

Site Reliability Engineering: Improving Performance Transparency in a Trading Platform

(Technical Paper)

The Effect of Data-Based Economic Metrics on Marginalized Identity Groups

(STS Paper)

A Thesis Prospectus Submitted to the
Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia
In Partial Fulfillment of the Requirements of the Degree
Bachelor of Science, School of Engineering

Khushi Chawla

Fall, 2022

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

Catherine D. Baritaud, Department of Engineering and Society

Briana B. Morrison, Department of Computer Science

As our world moves toward increasingly running on data, the importance, and dangers of metrics increase. Metrics can benefit groups by condensing data in a more readable and digestible way or provide transparency to inform future decisions. Key performance indicators, or KPIs, drive business decisions as they represent how close the company is to reaching performance goals (Lynch, 2015). Metrics can allow managers and executives to find inefficiencies in output, keep employees focused on executable goals and help focus attention (Lynch, 2015). However, like all technologies, metrics also have their drawbacks; metrics based on skewed data can improperly reflect the composition of datapoints (Atler et. al., 2016). People's biases are reflected in the data used to create Machine Learning (ML) models and Artificial Intelligence (AI) algorithms, which when used to make decisions in the real world can unfairly target certain identity groups (Coté et al., 2021).

The tightly coupled technical and STS papers center around the use and outcome of metrics based on financial data. The technical report focuses on a beneficial use of metrics, analyzing how using they can be used to increase speed and accuracy of a financial data platform to increase revenue and client productivity. The product of the technical project will be beneficial in maturing a new platform by adding a way to easily track performance in a consistent manner at scale. Without this, it is difficult to justify migrating from the old system to the new platform, despite the non-quantifiable benefits of the data platform. The STS report analyzes a potential harmful effect of metrics, examining how using algorithms to analyze people's credit score and loan eligibility to speed up bank's decision process can marginalize identity groups. The results found from the STS research can point out one of the drawbacks of using metrics without vetting the process with which they were calculated. It is important to analyze how biases in society can be displayed in data, which is used to create metrics that

inform decisions that impact society, creating a self-perpetuating cycle of marginalization.

The work for the projects began in June 2022 and will continue through May 2023, as seen in Figure 1. The technical project was completed in the Summer of 2022 at a summer internship at a financial software company under the guidance of a mentor and manager (specifics cannot be divulged due to a non-disclosure agreement). The technical report will be written in Fall 2022 under the technical supervision of Brianna Morrison, UVA Computer Science associate professor. The STS project will begin in Fall 2022 and continue through Spring 2023 under the supervision of Catherine Baritaud and Travis Elliot, professors in the department of Science, Technology and Society.

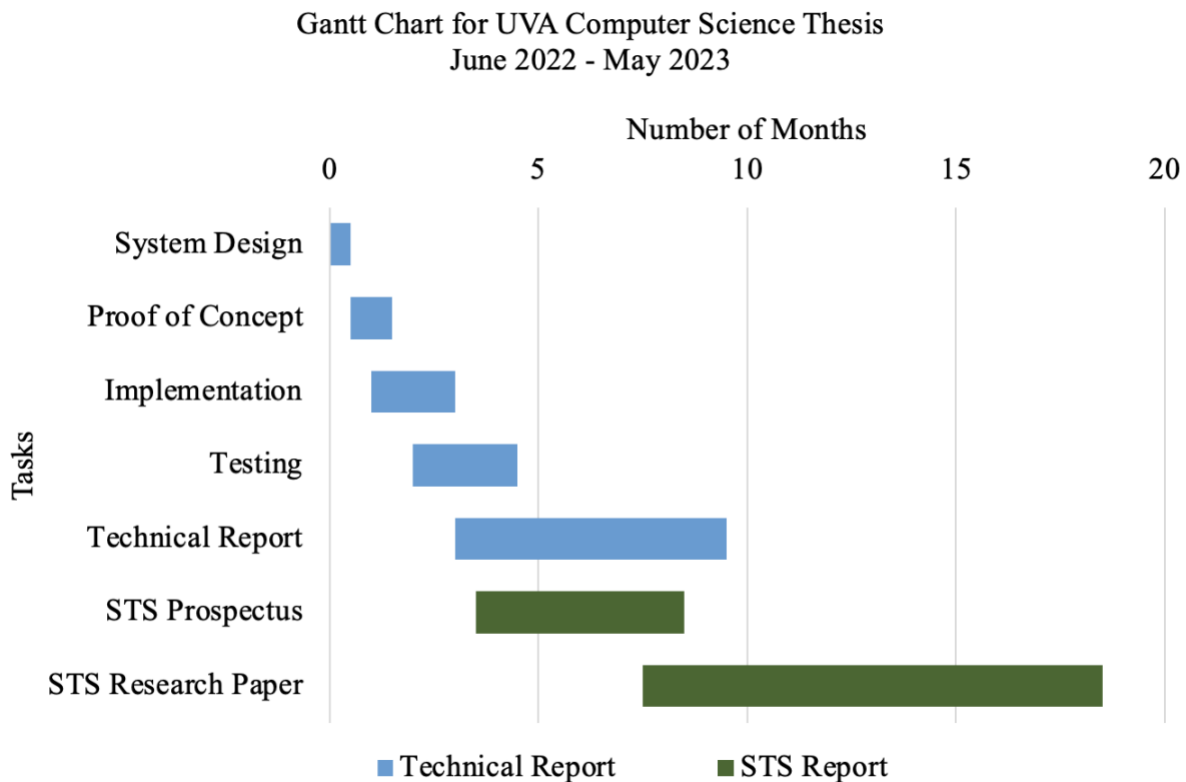


Figure 1: Gantt chart UVA Computer Science Thesis. This figure visualizes the expected timeline for the major milestones achieved for the technical project and STS project (Chawla, 2022).

IMPROVING PERFORMANCE TRANSPARENCY IN A TRADING PLATFORM

A global New York-based financial software company's trading system's data platform lacked the transparency necessary to inform decisions on resource allocation and outage response. The platform is comprised of multiple components, called domains or services, which supports data storage functionality. Other teams within the company rely on this platform to store their data in a standardized way to reduce duplication, mitigate mismatched data between teams, add division of concerns, and increase speed and efficiency. Because the platform is relatively new, there is no standardized and scalable way to track and publish performance metrics, which are crucial in pointing out inefficiencies with the platform or the teams using it, patterns in the platform's usage, and resource insufficiencies (Watts, 2022). Platforms like this one support international clients trading at high volumes at all times of day. Having information on the platform's performance can inform decisions that prevent downtime, can prevent the loss of millions of dollars an hour to client companies (Hersh, n.d.). The process taken to create a solution to this problem, seen in Figure 2 on page 4. First, determining what could feasibly be accomplished in the span of 2.5 months and what requirements must be met to move on to the next step. The next steps involve researching both open source and internally supported frameworks, architectures, and coding languages to design an architecture to be implemented. Once implementation is complete, the last stage involves incorporating the product into a testing environment to verify that the product is functional and iterate on improvements. Once improvements were established, starting from the third step was revisited to incorporate the changes into the product.




Figure 2. Technical Solution Method. This figure outlines the different steps taken to create a solution to the technical problem (Chawla, 2022).

According to Red Hat (2020), Site Reliability Engineering (SRE) helps system administrators and developers manage large systems through code. Implementing such a framework could allow for more transparency in the growing data platform. Incorporating Site Reliability Engineering (SRE) guidelines, which state that services should perform above an established threshold (Hall, 2021), requires publishing metrics to evaluate how well the domains are achieving goals. These thresholds are called Service Level Objectives, or SLOs and are associated with performance goals specified by the client, called Service Level Agreements, or SLAs (Beyer et al., 2016).

According to Saguier (2022), a variety of classes of objectives fall under the umbrella of SRE, the biggest three being speed, load, and accuracy. Ensuring that the associated Service Level Metrics, or SLMs, of latency, throughput and error rate of requests are performing better than predetermined thresholds is vital in keeping a system running efficiently. To achieve this, a pipeline to publish latency, throughput, and error rate metrics to a telemetry management system would be necessary.

Discussions regarding the architecture design were conducted internally and included consultations from other teams for what internally supported software/frameworks should be used for specific components, however more information can't be shared due to a non-disclosure agreement. In accordance with the SRE guidelines found and discussions regarding architecture, a system was designed to solve the problem, as seen in Figure 3 on page 5. To ensure that the platform was running at high speeds and with accuracy, aggregated metrics from each domain in

the system where our engineering users manage their domains were logged. Metrics were published to an internal management system, upon which a dashboard was created that queried the storage and displayed request quantity, type, time, and frequency that was incorporated in a client-facing web portal.




Figure 3. Technical System Design. This figure outlines the different components in the technical solution system design (Chawla, 2022).

The system is still in the process of being implemented and used, but some of the more immediate results already showed promise. Consistent spikes for response time for one of the domains were seen from the middle of the night through the time-series graphs on the dashboard. This indicated that it was likely being heavily used internationally, as countries' markets are open at different times, and resources would likely need to be dynamically allocated for that timeframe. Additionally, a newer domain was confidently rolled out to the public with the performance information about internal use through the metric monitoring system. This allowed performance promises to be made to external clients. Some anticipated outcomes include seeing a decrease in downtime and shorter outages. We also hope to see more teams adopt the platform, as we can now concretely that it can satisfy their performance goals.

The technical paper will be in a state-of-the-art format, discussing the project's related work, the design and development process and results.

THE EFFECT OF ECONOMIC METRICS ON MARGINALIZED IDENTITY GROUPS

Machine learning and data algorithms allow companies to reduce processing bottlenecks, make workers more efficient and increase revenue and profit, making them advantageous from a

business operations perspective (Atler et al., 2016, Predictive Maintenance section). One such application of this is in the banking industry: banks have started using machine learning models to generate metrics which can determine who should and should not qualify for certain loans and what percent interest to allocate for an approved loan based on their credit history and financial data. These algorithms allow bankers to make faster decisions, reduce manual labor, add consistency, reduce human error, and dynamically improve and update (Dakin, 2022, Credit Decisions section).

However, with the adoption of models to achieve faster and more accurate decisions may come the tradeoff of fair decisions. Machine learning models bias arises when the result of a model is systematically prejudiced due to assumptions made in the process (Pratt, 2020). In a study by Coté et al., biases can appear in the decision-making process when creating and deploying the model and in the data set or output (2021, p. 74). Their research paper argues that an algorithm's bias is discriminatory if the differences between its result and the expected outcome given the context in which its employed are due to protected attributes, which are encoded in legislation and include attributes like sex, gender, and ethnicity, or proxy attributes, which are closely correlated with protected attributes (Coté et al., 2021, p. 75). When such biases are present in a model deployed in a real world, it can perpetuate systematic racism, sexism, homophobia and more. Currently, little research has been done on the effect of the use of algorithms and metrics on different identity groups in making loan decisions. Using the handoff model seen in Figure 4 on page 7, we can determine the different groups involved in the creation and use of the Machine Learning model and how bias can be propagated through the model from each group's actions. Understanding how the end-to-end model development process flows using this model, which may vary from company to company or even model to model, will provide a

foundation for how bias can be prevented. Additionally, using the handoff model to create an Actor Network Theory map could provide insight into disconnects between different groups and information flow and transparency can be improved to reduce biased models from being used.




Figure 4. Loan Metrics Handoff Model. This figure outlines the different groups involved in the development and deployment in a loan approval ML model (Chawla, 2022).

Though banks are prohibited from considering protected attributes like race and ethnicity in making decisions, the use of proxy attributes like shorter credit histories, poorer neighborhoods and lower income jobs associated with minority communities in training models means biases can arise in likelihood of rejecting minority loan applications (Andrews, 2021, Root of the Problem section). This in turn can aggravate credit inequality and further the divide, causing the models to become even more biased and further marginalize traditionally lower income identity groups (Andrews, 2021, Misallocation of Credit section). The Actor Theory Network (ANT) map, seen in Figure 5 located on page 8, can be used to analyze how disconnects in information and collaboration can bring rise to bias and which identity groups and actors can be affected by results of a biased algorithm (Callon, 1984). Data scientists, model designers and data collectors all work together to create ML algorithms, but their workings are a black box to the other actors in the network. Banks use the end-product algorithm without much knowledge about the creation and testing processes and regulatory bodies are even further removed, only primarily interacting with banks directly. The minority groups impacted by these models aren't consulted by any of the discussed actors in this network yet are the receivers and users of these biased models. To mitigate the propagation of bias of the pipeline established in

the handoff model, the existing interactions outlined in blue in the figure should be supplemented with those in red. In involving the minority end users in the model creation process, the development of biased models can be prevented. Increasing information flow between banks and regulatory bodies, and the model creation process and the end users could prevent the use of already biased models and ensure that decisions made by algorithms don't disproportionately impact minority groups. By using ANT to understand the influence regulatory bodies have on the creation of these models, how banks business goals are achieved through these models and how the models themselves impact certain identity groups, a more structured and transparent approach to creating models with minimal bias can be developed and implemented.

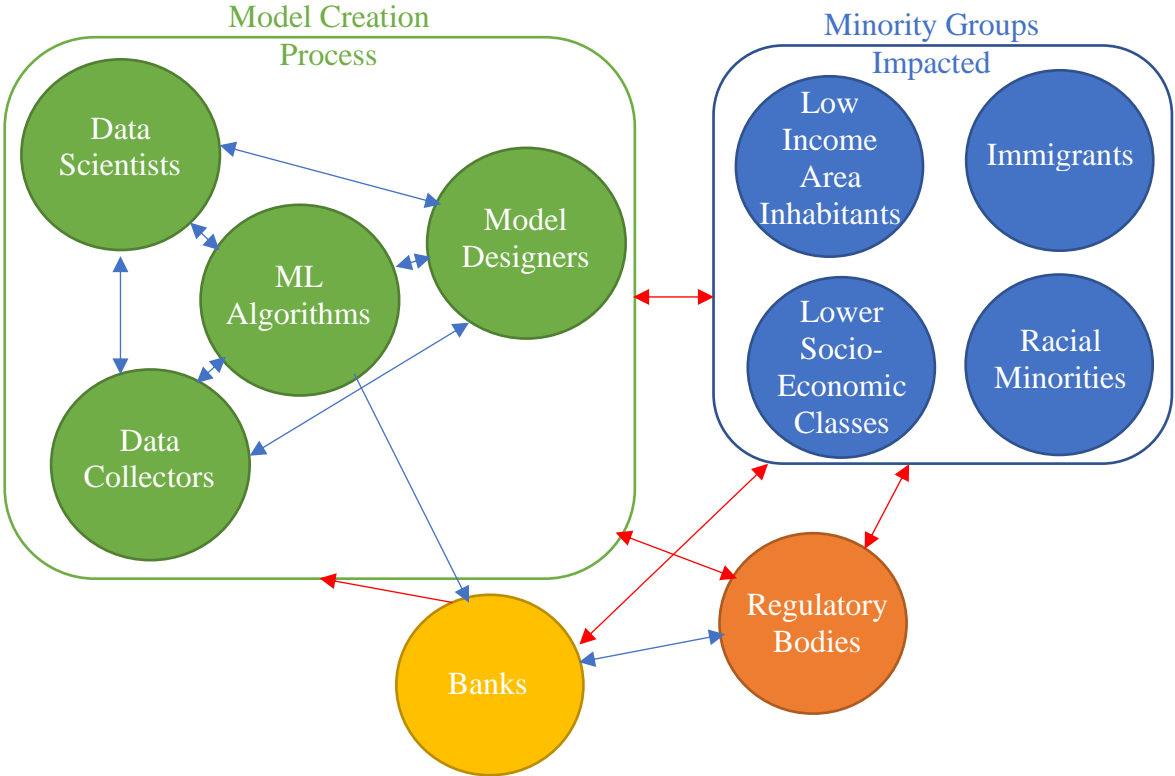


Figure 5. Loan Metrics Actor Network. This figure outlines the different groups affected by a loan approval ML model and interactions between them (Chawla, 2022).

The research paper will be in the form of a scholarly article, outlining the relationships between groups involved in creating loan approval ML models and metrics and its outcome and

the relationship between those involved in creation and those who are affected by the model. By determining whether the outcomes of the loan approval models result in differences in outcomes due to protected or proxy attributes using the groups in the ANT map in Figure 5 on page 8, we can establish the bias in ML models. Then, the handoff model outlined in Figure 4 on page 7 can be used to determine where these biases arose and be used to mitigate such biases in the future.

THE POSITIVE AND NEGATIVE EFFECTS OF ECONOMIC METRICS

Metrics can be beneficial or harmful depending on the data they're based on and the context in which they're used, which the tightly coupled technical and STS reports focus on. The technical project outlines a valuable use of metrics to improve speed and accuracy for a data platform that deals with large volumes of financial trades. The STS paper centers around the use of similar metrics, specifically the potentially harmful effects of using peoples' financial data to inform bank decisions like loans and interest rate. While both the technical and STS papers deal with metrics in the financial world, the impact of the metrics are very different as the metrics from the technical project can increase revenue for a financial software company while those from the STS project can be used to deny loan requests from those in marginalized identity groups.

REFERENCES

- Andrews, L. E. (2021, August 6). How flawed data aggravates inequality in credit. *Stanford University Human-Centered Artificial Intelligence*. <https://hai.stanford.edu/news/how-flawed-data-aggravates-inequality-credit>
- Atler, A., Sachdev, S., Wilson, H. J. (2016, May 3). How companies are using machine learning to get faster and more efficient. *Harvard Business Review*. <https://hbr.org/2016/05/how-companies-are-using-machine-learning-to-get-faster-and-more-efficient>
- Baritaud, C., Carlson, B. (2009). STS Handoff Model. *Class handout* (Unpublished). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.
- Beyer, B., Jones, C., Murphy, N. R., & Petoff, J. (2016). Site reliability engineering. *Site Reliability Engineering*. O'Reilly media, inc. <https://sre.google/sre-book/table-of-contents/>
- Bijker, W. E., & Pinch, T. J. (1984). The social construction of facts and artifacts. *Social Studies of Science*, 14, 399–441. <https://doi.org/10.1177/030631284014003004>
- Callon, M. (1984). Some elements of a sociology of translation: Domestication of the scallops and the fishermen of St Brieuc Bay. *The Sociological Review*, 32, 196-233. <https://doi.org/10.1111/j.1467-954X.1984.tb00113.x>
- Chawla, K. (2022). *Gantt Chart UVA Computer Science Thesis*. [Figure 1]. *Prospectus* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.
- Chawla, K. (2022). *Loan Metrics Handoff Model* [Figure 4]. *Class handout* (Unpublished). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.
- Chawla, K. (2022). *Loan Metrics Actor Network*. [Figure 5]. *Prospectus* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.
- Chawla, K. (2022). *Technical Solution Method*. [Figure 3]. *Prospectus* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.
- Chawla, K. (2022). *Technical System Design*. [Figure 2]. *Prospectus* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.
- Dakin, R. (2022, January 15). How banks are using machine learning. *Bits in Glass*. <https://bitsinglass.com/how-banks-are-using-machine-learning/>

- Coté, M., Criado, N., Ferrer, X., Nuenen, T. van, & Such, J. M. (2021, June 2). Bias and discrimination in AI: A cross-disciplinary perspective. *IEEE Technology and Society Magazine*, 40(2), 72–80. <https://doi.org/10.1109/MTS.2021.3056293>
- Hall, J. (2021, November 16). SRE principles: The 7 fundamental rules. *Dotcom Monitor*. <https://www.dotcom-monitor.com/blog/2021/11/16/sre-principles-the-7-fundamental-rules/>
- Hersh, I. (n.d.). The financial impact of downtime on the trading floor: \$9 million/hour. *IPC*. <https://www.ipc.com/insights/blog/the-financial-impact-of-downtime-on-the-trading-floor-9-million-an-hour/>
- Lynch, S. (2015, June 26). The benefits of having the right metrics (key performance indicators). *LinkedIn*. <https://www.linkedin.com/pulse/benefits-having-right-kpis-key-performance-indicators-stephen-lynch/>
- Pratt, M. K. (2020, June). Machine learning bias (AI bias). *Tech Target*. <https://www.techtarget.com/searchenterpriseai/definition/machine-learning-bias-algorithm-bias-or-AI-bias>
- Red Hat. What is SRE? (2020, May 4). <https://www.redhat.com/en/topics/devops/what-is-sre>
- Saguier, E. (2022, August 12). SRE signals: 3 types of metrics for site reliability engineering. *Invgate*. <https://blog.invgate.com/sre-signals>
- Watts, S. (2022, March 18). SRE metrics: Four golden signals of monitoring. *Splunk*. https://www.splunk.com/en_us/blog/learn/sre-metrics-four-golden-signals-of-monitoring.html