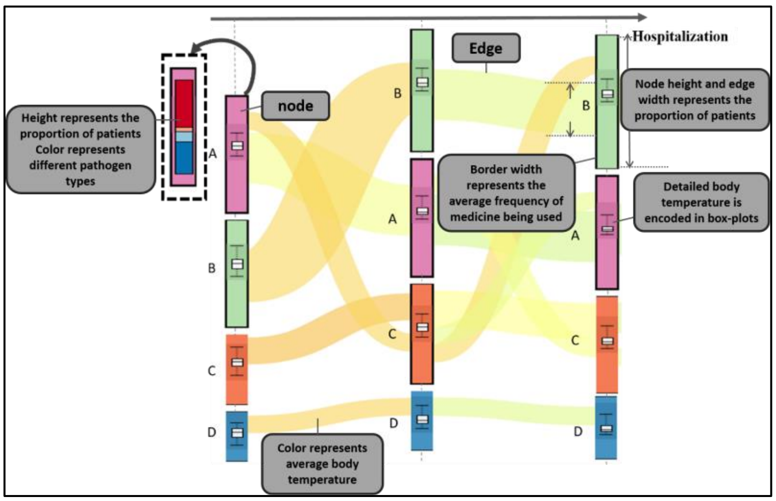

Figure 2.Nodes represent treatments and are positioned horizontally in temporal order. Edges connect treatments in each day and are color-coded to represent the average body temperature.

Nodes (Treatment): Nodes represent treatments and are positioned along the horizontal axis indicating treatment sequence over time. With different filling color representing different kind of medicine and the height representing the proportion of patients taking a given medicine (**T3**). And the border width of each node represents the frequency of this medicine being taken in a day. For example, as illustrated in Figure 2, there is a total of four types of medicine being taken during hospitalization, A, B, C and D. Medicine A and B are used by large and almost equal proportion of patients, while medicine D is used by the smallest cohort.

Layer: The overall view is horizontally divided into several layers according to the maximum days of hospitalization. Each layer represents one day of hospitalization, showing the treatment information from the first day in the hospital until the last. And the nodes of each layer represent the treatments being taken on that day. For example, as shown in Figure 2, the longest hospitalization of this group of patients is three. And for each day, there are all four types of medicine being taken. The vertical adjacency relation of each pair of nodes represents the correlation between medicines, which was implemented using force-directed layout with back and forward propagations. In other words, medicines that are frequently used as combination are more likely to be shown in the neighborhood. For example, the position of medicine A and B switches vertically in day 2, probably because the amount of people using medicine A and C slightly increases edging out the combination of medicine B and C.

Edge (Temperature): The edges are colored according to the average body temperature of all patients represented by the corresponding source node. Elements that are colored red represents parts of the care plan where patients are in the state of fever, whereas elements colored green are care plans where the patients' body temperature becomes normal. The width of edges is correlated with the source nodes and the target nodes, representing the number of patient flows through this particular edge.

Node Encoding: There are two types of filling for each node and doctors can switch one to another through interaction. The first is box-plot, which shows the variance of body temperature in the group of patients and to some extent avoids the temperature from being balanced through averaging. The box-plot is displayed in a shadowed area with certain height, and is visible only when the cohort size reach the threshold in order to maintain its comparability. The other filling is the pathogen proportion. Each color represents a type of pathogen and the height is proportional to the number of patients (**T4**).

4.4 User Interactions

As expressed in our design tasks (**T5**), user interactions are key to data and pattern exploration. Our visualization allows doctors to interactively perform the following actions.

Highlighting: As shown in Figure 3, doctors can highlight the overlap edges of the corresponding group of patients through hovering mouse over an edge, a node or a type of pathogen. This helps doctors to ascertain the actual path that the cohort goes through.

Hovering: Hovering also triggers the display of tooltips which provides the exact name of a type of medicine or pathogen as well as the specific number of patients (Figure 3).

Filtering: Doctors can filter data in two ways. First, doctors can select a cohort of patients and construct a new Sankey Diagram through double clicking on edges and nodes. This can help doctors to get a better analysis on the particular cohort they are interested in. Second, doctors can use the bidirectional slider to filter both nodes and edges on the number of patients to remove small subgroups. The graph will then show only the care plans that are mostly taken by removing thin edges and unconcerned nodes. Also the doctors can filter thick edges to observe the exceptional cases.

Raw data exploration: Doctors can filter the items displayed in raw data list through the filtering box on the top, or by clicking an edge or a node to select a cohort of patients.

Encoding Switching: As described in 4.3, doctors can choose whether to display pathogen information or not. And the fillings of the nodes can be changed through the switch on top.

5. Analysis of Pneumonia Care Plans

5.1 Domain Expert Interview

Based on the above visualization design, we demonstrate our system in front of a medical expert team based on a core dataset cleaned from the raw data. The medical expert team, led by the director of the department of respiratory disease, has rich experiences in both CAP clinical and treatments. The core dataset used for the system evaluation and expert interview consists of 953 patient records with a complete using records of 24 types of medicines collected in one year. The interview starts with a tutorial covering the visualization design and interactions. We then ask doctors to use the system on their own for exploring the core dataset. After a full understanding of the system's capability, each doctor was given a list of questions as the guideline of system evaluation. Doctors are asked to provide their feedbacks and suggestions or raise any questions during the process of using the system. The interview lasted approximately 2 hours. We recorded the entire conversation, and took notes of their comments. The rest of the section will give a brief summary of doctors' comments. And we list some of the medical related findings in section 5.2 and doctors' suggestions on system improvement in section 5.3.

Both doctors are very impressed by the volume of information that the visualization provided. They commented that "problems can be displayed intuitively through the view", and "Comparing to traditional statistical analysis method, this aggregates different types of statistics data with more detailed value in each hospitalization day". Even though they mentioned that "since the system provides such huge amount of information, it will take some learning efforts to master the use of the system". However, they believe "this tool can be very convenient and efficient in exploring large datasets once you get used to it". In addition, the first doctor suggested that "we should popularize this system to a larger platform in response to the trend of 'Precision Medicine'". He also expressed his alacrity of providing more data for comparison to see whether the patterns of care plan evolves through years. The second doctor is particularly fond of the idea of displaying care plans and all sorts of statistical data in this way. She said, "We used to display statistical data in forms of pie charts or histograms. However, this system is much more powerful. It is not only capable of displaying multiple types of statistical data integrally, but even the way they change over time."

5.2 Preliminary Analysis Result

In this section, we describe the doctors' feedback and demonstrate several interesting findings detected by the doctors.

F1 The system is effective in displaying major care plans. As shown in Figure 3, when the doctor hovers on the edge from medicine A (t1) to medicine A(t2), edges from and to medicine B are also highlighted. Considering the width of highlighted edges, they found that nearly 50% of the patients taking medicine A at day 2 and 3 are also taking medicine B at the same time. In particular, medicine A is Ceftriaxone and medicine B is Azithromycin. The combination of Cephalosporin drugs and Azithromycin is a common approach for dealing with children respiratory diseases and works pretty well. The same pattern is also found with medicine C, which is another type of Cephalosporin drug.

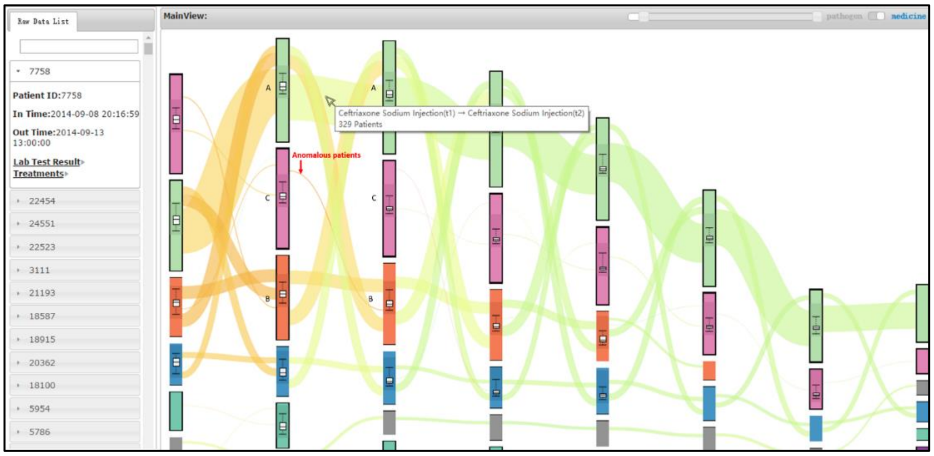

Figure 3. Overlap edges are highlighted after hovering, showing how the care pathways flow through the rest of the graph. Edges between node A and B are also highlighted after hovering the edge from A to A, indicating that medicine B are mostly taken along with medicine A. And few edges from medicine C is highlighted indicating medicine A and C are usually not taken together. Doctors can refer detailed information of the anomalous patient through raw data list on the left by clicking the corresponding edge.

F2 Anomalous care plans. The doctors also found several suspicious care plans. For example, as shown in Figure 3, medicine B is seldom used. Sometimes a thin edge comes from medicine B, according to our highlight mechanism, the doctors found a small group of patients take both medicine A and B at the same time. However, both A and B are Cephalosporin drugs, the doctors believed that the mixed used of these medicines are problematic, which worth a further inspection of the raw data.

F3 Verifying statistical results. The doctors are very familiar with the dataset and have already conducted some simple statistical analysis and illustrate the results base on pie charts and histogram. The result showed that 31% of the patients are taking Cefuroxime (medicine C) while 37.15% of the patients are taking Ceftriaxone (medicine A). From the Figure 3, we can see that medicine A and C are used by a large proportion of patients and their heights are substantially equal. This fully coincides with the previous statistics. However, our visualization design also shows the variety of basic statistics over time, which was preferred by the doctors.

F4 Revealing the correlation between pathogen and treatments. As illustrated in Figure 4, the doctors switched fillings of the nodes into pathogens proportions and they found a large proportion of mycoplasma patients taking Cephalosporin drugs. This is a surprising finding as Cephalosporin drugs take no effect in killing mycoplasma viruses. This also suggested problematic care plans as many doctors didn't take pathogens into consideration while making a prescriptions and most mycoplasma patients are not taking the right medical care plans at the very beginning. Thus, it is not hard to explain why the portion of mycoplasma patients maintain fever until the fourth day of hospitalization, while the patients taking Amoxicillin Sodium and Potassium recover in the second and third days (Figure 5).

F5 Revealing the care path of specific cohorts. The doctors also inspected the treatments used by a small group of patients used Meropenem and had extremely high fever and were cured after a long treatment. As shown in Figure 6, most patients recovered on the 7th day and one patient was still in the fever. In addition, Meropenem is an advanced antibiotic drug, which were only used on patients with very severe symptoms. Thus, this specific cohort must have very serious disease and are not easy to recover.

5.3 Discussion

Apart from patterns observed in the system, doctors also provide their suggestion on improving the system. First, they think it's necessary to provide filter for doctors to select patients with a certain type of pathogen. Because the symptom of pneumonia is very much pathogen-related. Sometimes doctors are more interested in analyzing the features of one specific pathogen in order to improve its corresponding care plans. Second, they believe the pre-selection of medicines is also needed. Since in most cases, doctors are only interested in a particular group of medicine. For example, during the interview, both doctors are very eager to find out the usage of Cephalosporin medicines, however, there are two Cephalosporin medicines and the system does not provide medicine selection. Even though they can capture most of the information from the general view, they believe patterns will be revealed more clearly if those two medicines were displayed alone.

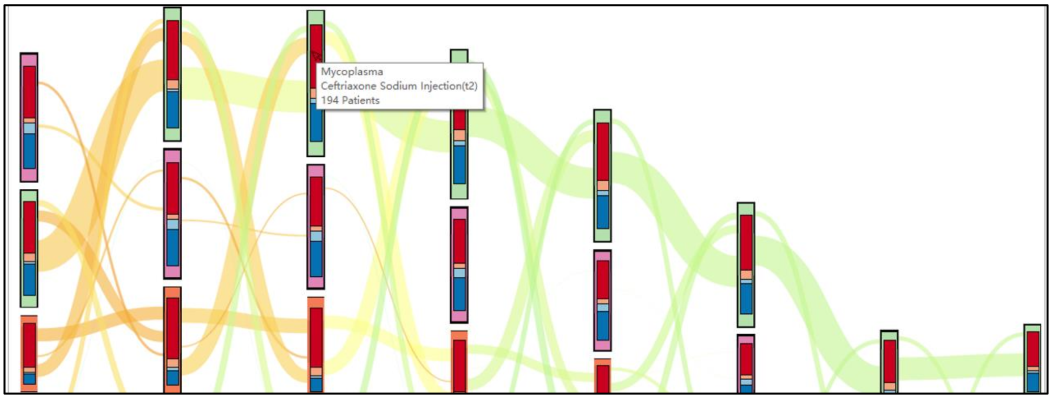

Figure 4. Mycoplasma patients account for a large proportion of each medicine type and usually hospitalize with high body temperature.

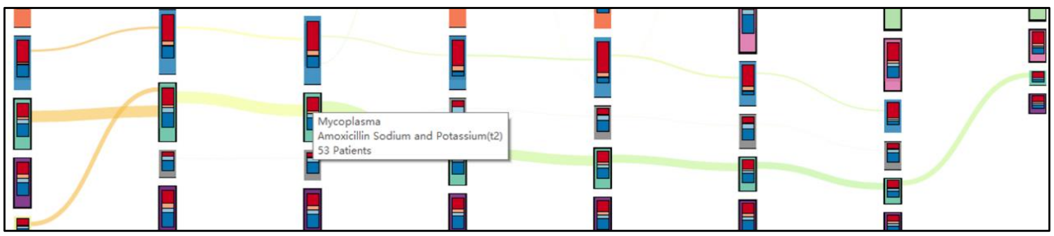

Figure 5. Mycoplasma patients taking Amoxicillin Sodium and Potassium turns out to recover faster than those taking Cefuroxime comparing to figure 4 due to mycoplasma virus's immunity to Cephalosporin drugs.

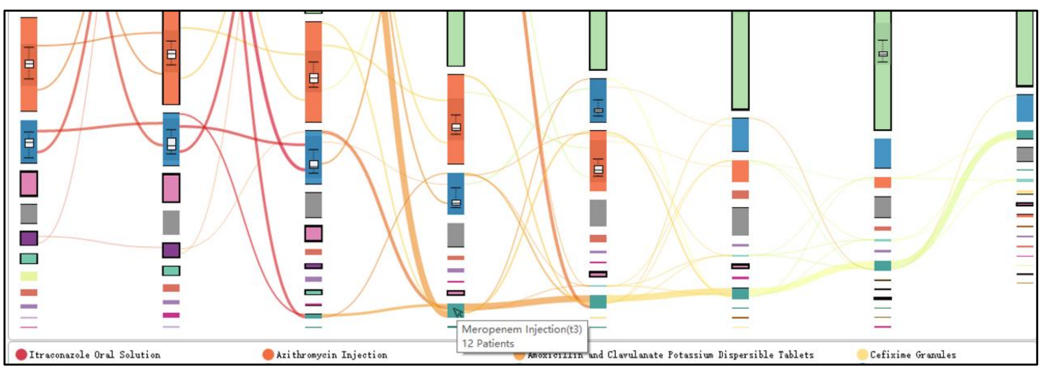

Figure 6. Patients taking Meropenem trends to have extremely high body temperature when hospitalize, and take long time to recover. The observation is confirmed by the fact that Meropenem is an advanced type of antibiotic and is only taken by patients with very severe symptoms.

6. Conclusion

In this paper, we presented techniques for representing medical care plans based on Sankey diagram. In our design, we aggregate different treatments of a focal disease into a layered directed graph. In this graph, each node represents a medicine (i.e., the treatment) and links are used to connect the same medicines used at different time and encode the corresponding outcomes (e.g., body temperatures). Our visualization has been used to summarize care plans of community-acquired pneumonia based on a patient dataset collected from a children's hospital in China. Based on the visualization, many findings such as anomalous care plans were detected by doctors, which verifies the usefulness of the tool. The future work includes proposing advanced algorithm for visual clutter reduction and improving the design to illustrate the situation in which multiple medicines are taken.

Acknowledgements

This work is supported by the Shanghai Committee of Science and Technology, China (Grant Nos. 15JC1401700).

References

1. Gotz, D., et al., GapFlow: Visualizing Gaps in Care for Medical Treatment Plans. IEEE VIS Poster (2013), Atlanta, Georgia.
2. Perer, A. and D. Gotz. Data-driven exploration of care plans for patients. in CHI'13 Extended Abstracts on Human Factors in Computing Systems. 2013. ACM.
3. Kosara, R., et al., AsbruView: capturing complex, time-oriented plans—beyond flow charts, in Diagrammatic Representation and Reasoning. 2002, Springer. p. 535-549.
4. Aigner, W. and S. Miksch, CareVis: integrated visualization of computerized protocols and temporal patient data. Artificial intelligence in medicine, 2006. 37(3): p. 203-218.
5. Aigner, W., et al., Visualization of time-oriented data. 2011: Springer Science & Business Media.
6. Plaisant, C., et al. LifeLines: visualizing personal histories. in Proceedings of the SIGCHI conference on Human factors in computing systems. 1996. ACM.
7. Harrison, B.L., R. Owen, and R.M. Baecker. Timelines: an interactive system for the collection and visualization of temporal data. in Graphics Interface. 1994. Citeseer.
8. Monroe, M., et al., Temporal event sequence simplification. IEEE transactions on visualization and computer graphics, 2013. 19(12): p. 2227-2236.
9. Aigner, W., et al. PlanningLines: novel glyphs for representing temporal uncertainties and their evaluation. in Ninth International Conference on Information Visualisation (IV'05). 2005. IEEE.
10. Bade, R., S. Schlechtweg, and S. Miksch. Connecting time-oriented data and information to a coherent interactive visualization. in Proceedings of the SIGCHI conference on Human factors in computing systems. 2004. ACM.
11. Cousins, S.B. and M.G. Kahn, The visual display of temporal information. Artificial intelligence in medicine, 1991. 3(6): p. 341-357.
12. Karam, G.M. Visualization using timelines. in Proceedings of the 1994 ACM SIGSOFT international symposium on Software testing and analysis. 1994. ACM.
13. Wongsuphasawat, K., et al. LifeFlow: visualizing an overview of event sequences. in Proceedings of the SIGCHI conference on human factors in computing systems. 2011. ACM.
14. Burch, M., F. Beck, and S. Diehl. Timeline trees: visualizing sequences of transactions in information hierarchies. in Proceedings of the working conference on Advanced visual interfaces. 2008. ACM.
15. Vrotsou, K., K. Ellegard, and M. Cooper. Everyday life discoveries: Mining and visualizing activity patterns in social science diary data. in Information Visualization, 2007. IV'07. 11th International Conference. 2007. IEEE.
16. Fails, J.A., et al. A visual interface for multivariate temporal data: Finding patterns of events across multiple histories. in 2006 IEEE Symposium On Visual Analytics Science And Technology. 2006. IEEE.
17. Vrotsou, K., J. Johansson, and M. Cooper, Activitree: Interactive visual exploration of sequences in event-based data using graph similarity. IEEE Transactions on Visualization and Computer Graphics, 2009. 15(6): p. 945-952.
18. Wongsuphasawat, K., et al., Querying event sequences by exact match or similarity search: Design and empirical evaluation. Interacting with computers, 2012. 24(2): p. 55-68.
19. Wongsuphasawat, K. and B. Shneiderman. Finding comparable temporal categorical records: A similarity measure with an interactive visualization. in Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on. 2009. IEEE.
20. Gotz, D., F. Wang, and A. Perer, A methodology for interactive mining and visual analysis of clinical event patterns using electronic health record data. Journal of biomedical informatics, 2014. 48: p. 148-159.
21. Gotz, D. and H. Stavropoulos, Decisionflow: Visual analytics for high-dimensional temporal event sequence data. IEEE transactions on visualization and computer graphics, 2014. 20(12): p. 1783-1792.
22. Frankovich, J., C.A. Longhurst, and S.M. Sutherland, Evidence-based medicine in the EMR era. N Engl J Med, 2011. 365(19): p. 1758-1759.
23. Ebadollahi, S., et al. Predicting patient's trajectory of physiological data using temporal trends in similar patients: A system for Near-Term prognostics. in AMIA Annu Symp Proc. 2010.
24. Gotz, D. and K. Wongsuphasawat. Interactive intervention analysis. in AMIA annual symposium proceedings. 2012. American Medical Informatics Association.
25. Wang, T.D., et al. Aligning temporal data by sentinel events: discovering patterns in electronic health records. in Proceedings of the SIGCHI conference on Human factors in computing systems. 2008. ACM.
26. Bartlett, J.G. and L.M. Mundy, Community-acquired pneumonia. New England Journal of Medicine, 1995. 333(24): p. 1618-1624.
27. McIntosh, K., Community-acquired pneumonia in children. New England Journal of Medicine, 2002. 346(6): p. 429-437