Data-Driven Event Sequence Visualization for Rectal Cancer Outcomes

Taxonomizing the Role of Mobile Applications in Sociotechnical Feedback Loops

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By
Andrew Balch

November 8, 2024

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

ADVISORS

Rider Foley, Department of Engineering and Society

Afsaneh Doryab, Department of Computer Science

Introduction

Colorectal cancer is one of the most common forms of cancer, and was responsible for 935,173 deaths worldwide in 2020 (PDQ® Adult Treatment Editorial Board, 2024). As with most cancers, there are a variety of factors that impact a patient's course of treatment and chances of recovery. The primary treatment for rectal cancer is surgical resection of the tumor, but the surgical approach, intent, and whether the surgery is preceded and/or followed by chemotherapy depends on the stage of cancer and individual patient characteristics (PDQ® Adult Treatment Editorial Board, 2024). Therefore, it is crucial that clinicians analyze as much about the patient and their potential treatment paths as possible to maximize treatment effectiveness and minimize negative outcomes. This has motivated 6 cancer hospitals to share data and create the US Rectal Cancer Consortium (RCC), a dataset compiled from over a decade of rectal cancer patient electronic health record (EHR) data (Ali et al., 2024).

The RCC includes 408 unique data features that span the treatment path of 1881 patients, broken up into 13 different observational and treatment stages. A series of papers has been published that conduct statistical analyses on subsets of features in this dataset (e.g. Ali et al., 2024). While these papers are successful in exploring how different subsets of features impact one or more outcomes of interest, their scope is limited by the reliance on statistical multivariate analyses. There is a need to assess all 408 features in this dataset and discover where important features may have been overlooked, how treatment decisions impact subsequent treatments as well as outcomes over time, and convey these insights to clinicians in an intuitive way. Machine learning (ML) techniques and interactive data visualizations are uniquely suited to address this gap and yield more complex, comprehensive, and interpretable insights. Within the problem area of colorectal cancer treatment, my technical topic focuses on a data-driven event sequence visualization for the incidence of an anastomotic leak (AL) in rectal cancer surgery, which is a critical, yet overtreated, complication (Shao et al., 2023).

In building a technological system that guides and informs human decisions, I must recognize that I am engaged in an action that is fundamentally political (Green, 2021). My tool encodes normative understandings of the world and applies them to individuals (patients and clinicians). This is an example of how algorithmic systems are actors, which shape and are shaped by social processes (Joyce et al., 2021). An interplay between society and algorithmic systems is evidenced by the social harms, biases, and injustices that are perpetuated by these systems (Burrell & Fourcade, 2021; Eubanks, 2018). Many approaches to addressing these issues have been proposed, with the most promising being algorithmic justice initiatives which seek to understand and ameliorate oppression in its varied yet intersecting forms (Chordia et al., 2024; Dombrowski et al., 2016). Theories, debates, and analyses within algorithmic justice often focus most on high-risk, high-profile cases (e.g. recidivism prediction). In the second half of this prospectus, I will apply actor-network theory (ANT) (Latour & others, 1992) to analyze the role of algorithms as social mediators, and argue for a deeper understanding of how algorithms influence our day-to-day lived experience, and how they may do so (un)justly.

Technical Topic

Clinical decision support systems (CDSS) are increasingly using artificial intelligence (AI) and ML to determine optimal treatment decisions based on clinical data, but are still limited by the hesitancy of clinicians to trust on them (Sivaraman et al., 2023). Prior work has also explored the use of ML classifiers to predict the incidence of AL, but have done so with far less detailed datasets and have not taken meaningful steps to provide utility to the clinician beyond a binary prediction (Shao et al., 2023). In my technical project, I focus on informing the clinician of critical treatment associations and outcomes in an intuitive manner, allowing them to use their own judgment to come to a final decision. I aim to develop a visualization tool for treatment event sequences, where clinicians can optionally leverage ML to determine the individual clinical variables during each stage of rectal cancer treatment that are most relevant to an outcome. This approach is especially useful in scenarios where clinicians may not want to rely on a binary prediction, or where the treatment sequence is too large to feasibly visualize every feature.

An event sequence is a time-ordered series of discrete events, and is a popular way to model the treatment of a disease over time (Guo et al., 2022). Timelines, hierarchies (e.g. trees), and Sankey diagrams are commonly used to visualize event sequences. In clinical use-cases, these charts are useful to compare different patient cohorts, visualize disease outcomes, and analyze treatment prognoses. Clinical event sequence visualizations incorporate complex dashboards to facilitate data selection, event pattern/rule mining, and/or clustering of patients into cohorts. However, the sophistication of these tools becomes their limitation, with studies noting a steep learning curve for clinical stakeholders, who prefer a more straightforward, at-a-glance style of communicating patient information (Guo et al., 2022). Some works have had success with a more subtle approach, where AI tools enhance a visualization instead of being the focus of it (Yang et al., 2019).

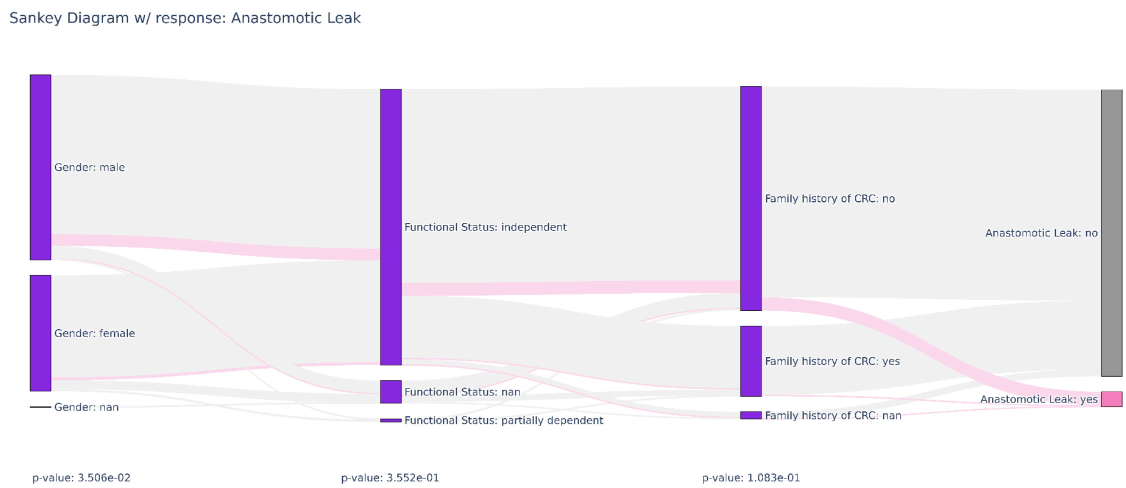

Figure 1: An example Sankey diagram-based visualization of RCC patient data. The visualization demonstrates how patient demographic features (purple) are connected to the incidence of AL (gray and pink). The relative width of each vertical bar and connecting line illustrates the frequency of those values within the data.

Therefore, my technical project will focus on simplicity and exploration. Working closely with an expert stakeholder in rectal cancer surgery, I will build a Sankey diagram-based visualization that represents each data point in the RCC as a node in a flowchart that culminates at a target outcome (AL) (Figure 1). Clinicians will be able to specify a set of treatment features they would like to investigate in the context of AL, and a Sankey diagram will be built from the RCC patient data. Different feature selection approaches will be experimented with to assess their usefulness in finding treatment variables that are strongly associated with AL, without requiring a surgeon to visualize all 408 features. Feature selection works by iterating over each feature in a dataset and selecting it only if it improves the performance of a predictive ML model (scikit-learn developers, 2024). Since the goal of the visualization is to present the user with the most important variables at each event in the sequence, an ideal feature selection method should evenly represent variables from each event without being biased towards events at the beginning or end of the chain. To improve the robustness of the feature selection techniques, I intend to take an ensemble approach where multiple predictive models will vote to determine which of the selected features will be visualized in the diagram.

Actor-Network Theory as a Means to Understand Algorithmic Injustice

This act of applying an artifact to inform and guide human action can be analyzed using the STS framework of ANT provided by Latour (1992). For Latour, five main aspects are key to understanding sociotechnical systems: programs of action, delegation, prescription, discrimination, and mediation. A technology's program of action is the set of values that are ascribed to it and that define its intended functionality. When a technology enters society, users apply it to do work for them (delegation). In doing so, the user can only act in a manner defined by the technology itself (prescription), and is therefore shaped by the technology (discrimination). Altogether, ANT describes the mediating push and pull between a technology and society.

In the context of my technical topic on visualizing how cancer treatments lead to oncologic outcomes, I inscribe into this technology a program of action that aligns with a value for quantifiable clinical outcomes. The task of determining which treatment decisions are most associated with a target outcome is delegated to my technology, away from the clinician. In doing so, the technology prescribes information to the human doctor that will influence their behavior when developing a treatment plan. By prescribing only some information to the clinician, my tool discriminates against all other treatment factors that are not identified as "relevant" to the given outcome. By recognizing how my technology prescribes, delegates, and discriminates, it is clear that it assumes a mediating role in the task of cancer treatment. When a clinician is considering how a course of treatment may influence a patient's outcome, my tool will influence the clinician's final decision.

My technical project is not the only algorithmic system that assumes such a mediating role. All automated processes which are enabled to act within society in accordance with values imbued by their design are mediators which prescribe, delegate, and discriminate. What has become increasingly prevalent with the rise of artificial intelligence is algorithms that are unjust mediators (Eubanks, 2018). The root issue of these unjust algorithmic systems can be seen as a power asymmetry, where economic power has more control over how algorithms are researched, developed, and deployed as actors in society than social power (Balch, 2024). Thus, I am concerned with when and why algorithms as mediators perpetuate social injustices and, most importantly, what engineers can do to combat this increasingly prevalent issue.

Several research areas have surfaced to combat issues of algorithmic injustice. Within the technological sphere, this is primarily characterized by fairness, responsibility, and trustworthiness (Kaur et al., 2022). These approaches have a limited effectiveness in confronting the root cause of injustice, because they center around the algorithm itself instead of considering the surrounding sociotechnical factors (Polack, 2020). Such a tendency is evident in recent responsible large language model (LLM) prototyping tools, where the responsibility of avoiding social harms is pushed onto the individual user rather than the owner of the AI model itself (tech companies) (Wang et al., 2024). By contrast, algorithmic justice initiatives provide a diverse set of tools that take social structures and systems of inequality into account (Chordia et al., 2024; Davis et al., 2021; Dombrowski et al., 2016). To apply algorithmic justice-based approaches, it is necessary to first deconstruct and analyze the mechanisms by which injustices may appear within the sociotechnical world.

Research Question and Methods

Injustices have been illuminated in high-profile cases like recidivism prediction (determining whether incarcerated persons will commit another offense), but smaller-scale, day-to-day influences remain under the radar and, therefore, unaddressed. With this gap in mind, I ask the following question: How do mobile applications leverage our data to invisibly influence the ways users see and move through the world in their daily lives?

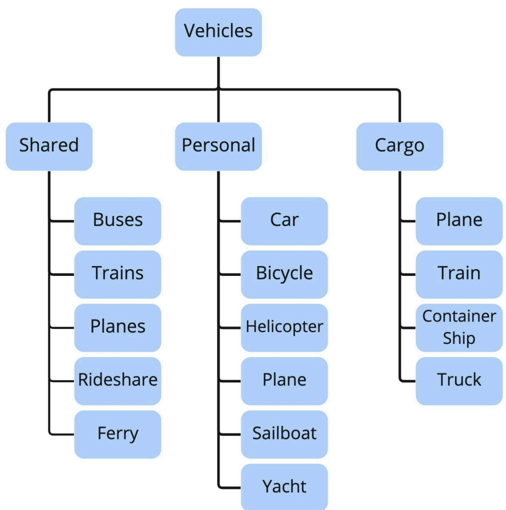

Figure 2: An example taxonomy for vehicles. Structuring information in this manner helps to describe each vehicle and classify them by their characteristics. Figure source: (Laubheimer, 2022)

I will answer this question by studying the popular applications on the Google Play Store to build a taxonomy (Figure 2) of how each app collects data and uses it to influence user interactions. App listings on Google Play contain a wealth of relevant data to support this analysis. First, developers self-report what data the app collects from its users. Second, app descriptions provide insights into the app's key features and how they function. Third, user reviews can further reveal the app's functionality, how users interact with the app, and systematic issues therein. After analyzing all three sources of evidence, I will characterize the program of action of each app, as well as how it prescribes, delegates, and discriminates during interactions with the user. From this, it will be possible to identify apps that share similar characteristics and organize them into a descriptive taxonomy of algorithmic influence. Last, I will theorize about how algorithmic justice can be applied to alleviate possible avenues for injustice as revealed by the taxonomy, and ultimately empower us as users of algorithmic systems.

Conclusion

For AI and ML tools to help clinicians, they must not aim to make decisions for the clinician. Rather, AI and ML tools are best utilized as a means to consider hundreds of variables about how patient demographics and treatment decisions may influence clinical outcomes and negative side-effects. My technical topic aims to design a visualization tool that can communicate these influences to a clinician, and utilize ML to determine the most important treatment variables.

My STS topic is formed around a recognition that algorithmic systems influence and are influenced by a variety of complex social factors. This feedback loop imparts a program of action into the algorithmic system as an artifact, which can (and does) lead to the reapplication of social injustices. In my research, I will analyze and categorize the role played by mobile applications as actors in a sociotechnical system. A well-rounded understanding of these opaque algorithmic influences will undoubtedly inform future work in algorithmic justice by pinpointing when and why harms may arise, as well as what design changes could be made to combat them and empower users.

References

Ali, D., Syed, M., Gamboa, A. C., Hawkins, A. T., Regenbogen, S. E., Holder-Murray, J., Silviera, M., Ejaz, A., Balch, G. C., & Khan, A. (2024). Association of omental pedicled flap with anastomotic leak following low anterior resection for rectal cancer. *Journal of Surgical Oncology*. https://doi.org/10.1002/jso.27572

Balch, A. (2024). *Why Algorithms Remain Unjust: Power Structures Surrounding Algorithmic Activity* (arXiv:2405.18461). arXiv. https://doi.org/10.48550/arXiv.2405.18461

Burrell, J., & Fourcade, M. (2021). The Society of Algorithms. *Annual Review of Sociology*, *47*(Volume 47, 2021), 213–237. https://doi.org/10.1146/annurev-soc-090820-020800

Chordia, I., Baltaxe-Admony, L. B., Boone, A., Sheehan, A., Dombrowski, L., Le Dantec, C. A., Ringland, K. E., & Smith, A. D. R. (2024). Social Justice in HCI: A Systematic Literature Review. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–33. https://doi.org/10.1145/3613904.3642704

Davis, J. L., Williams, A., & Yang, M. W. (2021). Algorithmic reparation. *Big Data & Society*, *8*(2). https://doi.org/10.1177/20539517211044808

Dombrowski, L., Harmon, E., & Fox, S. (2016). Social Justice-Oriented Interaction Design: Outlining Key Design Strategies and Commitments. *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, 656–671. https://doi.org/10.1145/2901790.2901861

Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.

Green, B. (2021). Data science as political action: Grounding data science in a politics of justice. *Journal of Social Computing*, *2*(3), 249–265.

Guo, Y., Guo, S., Jin, Z., Kaul, S., Gotz, D., & Cao, N. (2022). Survey on Visual Analysis of Event Sequence Data. *IEEE Transactions on Visualization and Computer Graphics*, *28*(12), 5091–5112. https://doi.org/10.1109/TVCG.2021.3100413

Joyce, K., Smith-Doerr, L., Alegria, S., Bell, S., Cruz, T., Hoffman, S. G., Noble, S. U., & Shestakofsky, B. (2021). Toward a Sociology of Artificial Intelligence: A Call for Research on Inequalities and Structural Change. *Socius*, *7*. https://doi.org/10.1177/2378023121999581

Kaur, D., Uslu, S., Rittichier, K. J., & Durresi, A. (2022). Trustworthy Artificial Intelligence: A Review. *ACM Computing Surveys*, *55*(2), 39:1-39:38. https://doi.org/10.1145/3491209

Latour, B. & others. (1992). Where are the missing masses? The sociology of a few mundane artifacts. *Shaping Technology/Building Society: Studies in Sociotechnical Change*, *1*, 225–258.

Laubheimer, P. (2022, July 3). *Taxonomy 101: Definition, Best Practices, and How It Complements Other IA Work*. Nielsen Norman Group. https://www.nngroup.com/articles/taxonomy-101/

PDQ® Adult Treatment Editorial Board. (2024, September 6). *PDQ Rectal Cancer Treatment*. National Cancer Institute. https://www.cancer.gov/types/colorectal/hp/rectal-treatment-pdq

Polack, P. (2020). Beyond algorithmic reformism: Forward engineering the designs of algorithmic systems. *Big Data & Society*, *7*(1), 2053951720913064. https://doi.org/10.1177/2053951720913064

scikit-learn developers. (2024). *SequentialFeatureSelector*. Scikit-Learn. https://scikit-learn.org/1.5/modules/generated/sklearn.feature_selection.SequentialFeature

Selector.html

Shao, S., Zhao, Y., Lu, Q., Liu, L., Mu, L., & Qin, J. (2023). Artificial intelligence assists surgeons' decision-making of temporary ileostomy in patients with rectal cancer who have received anterior resection. *European Journal of Surgical Oncology*, *49*(2), 433–439. https://doi.org/10.1016/j.ejso.2022.09.020

Sivaraman, V., Bukowski, L. A., Levin, J., Kahn, J. M., & Perer, A. (2023). Ignore, Trust, or Negotiate: Understanding Clinician Acceptance of AI-Based Treatment Recommendations in Health Care. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–18. https://doi.org/10.1145/3544548.3581075

Wang, Z. J., Kulkarni, C., Wilcox, L., Terry, M., & Madaio, M. (2024). Farsight: Fostering Responsible AI Awareness During AI Application Prototyping. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–40. https://doi.org/10.1145/3613904.3642335

Yang, Q., Steinfeld, A., & Zimmerman, J. (2019). Unremarkable AI: Fitting Intelligent Decision Support into Critical, Clinical Decision-Making Processes. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–11. https://doi.org/10.1145/3290605.3300468