GapFlow: Visualizing Gaps in Care for Medical Treatment Plans

David Gotz *
IBM T.J. Watson Research Center

Nan Cao †
IBM T.J. Watson Research Center

Esther Goldbraich ‡
IBM Research Haifa

Boaz Carmeli §
IBM Research Haifa

ABSTRACT

Despite the widespread adoption of clinical guidelines (i.e. canonical treatment plan templates that represent generally accepted best practices), significant variations in care are often found across a population of patients. Gaps between the actual treatment programs performed on patients and the recommended guidelines are inevitable given the complexity of disease, differences between patients, and the individualized patient-centered decisions made by clinicians during each encounter. This poster presents a visualization tool designed to help clinical organizations better understand these gaps in care. We describe the input data, our analysis technique to classify individual gap events, and an interactive visualization technique which aggregates and summarizes the results for clinical interpretation.

Index Terms: H.5.2 [Information Systems]: Information Interfaces and Presentation—User Interfaces; I.3.8 [Computing Methodologies]: Computer Graphics—Applications; J.3 [Computer Applications]: Life and Medical Sciences—Health

1 INTRODUCTION

Heathcare providers routinely follow clinical guidelines when designing patient treatment plans as a way to standardize care across a population of patients. However, even when all patients are exhibiting the same disease presentation, variations in care appear between individuals due to differences in disease progression, patient comorbidities, and other factors. Doctors, when choosing which treatment to try next for a given patient, take all of these factors into account.

The complexity of the treatment planning process means that despite widespread adoption of best practice guidelines, gaps in care—differences between a guideline's suggested treatment program and the actual observed treatment program given to an individual patient—are commonplace. Given that these differences exist, medical institutions see major value in learning from past treatment patterns, identifying gaps in care, and understanding how they impact patient outcomes. Moreover, the growing use of electronic medical records within the healthcare industry means that more data than ever is available for analysis to help answer these questions.

However, deriving meaningful insights from raw electronic medical data is challenging. These patient records contain large sets of individual events (e.g., procedures, medications, diagnoses) but do not explicitly capture which guideline steps were followed, altered, omitted, or rearranged. Moreover, often multiple guideline recommendations are applicable for a given condition. This complex information space makes it difficult to extract the desired insights.

We are developing a visual analytics system called GapFlow designed to help address these challenges. Our approach combines a data analytics component with interactive visualization techniques to (1) align actually performed patient treatment events with recommended clinical guidelines, (2) discover and classify gaps in care, and (3) aggregate care gap information across a population of patients and visualize the results for clinical interpretation.

*e-mail: dgotz@us.ibm.com
†e-mail: nancao@us.ibm.com
‡e-mail: esthergold@il.ibm.com
§e-mail: boazc@il.ibm.com

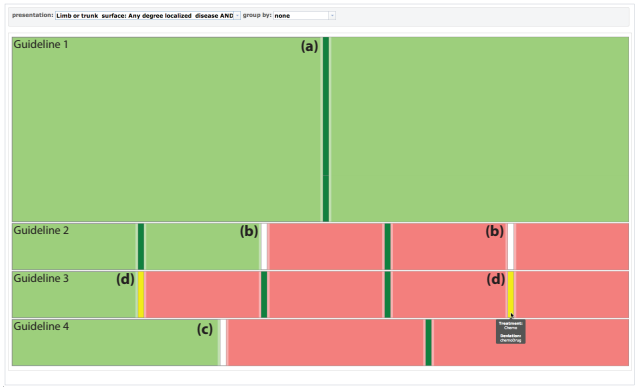

Figure 1: GapFlow showing treatments for a single presentation with four possible guidelines. (a) Roughly half of the patients were given a single treatment according Guideline 1 with no deviations. Meanwhile, missing treatments were found for the smaller (b) Guideline 2 and (c) Guideline 4 groups. Guideline 3 showed (d) two treatments with minor deviations. other groups.

2 METHODOLOGY

In this section we first review the underlying data model and basic terminology for the gap analysis domain. We then briefly overview the gap analysis algorithm before presenting our approach to data aggregation and visualization.

2.1 Data and Terminology

GapFlow visualizes treatment data for a group of patients, which we call a *cohort*. Each patient in the dataset has a single primary disease *presentation* which we use to define specific cohorts. For example, a group of cancer patients might have multiple cohorts, one for each type of cancer (i.e. each presentation) that is observed in the population.

Moreover, for a given presentation, there exist one or more recommended treatment *guidelines*. Guidelines consist of a sequence of individual *treatments* that are recommended for a particular class of patients. For example, a guideline for one presentation of cancer might suggest starting with a surgical procedure before progressing to multiple stages of chemotherapy.

At the time of treatment, a care provider considers the patient's presentation, the recommended guidelines, and the patient's individual circumstances to arrive at a treatment decision. Over time, a patient might receive a sequence of individual treatments as part of his/her *actual treatment program* (ATP).

When a guideline is followed exactly, the patient's ATP will be the same as the guideline. In practice, however, clinicians often customize a patient's ATP when they feel that changes would result in better outcomes. This results in *gaps* between the actual treatments performed on a patient and those recommended by the guideline.

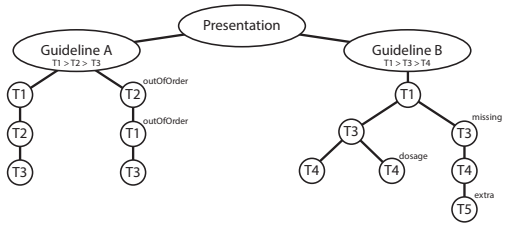

Figure 2: Treatments are first grouped by presentation and guideline. Then, a tree structure is used to aggregate statistics for all observed variations in care.

2.2 Gap Analysis

The raw input data is analyzed to first match each patient's ATP to the closest guideline. This match takes into account the patient's presentation as well as differences between the ATP and each applicable guideline for the given presentation. Then, each individual treatment in a patient's ATP is compared against the identified guideline to detect gaps.

During this comparison, some treatments are found to match exactly with the guideline's recommendations. Otherwise, a gap is recorded. Several types of gaps are identified during the analysis including extra treatments (those performed in addition to what was recommended), missing treatments (recommended treatments that were not performed), out of order treatments (administered in a different order than called for by the guideline), and various types of modified treatments (e.g., a different drug or a different dosage).

More details regarding the gap analysis algorithm itself can be found in prior work [1].

2.3 Data Aggregation and Visualization

After gap analysis, we use the GapFlow visualization to visually summarize patterns in care for groups of patients with the same presentation. The GapFlow visualization, shown in Figure 1, is a temporal visualization technique that illustrates both (1) variations in the ATPs for a population of patients and (2) gaps between ATPs and the recommended treatment programs found in the corresponding guidelines.

The first step is data aggregation. For a cohort of patients with the same presentation, GapFlow builds a tree-based data structure as shown in Figure 2. After segregating by guideline, the tree represents all observed permutations of treatments in the patients' ATPs. The nodes in the tree correspond to individual treatments labeled by gap type if a gap exists. We store the number of patients at each node to capture how frequently each treatment path is traversed.

Once the tree structure has been constructed, the next step is visualization. We employ a space-filling technique in which the X axis represents treatment progression and the Y axis represents the proportion of patients. First, the space is divided into horizontal layers with one for each guideline in the tree. The height of each layer is proportional to the number of patients assigned to the corresponding guideline.

Within each guideline layer, we plot a series of evenly spaced vertical bars representing the recommended steps in the guideline (including optional steps). The number of bars depends on the number of steps in the guideline.

We then traverse the aggregated data structure shown in Figure 2, plotting rectangles for each node in the tree. The rectangles are positioned horizontally and colored based on the gap type associated with the corresponding node. If a treatment is according to guideline, green rectangles aligned with the guideline bars are used. If a node represents an 'extra' treatment, a rectangle is plotted between guideline bars to emphasize that the treatment was above and beyond the guideline recommendation. Missing treatments are visualized as an empty rectangle aligned with the missing guideline step. Other gaps, such as out of order treatments and changes to medications are visualized as red or yellow rectangles (depending on the gap type) aligned with the corresponding guideline bars.

Figure 3: Most patients were given (a) extra treatments including all patients treated under Guideline 1. In addition, (b) the guideline recommended treatments deviated in minor ways roughly half the time regardless of which guideline was followed. Interesting, (c) clinicians universally skipped the second recommended treatment that came second in the guideline. This was an optional treatment, however, which meant a small subgroup in Guideline 2 was officially treated without any deviation.

The height of each rectangle is proportional to the number of patients. The background color between treatment indicators changes from green to red after the first deviation. Sorting to place early deviations at the bottom produces a cumulative distribution chart showing the rise in treatment programs with care gaps over time.

The result is a flow-like chart as seen in Figure 3. GapFlow charts have some similarities to previous work visualizing aggregate medical event data [2, 3]. However, prior work does not handle the overlay of guideline information which introduces several complexities. These include the need for multiple alignment points (one for each guideline event) and a visual representation of gap types (e.g., missing vs. extra vs. out of order).

3 Future Work

The work described above is preliminary. Plans for future work include (1) additional interaction capabilities to surface additional treatment information, (2) formal evaluations with clinical users to measure the technique's efficacy, and (3) applications for additional use cases to generalize our approach.

Acknowledgements

The authors thank our clinical collaborators at the Instituto Nazionale dei Tumori in Milan, Italy for their support. Their clinical expertise and data were essential prerequisites for this work.

References

[1] E. Goldbriach, N. Parush-Shear-Yashuv, A. Farkash, M. Monti, M. Torresani, R. Bertulli, P. G. Casali, B. Carmeli, and Z. Waks. Understanding deviatons from clinical practice guidelines in adult soft tissue sarcoma. In *Submitted to AMIA Annual Symposium Proceedings*. American Medical Informatics Association, 2013.

[2] K. Wongsuphasawat and D. Gotz. Exploring flow, factors, and outcomes of temporal event sequences with the outflow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2659–2668, 2012.

[3] K. Wongsuphasawat, J. Guerra Gómez, C. Plaisant, T. Wang, M. Taieb-Maimon, and B. Shneiderman. Lifeflow: visualizing an overview of event sequences. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, pages 1747–1756. ACM, 2011.