PLATFORM FOR ADVERSARIAL NATURAL LANGUAGE PROCESSING

NATURAL LANGUAGE PROCESSING IN FINANCIAL SERVICES

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
Soukarya Ghosh

November 22, 2020

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.

Signed: Soukarya Ghosh Date: November 22, 2020

Approved: Date:
Yanjun Qi, Department of Computer Science

Approved: Date:
Catherine D. Baritaud, Department of Engineering and Society

Natural language processing (NLP) is a subfield of machine learning and artificial intelligence that deals with the interaction between computers and humans using natural human languages. The goal of NLP is to read, decipher, understand, and piece together natural language in a manner that is more valuable in terms of quantitative analysis and decision making (Belohlavek, Platek, Straka, 2015). However, this being the ideal NLP model, there exists many shortcomings, which malicious users can exploit in order to render existing models useless and augment data in a way that causes harm to certain parties. Such practices result in adversarial examples of NLP which can be used to attack existing systems to modify their behavior and the end result (Chakraborty, Alam, Dey, Chattopadhyay, Mukhopadhyay, 2018). The most naive example of this includes changing a few characters of a word in a sentence to change the computer's perceived meaning of said sentence (Wang, Wang, 2020). Further studies in these fields have produced valuable insights into how to make current and upcoming models more secure and less prone to adversarial attacks. In order to aid this effort, I plan to design, build, and deploy a web application, by the start of 2021, that will ease the study of adversarial examples by allowing the TextAttack source code to be accessed through an approachable user interface.

With research and applications expanding in the field, NLP has caught the attention of many established industries that deal with qualitative data and content with little to no structure (Huq, Pervin, 2020). As the parsing and extracting of knowledge from qualitative data can be costly to firms, NLP has provided a leverage to automatically expedite the process here. The most interesting applications come up in the financial markets, where such techniques are deployed in order to increase speed to market, increase the volume of data to base decisions off of, and ultimately, stimulate the economy by increasing the flow of money throughout the markets (Keith, Stent, 2019). Since the condition of the economy is so closely tied to the status

of society, technologies such as natural language processing have significant impact on the lives of even those that are uninformed about its existence.

PLATFORM FOR ADVERSARIAL NATURAL LANGUAGE PROCESSING

TextAttack is a Python framework for adversarial attacks, data augmentation, and adversarial training in natural language processing (Morris, Lifland, Yoo, Grigsby, Jin, Qi, 2020, p.1). The extensive library, developed by Morris, a machine learning researcher at the University of Virginia, and managed by Dr. Yanjun Qi, is the state of the art in its field, housing over sixteen adversarial attack implementations from published and accredited literature. The ultimate goal of the project is to provide a tool to study adversarial attacks in order to improve existing model performance and robustness (Morris, Lifland, Yoo, Grigsby, Jin, Qi, 2020, p.2). As of today, in order for fellow academics and machine learning practitioners to access these powerful tools, they must install the memory heavy package locally and run costly computation on their graphics processing unit (GPU) in order to use the services. In worst case scenarios, where the user does not have a GPU to allow such intensive calculations to run, the process resorts to using cloud computation which takes significantly longer.

In order to alleviate these issues and increase usability of the service by both data science veterans and newcomers, a new method of accessibility must be introduced. The objective of this technical research is to leverage the industry full stack experience I have gathered in order to build a scalable web application with an accessible interface to allow the TextAttack library to be both demoed and, more ideally, be provided as a service for data scientists. The web application, in its final implementation, should enable users to run all attacks on both custom and pre-fetched datasets, perform data augmentation in real time with the web application's interface, and be able

to use an application program interface hosted on the web application's endpoint to run various commands that would otherwise require the installation of the command line interface. In other words, the goal of this project is to perform a full migration of the TextAttack functionalities and services to a URL endpoint for global accessibility.

There is much to be considered when designing the architecture of any full stack service, such as the language and framework to develop it in, the servers to host it on, and much more. Web applications have evolved to mature solutions with a plethora of frameworks to build the foundations of a project, which provide sophisticated architecture and base code. At the same time, the introduction of multiple frameworks to achieve the same goal has introduced increasing complexity with respect to getting different frameworks to perform as intended in unison (Kersten, Goedicke, 2010). The choices come down to the need to prioritize certain aspects over another. These include metrics such as load speeds, ease of development, language support, hosting capabilities, integration with tech stack native to TextAttack, and much more. Seeing as TextAttack is a framework only supported on Python at the moment, our search for the right web framework narrows to the single language.

The most popular Python web frameworks include Django, Flask, and Tornado (Guardia, 2016, p.31). These all have their particular strengths and weaknesses. However, when it comes to building a large-scale project, Django is the most comprehensive tool. This is because this framework is based on the model-view-controller (MVC) software design pattern which divides the related program logic into three interconnected, yet clearly separate, elements. This separation and abstraction of elements allow multiple person projects to thrive, as development on a single portion can be asynchronous of another (Qureshi, Sabir, 2014). Additionally, Django

allows more autonomy than most other frameworks, allowing developers to make key decisions on everything from code layout to system and data security.

The next big architecture resource to decide is the host of the platform and how to go about maintaining the service and keeping it scalable. With the rise of cloud computing and many large corporations migrating to cloud giants, such as Amazon Web Services (AWS), hosting on the cloud has become more user friendly than ever before. There are various offspring products of AWS that use its cloud computing capabilities to provide hosting services that prioritize ease of deployment and monitoring metrics. A prime example of this is Heroku, a cloud hosting platform that allows developers to deploy Django apps with the click of a single button and minimal setup. Nevertheless, due to the large array of tools AWS provides for scalability, including an intelligent load balancer, this will be the primary option for hosting for this platform. However, with hosting comes monetary cost that scales significantly as usage increases. In order to pay these costs, we will be applying for grants directly from AWS to either cover a portion or the entirety of the cost. As the service is going to be in development for the foreseeable future, large costs are not to be considered at the moment and will be dealt with as the application gets closer to production.

The system design shown in Figure 1 details the flow of the code from when it is pushed to the code repository on GitHub, all the way to when the users of the platform get to access said code. The bulk of the diagram focuses on the AWS suite of features, which allows the entire process of deployment, integration, and testing to be seamless and automated. Furthermore, the point of concern of scalability is handled by the Elastic Load Balancer instance in the diagram, which makes sure that users do not face any fault tolerance or single points of failure when trying to access the web endpoint (“Elastic Load Balancing”, 2020).

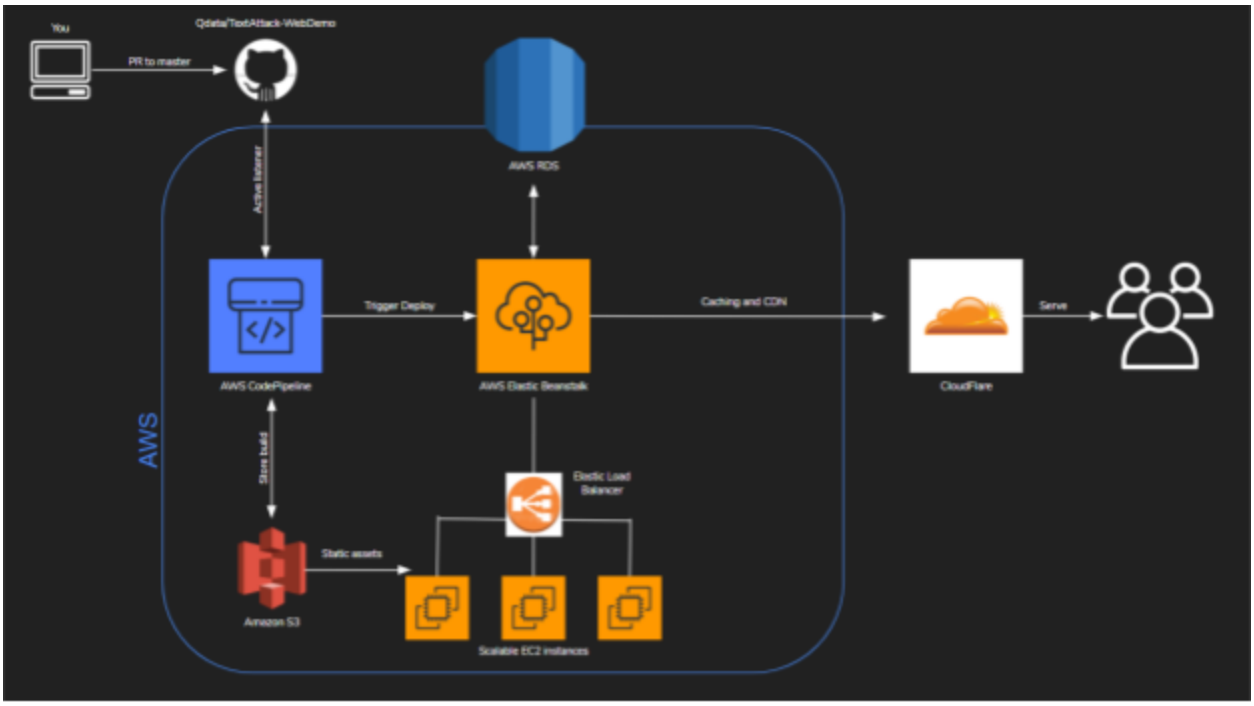

Figure 1. TextAttack Web Application System and Architecture Design. Outlines the flow data from initial user interaction all the way to the resulting interface. (Ghosh, 2020)

The initial offering will come with a small suite of features and applications, however, seeing as this will be an ongoing project after my departure, there will be a roadmap of future implementations and paths. After completing a minimal viable product phase of the web application, I hope to write a scholarly article detailing the implementation and usage of the platform and explaining the vision for this platform: an adversarial playground as a service for researchers and an educational tool for newcomers and those who aim to improve adversarial algorithms and advanced natural language processing technology for all industries.

NATURAL LANGUAGE PROCESSING IN FINANCIAL SERVICES

The financial markets and the economy have a symbiotic relationship. When one is thriving, it boosts the other in a cycle of positive gains. However, during downturns, one spirals with the other into a rock bottom status that ends up costing countless lives. The financial

markets in the United States have been the harbinger of the Great Depression as well as the most recent significant recession in 2008. Ultimately, this means, the investors and key actors of the financial markets end up having a large role in society, economy, and the trajectory of a nation's, and potentially the entire world's, future successes and failures.

With this responsibility to uphold the basis of societal life, many firms and actors in the space strive to discover and obtain the state of the art in terms of research, technology, and methodologies. As such, with natural language processing seeing unprecedented growth in all industries, as made apparent in Figure 2. One of the lead examples of this being the financial markets, NLP driven analysis has become a basis of decision making for financial firms (Yang, Rosenfeld, Makutonin, 2018, p.2). A significant problem of gaining an edge on the information war that takes place between competing firms is the accurate and timely analysis and assessment of qualitative, unstructured content.

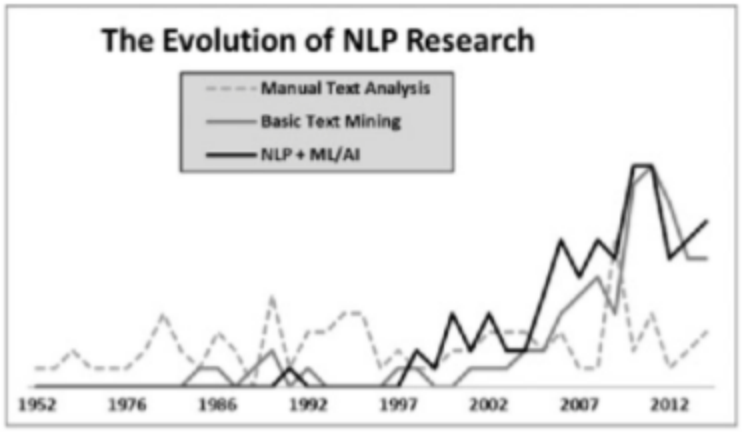

Figure 2. The Evolution of Analysis of Unstructured Content. Shows the rise of machine learning usage to analyze and understand text as opposed to manual analysis techniques. (Fisher, Garnsey, Hughes, 2016).

This can range from research reports by experts, public company filings, quarterly earnings reports, and much more. Prior to NLP approaches, all of this data required manual analysis, which funneled down the pipeline and created even more qualitative reports (Araci, 2019). There was a significant barrier to entry into the market for smaller firms that simply did

not house the manpower to crunch down reports across the market to make a smart and informed investment decision (Chen, Huang, Chen, 2020).

Natural language processing made it possible to review unstructured content and spot trends across world markets. The current use cases range from content enrichment, which tags and categorizes target data and literature, to sentiment analysis, which assigns emotion and sentiment score to literature and qualitative data (Xing, Cambria, 2017, p.3). These additional metrics help analysts and investors make better informed decisions and streamline risk management and compliance, which is a more relevant sector than ever before, due to the impact of COVID-19 on the financial markets (Sun, Belatreche, Coleman, McGinnity, Li, 2014). Other applications in the field are more directed towards extraction of information of live qualitative data and developing stories across all markets (Costantino, Morgan, Collingham, Carigliano, 1997, p.116).

Nevertheless, the applications of this technology in this sector is in its infancy and there remains much to be researched and applied to all relevant problems and subfields. Some examples include social media posts, which play a large role in the volatility of a company's valuation. One such example was when Mark Zuckerberg, the CEO and founder of Facebook, publicly stated, "We want Facebook to be somewhere where you can start meaningful

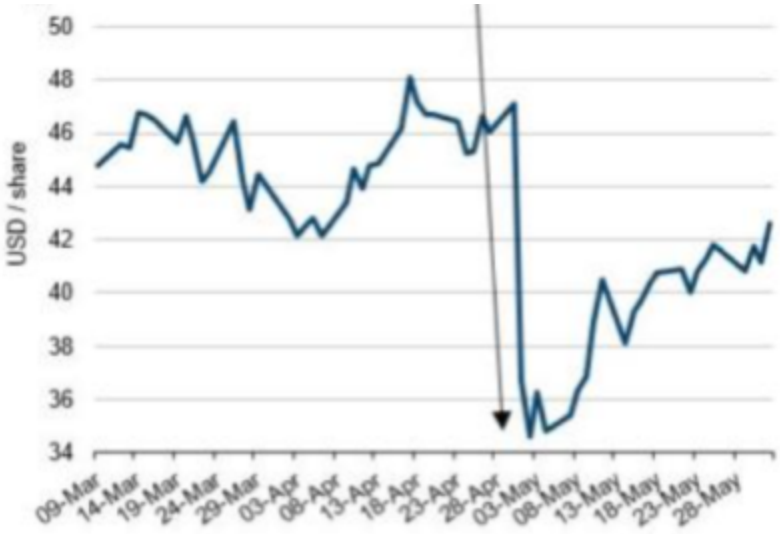

Figure 3. Share price of Macth.com in 2018. Price plunged on May 1, 2018 as Facebook announced that it would integrate a feature for online dating directly onto its app (Marinoc, AHL, 2019).

relationships." Within the hour of this event, stock prices of companies such as Match.com, Tinder, and Bumble showed plunges of over 20%, as shown in Figure 3 (Levin, 2018). These are unpredictable and volatile events that can be easily capitalized upon by firms that have the power to harness the utilities of NLP. Similar events have happened with CEOs of other notable market players, such as when Elon Musk caused a ten percent loss in the valuation of his own company's stock, TSLA, after claiming he believed the price was too high (Siddiqui 2020, p.1).

Natural language processing has a clear stake in this problem as a solution technique and although it is currently being utilized to the extent mentioned above, there exists a great deal of room for advancement. More than the firms crafting portfolios and managing risk with NLP, stakeholders on a more micro level exist. These examples are the academic institutions, small businesses, and even individual civilians who pass their wealth off to money managers utilizing the technologies. In these cases, for the money management firms to stay competitive, the technology has formulated out of necessity. Thus, we are able to say the social implications and pressures have induced the rise of the technical advancements. Hence, sociotechnical, more specifically Social Construction of Technology (SCOT), analysis is relevant to such an entity. In Figure 4, the roles of interactions of researchers, financial experts, general public, machine learning engineers, and the NLP technology itself. The engineers are in charge of implementing the research conducted by the scientists and researchers, which is influenced heavily by public opinion of existing technology. The implementations are also influenced by the input of financial experts, who utilized such technologies to increase the firm's profits. Finally, the financial experts also collect feedback, in the form of participation in the markets, from the general public to gauge the effectiveness of the NLP technology in the resulting portfolios.

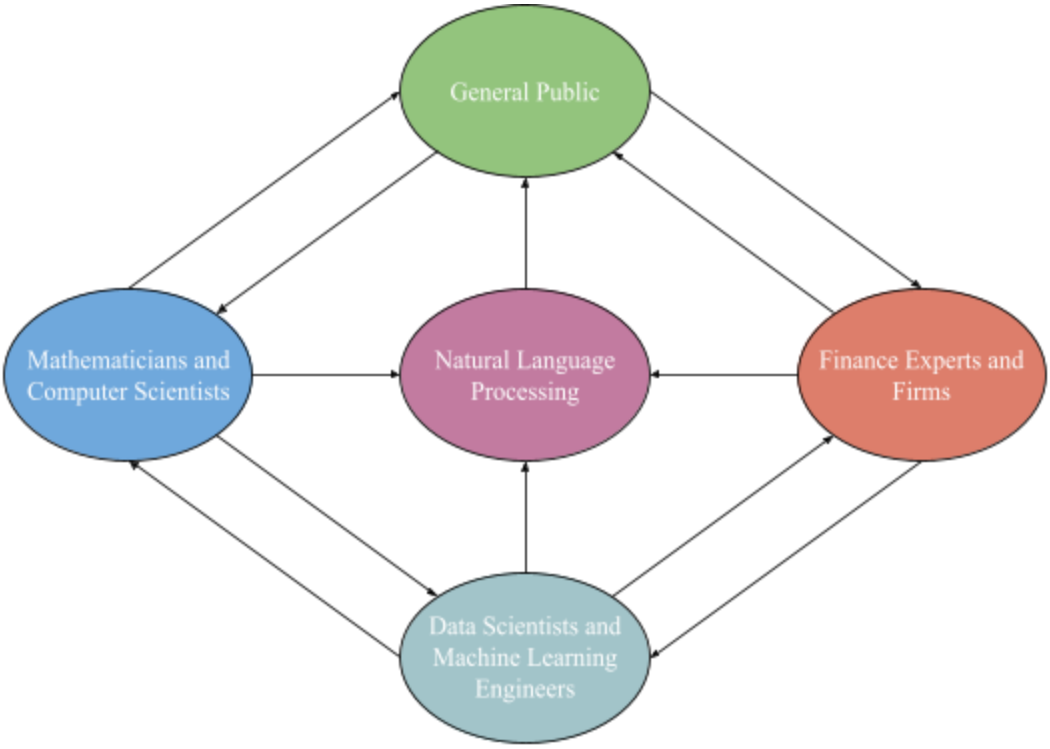

Figure 4. Natural Language Processing SCOT Model. Machine learning engineerings, academic researchers, and finance experts work together to produce the natural language processing algorithms and application, specialized towards applications of understanding unstructured data, which works to benefit the general public's wealth. The public voices approval by deciding to participate in the financial markets through firms that utilize this technology. This feedback fuels

In the analysis of NLP usage in the finance industry, I hope to clarify the impact of the interactions of the different stakeholders and how it ultimately formulates the machine learning models that are built and the effect these advancements have on society. To address and communicate my findings, I hope to write a scholarly article. Furthermore, I hope to address the impact of usage of such technologies on the economy as a whole.

WORKS CITED

Amazon. (2020). Elastic Load Balancing. https://aws.amazon.com/elasticloadbalancing/.

Araci, D. (2019, August). FinBERT: Financial Sentiment Analysis with Pre-trained Language Models. *Cornell University*. arXiv:1908.10063.

Belohlavek, P., Platek, O., Straka, M. (2015, January). Using Adversarial Examples in Natural Language Processing. *ACL*. https://www.aclweb.org/anthology/L18-1584.pdf.

Chakraborty, A., Alam, M., Dey, V., Chattopadhyay, A., Mukhopadhyay, D. (2018, August). Adversarial Attacks and Defences: A Survey. *Cornell University*. arXiv:1810.00069.

Chen, C., Huang, H., Chen, H. (2020, May). NLP in FinTech Applications: Past, Present and Future. *Cornell University*. arXiv:2005.01320.

Costantino, M., Morgan, R., Collingham R., Carigliano, R. (1997, March). Natural Language Processing and Information Extraction: Qualitative Analysis of Financial News Articles, Proceedings of the IEEE/IAFE 1997 Computational Intelligence for Financial Engineering (CIFEr). pp. 116-122. 10.1109/CIFER.1997.618923.

Fisher, I., Garnsey, M., Hughes, M. (2016). Natural Language Processing in Accounting, Auditing and Finance: A Synthesis of the Literature with a Roadmap for Future Research. xuebalib.com.5838.pdf

Freedberg, S. (2017). Artificial Stupidity: Learning To Trust Artificial Intelligence (Sometimes). Breaking Defense. https://breakingdefense.com/2017/07/artificial-stupidity-learning-to-trust-the-machine/.

Ghosh, Soukarya. (2020). *TextAttack Web Application System and Architecture Design.* [*Figure 1*]. *Prospectus* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

Ghosh, Soukarya. (2020). *Natural Language Processing SCOT Model*. [*Figure 4*]. *Prospectus* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

Guardia, C. (2016). Python Web Frameworks. *O'Reilly Media*. https://theswissbay.ch/pdf/_to_sort/O%27Reilly/python-web-frameworks.pdf.

Huq, A., Pervin, T. (2020, May). Adversarial Attacks and Defense on Texts: A Survey. *Cornell University*. arXiv:2005.14108.

Keith, K., Stent, A. (2019, June). Modeling Financial Analysts' Decision Making via the Pragmatics and Semantics of Earnings Calls. *Cornell University*. arxiv.org/pdf/1906.02868.

Kersten, B., Goedicke, M. (2010, September). Browser-based Analysis of Web Framework Applications. *Cornell University*. arXiv:1009.3714.

Levin, S. (2018). Facebook Announces Dating App Focused on 'Meaningful Relationships'. *The Guardian.* https://www.theguardian.com/technology/2018/may/01/facebook-dating-app-mark-zuckerberg-f8-conference

Morris, J. X., Lifland, E., Yoo, J. Y., Grigsby, J., Jin D., Qi Y. (2020, May). TextAttack: A Framework for Adversarial Attacks, Data Augmentation, and Adversarial Training in NLP. *Cornell University*. arXiv:2005.05909.

Qureshi, R. J., Sabir, F. (2014, August). A Comparison of Model View Controller and Model View Presenter. *Cornell University*. arXiv:1408.5786.

Siddiqui, F. (2020). Tesla Stock Plummets More Than 10 Percent After Elon Musk Tweets Valuation is ‘too high’. *Washington Post*. https://www.washingtonpost.com/technology/2020/05/01/musk-tesla-stock/.

Sun, F., Belatreche, A., Coleman, S., McGinnity T., Li, Y. (2014). Pre-processing Online Financial Text for Sentiment Classification: A Natural Language Processing Approach. *IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFEr), London, 2014*. pp. 122-129. 10.1109/CIFEr.2014.6924063.

Wang, Z., Wang, H. (2020, June). Defense of Word-level Adversarial Attacks via Random Substitution Encoding. *Cornell University*. arXiv:2005.00446.

Xing, F., Cambria, E. (2017, October). Natural Language Based Financial Forecasting: A Survey. *Artif Intell Rev*. p. 50, 49–73. https://doi.org/10.1007/s10462-017-9588-9.

Yang, S., Rosenfeld, J., Makutonin, J. (2018, August). Financial Aspect-Based Sentiment Analysis Using Deep Representations. *Cornell University*. arXiv:1808.07931.