Predicting Future Energy Consumption with Machine Learning
(Technical Paper)

Working From Home: How the Advent of Remote Work Impacts the Software Engineering Industry
(STS Paper)

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
Ruthvik Gajjala

December 6th, 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.

ADVISORS

Haiying Shen, Department of Computer Science

Bryn E. Seabrook, PhD, Department of Engineering and Society

Introduction

The clattering of keyboards and loud, overlapping gossip of coworkers; this is likely what the average software engineer associated with their line of work prior to 2020. Now, picture a different scene: a dimly lit, enclosed room, the only sounds coming from the occasional chirp of a bird. Though painted in extremes, this second scenario is the new normal for many developers.

In a broad sense, remote work "generally refers to organizational work performed outside of the normal organizational confines of space and time" (Olson, 1983). While typically associated with working in one's own home, remote work encapsulates everything from satellite work centers, in which employees commute to a more localized branch office, to hybrid models, in which employees commute to the main office some days. For this research, the main focus will be on the hybrid and work-from-home models.

For these two styles of remote work, the COVID-19 pandemic proved to be the spark plug for the transition. According to the New York Times, before "the pandemic, in 2019, about 4 percent of employed people in the U.S. worked exclusively from home; by May 2020, that figure rose to 43 percent" (Goldberg, 2022). When limiting the scope to white-collar workers, the trend is even more pronounced, with a transition from 6 percent to 65 percent; the focus of the proposed research, the software engineering job, is considered to be a white-collar profession. Now that many are siloed to their homes when doing work, is productivity affected? Moreover, how does remote work affect company output? As a result, questions of job-satisfaction and work-life balance also come into play. With the research paper, I would like to dissect the effectiveness of remote work and whether is truly beneficial, from both a business and work-life standpoint.

My capstone technical research was also done completely via the work-from-home model. The research itself had to do with attempting to predict future energy usage for various "smart" buildings across the UVA campus, presenting two motivations of the impacts of remote work on software engineering. First, I was motivated to determine how my personal experience doing development work was affected by the work-from-home model, such as my ability to communicate and code effectively. Second was sustainability; how would working remotely affect energy consumption of buildings? The capstone experience will hopefully reveal from a personal perspective the effectiveness and individual impact of remote work on my mindset; alongside the research paper, this technical project will relate my experience with remote work to the general consensus.

Technical

The expression "smart home" oftentimes refers to a term that encompasses the realm of applying sustainable practices to one of the largest consumers of global energy: the building sector. As it currently stands, the building sector will not only increase in gross energy consumption, but will also expand its share of total global energy consumption, with the U.S. Energy Information Administration estimating "that the building sector's share of the global delivered energy consumption would increase from about 20% in 2018 to 22% in 2050" (Kim et al., 2022). Applying "smart" technology to homes and buildings seeks to improve energy management and transparency. The technology can come in the form of tools that allow for accessible viewing of energy consumption from a building or even the strategic layout of a space to encourage lower energy use. Even the University of Virginia (UVA) has begun the transition of applying smart technology to buildings across grounds. The transition is apparent in one of the more recent additions to UVAs arsenal of lecture spaces, Rice Hall.

The analysis that the capstone project will focus on is attempting to predict future energy consumption of Rice Hall given current data, which comes from hundreds of installed "sensors for a variety of environmental parameters to keep tabs on electricity use, air quality, humidity, occupancy, temperature, plug power and lighting power" (Richards, 2010). Alongside attempting to predict energy usage, the research will strengthen my knowledge of generating machine learning models for prediction tasks, and performing the required exploratory analysis and data processing on real-world electricity usage data.

At the broadest level, machine learning is simply the act of giving a computer input until it is able to make predictions on its own. The "learning" is done through the help of mathematics such that a computer is "trained" until it is able to recognize certain patterns in old data, such that it can then extrapolate to new data. The task of predicting energy consumption fits the general scheme of machine learning perfectly. In the capstone research, the particular type of machine learning that will be used is referred to as the "neural network." The neural network model seeks to imitate the brain, with "neurons" that receive feedback from each other, update their own state, and ultimately provide a response. Just as humans learn through past experiences with touching a flame and responding with the swift movement of their hand, these neural networks learn via large amounts of past data and expected responses.

However, due to the volatile nature of electricity usage, it is often very difficult to make predictions into the future. To combat the existing volatility, the research will explore a relatively recent variant of the neural network known as the long short-term memory (LSTM) model, proposed by Hochreiter and Schmidhuber in 1997. As stated in their research proposal for the use of LSTM models, computer scientists Sepp Hochreiter and Jurgen Schmidhuber exclaimed that their new model "can learn to bridge time intervals in excess of 1000 steps even in case of noisy, incompressible input sequences, without loss of short time lag capabilities" (Hochreiter & Schmidhuber, 1997). In layman's terms, the structure of the LSTM model allows for recognizing and prioritizing the important patterns in the electricity usage data while ignoring the random fluctuations.

The relevance of energy consumption forecasting is undoubtedly also linked to the extent of remote work present in modern society. The capstone research will seek to assess predictions in energy consumption in a computer science building during the peak of COVID-19 – the research could reveal insights regarding the shifts in energy consumption on an industrial level, and whether it makes more sense from a sustainability perspective to encourage remote work.

STS Topic

In the current digitized age, all mobile applications, phones, computers, smart televisions, and more are based extensively on underlying programs and algorithms made possible by the work of software engineers. In fact, the U.S. Bureau of Labor Statistics projects a 25% growth rate in development related roles from 2021 to 2031 (U.S. Bureau of Labor Statistics, 2022). For reference, the "average growth rate for all occupations is 5 percent" (U.S. Bureau of Labor Statistics, 2022). Likewise, top consulting firm McKinsey explored the relationship between developer velocity index (DVI) - the idea of "empowering developers, creating the right environment for them to innovate, and removing points of friction" - and revenue growth of companies (Srivastava et al., 2021). In doing so, they found that "top-quartile DVI scores correlate with 2014–18 revenue growth that is four to five times faster than bottom-quartile DVI scores" (Srivastava et al., 2021). In other words, the ability for developers to code efficiently is highly correlated to business success. Enter the age of remote work: with many software engineers working from home and some in the office, questions arise regarding communication and productivity. With developers being an integral part of large businesses, it is of concern whether businesses are producing less than before.

The stakeholders involved in the rapid transition to remote work range from the individual developer to the entire software industry. From the individual perspective, the impact of remote work is highly unique. A potential example of the individuality is transgender software developers, who currently face disproportionate economic hardship in the job markets. Some transgender developers feel that remotely, the "advantage to being a software developer is their control of how they share their professional identity across platforms" (Ford & Serebrenik, 2019). The concept of identity disclosure is likely much more beneficial to certain marginalized groups than the average software engineer, and thus the impacts cannot be generalized to all. Tech companies are also affected by the remote transition. With software roles going remote, software adjacent roles, such as program managers who "define and design products and services, collaborating with many different stakeholders, including engineering teams" will also be impacted (Ford et al. 2022). It will be up to companies on how to decide whether these roles will also be transitioned to remote work, especially with concerns of equality and equity. Finally, the software industry will also be impacted. As it stands, the "typical retention period for permanent employees within the sector is between 18 and 36 months" (Scholarios & Marks, 2004). Such a short tenure may be drastically changed if engineers become satisfied with the benefits of remote work, also changing their power and negotiation ability in the job market.

To analyze the research, the social construction of technology (SCOT) framework will be used. In literary terms, SCOT argues that "the development of technology can be explained as a social process in which an open variety of relevant social groups participate" (Bijker 2015). Rather than humans being deterministically shaped by technology, SCOT puts forth the idea that humans are the social power that influence and define technology. Examples of human influence can be seen with the development of anything from medieval castles to modern day graphics cards. Different social groups have different preferences and present differing designs for technology; the clash of ideals itself is what generates evolution in existing and new artifacts of technology.

On the other hand, one of the main criticisms of the SCOT framework is that it makes a major assumption on the social groups that are able to influence a technology. This brings forth questions such as what social groups are considered to be “relevant” to the process of solving technical problems, and whether some groups are “suppressed or deliberately excluded” (Winner 1993). Lack of representation is an important factor, as it then suggests that silenced social groups then deterministically bear the consequences of the given technology. In the case of remote-work, this is particularly important, as two of the primary social groups are individual developers and large corporations - corporations likely favor output and employees favor benefits and pay. A silencing of workers in favor of more powerful social groups is a real possibility in the research, and I will account for the criticism by attaining data that represents both the individual and the industry. This will hopefully show that SCOT is the ideal framework to be used in the research, and that the root of influence on remote work is not defined wholly by output-oriented conglomerates. Another main criticism of the SCOT framework is that it focuses too much on the social influences that define a technology, and ignores the consequences of the social influences that are not well defined by the “social groups” in the framework. I will address this point by observing the impacts upon software engineers and related entities, both in a productivity and work-life sense, and frame them in a manner that shows a clear relationship to the social choices that are made in remote-work.

Methodologies

Research Question: How does working remotely affect the productivity and personal lives of software engineers and the success of their respective places of employment?

For my research question, I will be using the documentary research method. For the documentary approach, I will start by collecting scholarly sources regarding trends in the software industry prior to the work-from-home boom. These sources will primarily include studies concerning the keywords of retention and turnover in the software industry; these keywords will be important in assessing the market power of software engineers with the advent of remote work. Traditionally speaking, "high staff turnover has a negative impact on software development productivity and product quality" (Bass et al., 2018). From a research standpoint, it would be valuable to explore whether a company can retain a consistent workforce via remote-work policies, avoiding the cost and time of training new hires. It may also reflect the flexibility of engineers in finding new jobs and opportunities with other companies, potentially reducing stress and improving work-life-balance.

Second, I will research strategies employed by highly successful companies. An example of such a company is Google, in which the "working environment … is so comfortable so that employees will not think of it as a working room, with a full area of work, relaxation, exercise, reading, watching movies" (Tran, 2017). Parallels can be drawn from these case studies to qualities that are present when working from home, which can then be used to draw conclusions on what drives productivity for a developer. Third, I will then look into literature that employs survey-based methods on employees that are currently working from home. The surveys will reveal the sentiments of software engineers; the survey results will then be critical in determining whether sentiments regarding remote work are consistent across different positions of varying levels of experience. Then, I will be able to synthesize the information collected from the sources to determine whether survey-based methods are consistent with the larger overarching trends in the software industry. Finally, I will attempt to make a conclusion on whether remote work is a generally beneficial strategy for software engineers.

Conclusion

Overall, this prospectus aims to deliver the appropriate background information and motivation to precede an in-depth research paper on whether remote work is a beneficial strategy in the software industry. The technical capstone portion of the paper will focus on attempting to predict future energy consumption in Rice Hall by using neural networks and similar modeling approaches. The aim of the capstone is to draw any apparent observations in the data; in particular, the goal is to assess whether energy consumption can even be accurately predicted 1, 5, or even 10 days into the future. Due to the volatile nature of energy usage over a period of time, the anticipated outcome of the research will likely show that short-term predictions will be significantly more accurate than long-term predictions. Additionally, it will be anticipated that the ability to conduct research will be impacted by the work-from-home setting.

On the other hand, the STS research paper will explore trends in remote work and how they affect both the social and work lives of employees and their respective businesses. Utilizing firsthand accounts from software engineers and literature surrounding sentiments regarding the transition to remote work, the research will attempt to define in which scenarios remote work is effective. The supporting research will likely lead to the result that remote work is far from a one-size-fits-all solution. Rather, it will be highly dependent upon the experience and drive of the employee as well as the bureaucracy of the company.

References

Bass, J. M., Beecham, S., Razzak, M. A., & Noll, J. (2018). Employee retention and turnover in Global Software Development. *Proceedings of the 13th International Conference on Global Software Engineering*. https://doi.org/10.1145/3196369.3196375

Bijker, W. E. (2015). Technology, Social Construction of. *International Encyclopedia of the Social & Behavioral Sciences*, 135–140. https://doi.org/10.1016/b978-0-08-097086-8.85038-2

Ford, D., Milewicz, R., & Serebrenik, A. (2019). How remote work can foster a more inclusive environment for transgender developers. *2019 IEEE/ACM 2nd International Workshop on Gender Equality in Software Engineering (GE)*. https://doi.org/10.1109/ge.2019.00011

Ford, D., Storey, M.-A., Zimmermann, T., Bird, C., Jaffe, S., Maddila, C., Butler, J. L., Houck, B., & Nagappan, N. (2022). A tale of two cities: Software developers working from home during the COVID-19 pandemic. *ACM Transactions on Software Engineering and Methodology*, *31*(2), 1–37. https://doi.org/10.1145/3487567

Goldberg, E. (2022, March 10). *A two-year, 50-million-person experiment in changing how we work*. The New York Times. Retrieved September 25, 2022, from https://www.nytimes.com/2022/03/10/business/remote-work-office-life.html

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Kim, D., Yoon, Y., Lee, J., Mago, P. J., Lee, K., & Cho, H. (2022). Design and implementation of smart buildings: A review of current research trend. *Energies*, *15*(12), 4278. https://doi.org/10.3390/en15124278

Olson, M. H. (1983). Remote Office work. *Communications of the ACM*, *26*(3), 182–187. https://doi.org/10.1145/358061.358068

Richards, Z. (2012, April 17). *How rice hall will help trim buildings' energy consumption*. UVA Today. Retrieved October 20, 2022, from https://news.virginia.edu/content/how-rice-hall-will-help-trim-buildings-energy-consumption

Scholarios, D., & Marks, A. (2004). Work-life balance and the software worker. *Human Resource Management Journal*, *14*(2), 54–74. https://doi.org/10.1111/j.1748-8583.2004.tb00119.x

Srivastava, S., Trehan, K., Wagle, D., & Wang, J. (2021, August 31). *Developer velocity: How software excellence fuels business performance*. McKinsey & Company. Retrieved September 25, 2022, from https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/developer-velocity-how-software-excellence-fuels-business-performance

Tran, S. K. (2017). Google: A reflection of culture, leader, and Management. *International Journal of Corporate Social Responsibility*, *2*(1). https://doi.org/10.1186/s40991-017-0021-0

U.S. Bureau of Labor Statistics. (2022, September 9). *Software developers, Quality Assurance Analysts, and testers : Occupational outlook handbook*. U.S. Bureau of Labor Statistics. Retrieved October 4, 2022, from https://www.bls.gov/ooh/computer-and-information-technology/software-developers.htm

Winner, L. (1993). Upon opening the black box and finding it empty: Social constructivism and the Philosophy of Technology. *Science, Technology, & Human Values*, *18*(3), 362–378. https://doi.org/10.1177/016224399301800306