Data-Driven Exploration of Care Plans for Patients

Adam Perer
IBM T.J. Watson Research Center
Yorktown Heights, NY 10598
United States
adam.perer@us.ibm.com

David Gotz
IBM T.J. Watson Research Center
Yorktown Heights, NY 10598
United States
dgotz@us.ibm.com

CHI 2013 Extended Abstracts, April 27–May 2, 2013, Paris, France.
ACM 978-1-4503-1952-2/13/04.

Abstract
CareFlow is a novel visual analytics tool designed to help doctors devise a care plan for their patient. Using historical outcomes from clinically similar patients, CareFlow allows doctors to analyze which treatments have been effective for patients like theirs.

Author Keywords
healthcare visualization, visual analytics, care plans

ACM Classification Keywords
H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous.

General Terms
Visualization, Healthcare

Introduction
When a patient is diagnosed with a disease, their doctor will often devise a *care plan*, a sequence of medical treatments to help manage their disease or condition. When doctors devise care plans, they often must rely on their education, experience, and intuition [5]. In order to empower doctors, *care pathways* have been proposed as a promising direction in modern healthcare. Care pathways are a description of standardized, evidence-based practices to effectively manage specific diseases or conditions. However, in practice,

there is still limited access to care pathways, they are often not standardized across institutions, and may be described at too high of a granularity to directly implement. Furthermore, even if patients suffer from a common disease, they often have specific co-morbidities that can require modifications to standard, generically-defined care pathways.

The goal of our work is to leverage the rich longitudinal data found in Electronic Medical Records (EMRs) to empower clinicians with a new data-driven resource. We introduce CareFlow, a novel visual analytics tool designed to visualize the effectiveness of different potential care plans that a doctor may consider prescribing to a patient. Using the relevant clinical data of a specific patient, CareFlow mines the EMRs to find clinically similar patients using our patient similarity analytics. CareFlow then visualizes all of the different care plans that these similar patients have undergone, while providing context on which care plans were successful and which were not. The resulting visualization supports the identification of the most desirable and most problematic care plans.

Related Work

The use of electronic data has been gaining attention in recent years as a means for developing personalized evidence for making complex treatment decisions [1, 2]. The need to make sense of this complex data had led to a number of visual interfaces designed to help doctors and clinical researchers make sense of this complex data [7, 9, 3].

CareFlow is unique from the previous systems because it integrates with advanced patient similarity analytics [6] to limit the system to only display patients that are clinically similar. This provides a more contextual interface for doctors and also reduces the complexity of the visualization. Despite using a visual encoding similar to [3], Outflow was designed for symptom progression, whereas CareFlow uses a different aggregation technique to support the visualization of care plan.

Mining Care Plans from Data

CareFlow is designed to help doctors devise a personalized, data-driven care plan for specific patients. As the interface is personalized, CareFlow begins by analyzing the medical history of the target patient. Using historical medications, symptoms, diagnoses and lab results of the patient, CareFlow uses our patient similarity analytics [6] to find a population of clinically similar patients. The care plans and outcomes of this similar patient population will be made visible by CareFlow.

In order to model the care plans for the similar patient population, CareFlow mines the EMRs for relevant patient events. For each similar patient, CareFlow will extract records of performed treatments and their associated dates by querying the EMR database for relevant medical events. The result of this query is a complex dataset describing the details of various treatments given to the entire similar patient population.

Of course, each similar patient underwent treatments at different points in time. In order to unify them, CareFlow aligns all treatments by the time at which each patient was first diagnosed with the disease of interest. CareFlow defines the care plan as the sequence of treatments after diagnosis. In addition to deriving care plans, outcomes are also derived from the EMRs for each of these similar patients. By associating each care plan with an outcome, it is possible to infer which care plans lead to statistically better outcomes. CareFlow makes this outcome information visually prominent to inform medical decisions.

Visualizing Care Plans

While a doctor may be able to make sense of a care plan for a single patient (e.g. [4]), doing so for a similar patient population is much more challenging. Care plans may have a large number of different types of treatments, and the sequence of these treatments often varies as well. CareFlow was designed to summarize the temporal sequence of treatments in a comprehensive visual interface using an enhancement of the Outflow visualization technique [8].

As described in Figure 1, treatments are represented as nodes and positioned along the horizontal axis, which represents the sequence of treatments over time. The diagnosis of a disease occurs on the far left of the visualization, and treatments in the care plan extend to the right. For instance, in this illustrative example, patients diagnosed with a specific disease were then given either *Diuretics* or *Cardiotonics*. The height of each node is proportional to the number of patients that took a given treatment. So, in this example, twice as many patients were prescribed *Diuretics* than treatment *Cardiotonics*.

Each treatment node is also augmented with a time edge, whose width encodes the duration of a how long it took patients to transition to this treatment. In this example, patients were typically prescribed *Cardiotonics* more quickly than *Diuretics* after diagnosis. Link edges are also present to connect nodes from their previous and future nodes in the care plan.

The visual elements are colored according to the average outcome of all patients represented by the node or edge. Elements that are colored green represented parts of the care plan where patients remained healthy, whereas elements that are colored red indicate care plans of patients who ended up in poor health.

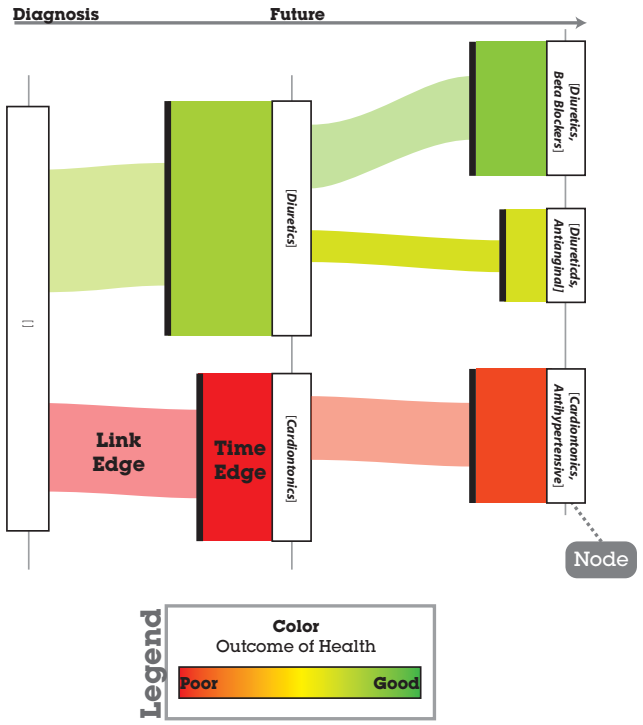

Figure 1: The visual encoding of care plans in CareFlow. The height of nodes represents the number of patients. The width of time edges represents the average duration of treatments. Color represents the average patient outcome.

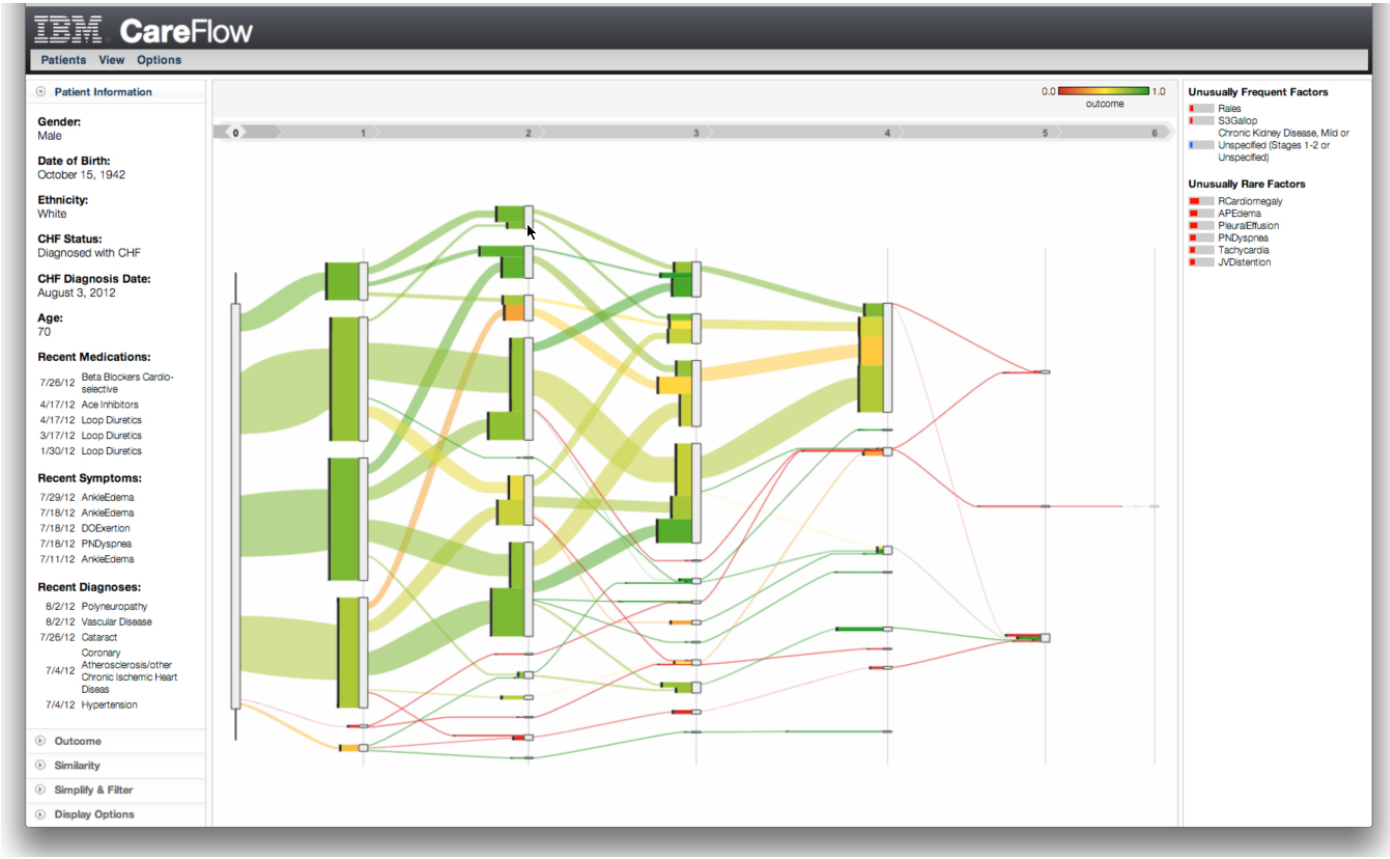

Figure 2: CareFlow's visual interface. The left panel displays a summary of the patient's relevant medical history. The center panel displays a visualization of the care plans of the 300 most similar patients. The right panel displays the factors associated with a selected subset of patients.

For example, for patients who initially took *Diuretics*, the next treatment that leads to the best outcome for patients is *Beta Blockers*, whereas the patients who are treated with *Antianginal Agents* end up with worse outcomes.

Scenario: Congestive Heart Failure

This scenario involves a doctor who has recently diagnosed a patient with congestive heart failure and wishes to use CareFlow to examine the historical outcomes of possible care plans. CareFlow connects to a longitudinal EMR database of over 50,000 patients with heart conditions spanning over 8 years.

On the left-hand side of Figure 2, a summary of the patient's relevant medical history is shown, including recent medications, symptoms, and diagnoses. In the center panel of Figure 2, a visualization of the care plans of the 300 most similar patients is shown. The left-most node represents these similar patients at their point of diagnosis with heart failure. As the visualization extends to the right, the various treatment sequences of similar patients are shown. The care plans are colored according to a continuous color scale, as indicated by the legend in the upper-right. By default, plans that are colored red implies most patients within that node ended up being hospitalized, whereas green plans means most patients managed to stay out the hospital. However, CareFlow empowers doctors to interactively customize the outcome measure. For example, it is possible for the doctor to use CareFlow to focus on care plans that lead to patient deaths, to better understand treatments associated with mortality.

In addition to gaining an overview of all care plans, a doctor can also focus on the most successful treatment plan. By selecting the appropriate button, the care plan that leads to the best outcomes for patients is highlighted, as shown

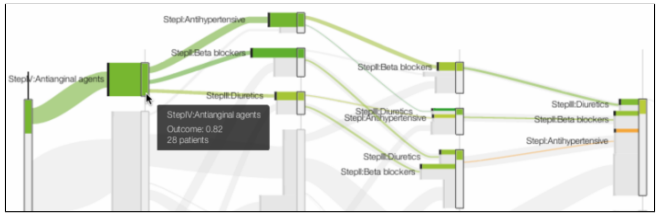

Figure 3: Doctors can choose to automatically highlight the care plan that leads to the best outcomes for patients.

in Figure 3. This particular care plan is intriguing to the doctor, as it suggests that similar patients have the best outcome when starting with a Stage 4 treatment (Antianginal Agents) , instead of a Stage 1 treatment (Antihypertensive Drugs) that the doctor's institutional clinical guidelines suggest.

CareFlow provides doctors with the ability to get more information about the patients who undertook a particular care plan. By selecting a Treatment node, doctors can view a precise count of the number of patients the node represents, as well as the average outcome for these patients. In addition, the right panel of the interface displays summary information about a set of patients by displaying factors common to this cohort, as well as factors rare in this group. For instance, the selected patients in Figure 2 typically exhibit symptoms of Rales and have more diagnoses of Chronic Kidney disease. On the other hand, these patients exhibit *Acute Pulmonary Adema* and *Pleural Effusion* less frequently than the rest of the similar patients.

Future Work and Conclusion

In order to validate the effectiveness of CareFlow, we aim to continue to deploy and evaluate the tool with doctors in clinical settings. In order to address some of the tool's

limitations, we plan to enhance the capabilities of the visualization and analytics with novel methods to handle large numbers of treatment types as well as the occurrence of concurrent treatments.

To conclude, CareFlow is a novel visual analytics tool designed to assist doctors when researching care plans for their patients. CareFlow provides an overview of the care plans for a cohort of similar patients to help identify the most desirable and most problematic plans.

Acknowledgements

We thank Jianying Hu, Robert Sorrentino, Harry Stravopolous, and Jimeng Sun for their assistance and feedback in designing the CareFlow system.

References

[1] Ebadollahi, S., Sun, J., Gotz, D., Hu, J., Sow, D., and Neti, C. Predicting patients trajectory of physiological data using temporal trends in similar patients: A system for near-term prognostics. In *American Medical Informatics Association Annual Symposium (AMIA)*, AMIA 2010 (2010).

[2] Frankovich, J., Longhurst, C. A., and Sutherland, S. M. Evidence-based medicine in the emr era. *New England Journal of Medicine* (Nov. 2011), 1758–1759.

[3] Gotz, D., and Wongsuphasawat, K. Interactive intervention analysis. In *American Medical Informatics Association Annual Symposium (AMIA)*, AMIA 2012 (2012).

[4] Plaisant, C., Mushlin, R., Snyder, A., Li, J., Heller, D., and Shneiderman, B. Lifelines: using visualization to enhance navigation and analysis of patient records. In *American Medical Informatics Association Annual Symposium (AMIA)*, AMIA 1998 (1998), 7680.

[5] Tracy, C. S., Dantas, G., and Upshur, R. Evidence-based medicine in primary care: qualitative study of family physicians. *BMC Family Practice 4*, 1 (2003), 6.

[6] Wang, F., Sun, J., and Ebadollahi, S. Integrating distance metrics learned from multiple experts and its application in inter-patient similarity assessment. In *SDM*, SIAM / Omnipress (2011), 59–70.

[7] Wang, T. D., Plaisant, C., Quinn, A. J., Stanchak, R., Murphy, S., and Shneiderman, B. Aligning temporal data by sentinel events: discovering patterns in electronic health records. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, ACM (New York, NY, USA, 2008), 457–466.

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

[9] Wongsuphasawat, K., Guerra Gómez, J. A., Plaisant, C., Wang, T. D., Taieb-Maimon, M., and Shneiderman, B. Lifeflow: visualizing an overview of event sequences. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, ACM (New York, NY, USA, 2011), 1747–1756.