Developing an Application Utilizing Artificial Intelligence to Generate Customized Short Stories for Use In Classrooms
Technical Topic

Examining the Societal and Environmental Impacts of AI Development: Inequities and Infrastructure
STS Topic

A Thesis Project 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
Ryan Chung

November 18, 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.

Signed: _____

Advisor

Dr. Coleen Carrigan, Department of Engineering and Society

Introduction

Developing an application like StorySprout was a deeply personal and exciting journey for me. StorySprout was created to use OpenAI's ChatGPT to generate customized children's stories for use in classrooms. The objective was straightforward: to motivate students and provide educators with a tool that enhanced learning via creativity and engagement. I was first enthralled by AI's potential because of its capacity for adaptation, scalability, and universal access to creativity. I was honored to be a part of a project that looked to have limitless possibilities.

However, one day during the development of StorySprout, I was shown an article stating that one ChatGPT prompt used a sixteen ounce bottle of water. Following that, I began to confront a side of artificial intelligence that I hadn't considered deeply enough before. Behind the light-hearted user generated stories lay a network of energy-hungry data centers, a growing carbon footprint, and societal costs that were impossible to ignore. Each line of code I contributed to StorySprout was powered by servers running 24/7, consuming resources that strained local communities and ecosystems. These realities raised urgent ethical questions about the long-term sustainability of my work.

My perspective shifted dramatically as I learned more about the broader implications of AI. I realized that while tools like StorySprout could provide a customizable and enjoyable experience for grade school children, they also perpetuated a cycle of environmental and societal harm. Data centers, the unseen catalysts of AI, disproportionately impact marginalized communities, raising energy costs and accelerating gentrification in areas like Northern Virginia. The very systems that made my app possible were also emblematic of AI's double-edged nature:

its benefits concentrated in the hands of a few, while its costs were distributed widely, often on the shoulders of those who couldn't bear the large cost of technological progress.

Ultimately, I made the difficult decision to leave StorySprout. This choice stemmed not from a lack of belief in technology's potential but from a growing recognition of its current trajectory. AI development, I realized, disproportionately benefits a narrow group of tech elites, such as Silicon Valley entrepreneurs and graduates of elite institutions, while the environmental and societal costs harm millions. This trend is driven by the profit motives of the tech sector and venture capital priorities.

Why, then, do we continue to push AI despite its harms? Who ultimately benefits from this relentless development? These questions have fueled my reflection and guided the direction of this prospectus. As I examine the nuances of AI's influence on society, I want to look at both the potential solutions to its most urgent problems as well as the injustices it perpetuates.

Case Study on YCombinator

YCombinator (YC) is known worldwide as a launchpad for innovation, fostering startups that shape the global technology landscape. From companies like Airbnb to Stripe, YC has consistently driven its cohorts toward rapid scalability and market disruption. However, behind this celebrated narrative lies a replicative ethos—one that sculpts an "ideal" entrepreneur, mirroring the myth of Pygmalion and pushes for profit regardless of the cost. Much like the sculptor who created his idealized Galatea, YC shapes its cohorts in the mold of male, elite-educated, and tech-savvy individuals, perpetuating structural biases in the venture capital ecosystem.

In recent years, YC has pivoted heavily toward artificial intelligence (AI), reflecting the broader enthusiasm of venture capitalists. In its Winter 2024 batch, 86 startups—nearly half the

cohort—focused on AI technologies (Shah, 2024, p.1). This shift underscores YC's cultural narrative: technological disruption is synonymous with progress, and AI stands as the leader of that transformation. However, this emphasis often sidelines initiatives addressing systemic inequalities or non-technological innovations, reinforcing a narrow definition of success that prioritizes rapid returns on investment over diverse societal needs.

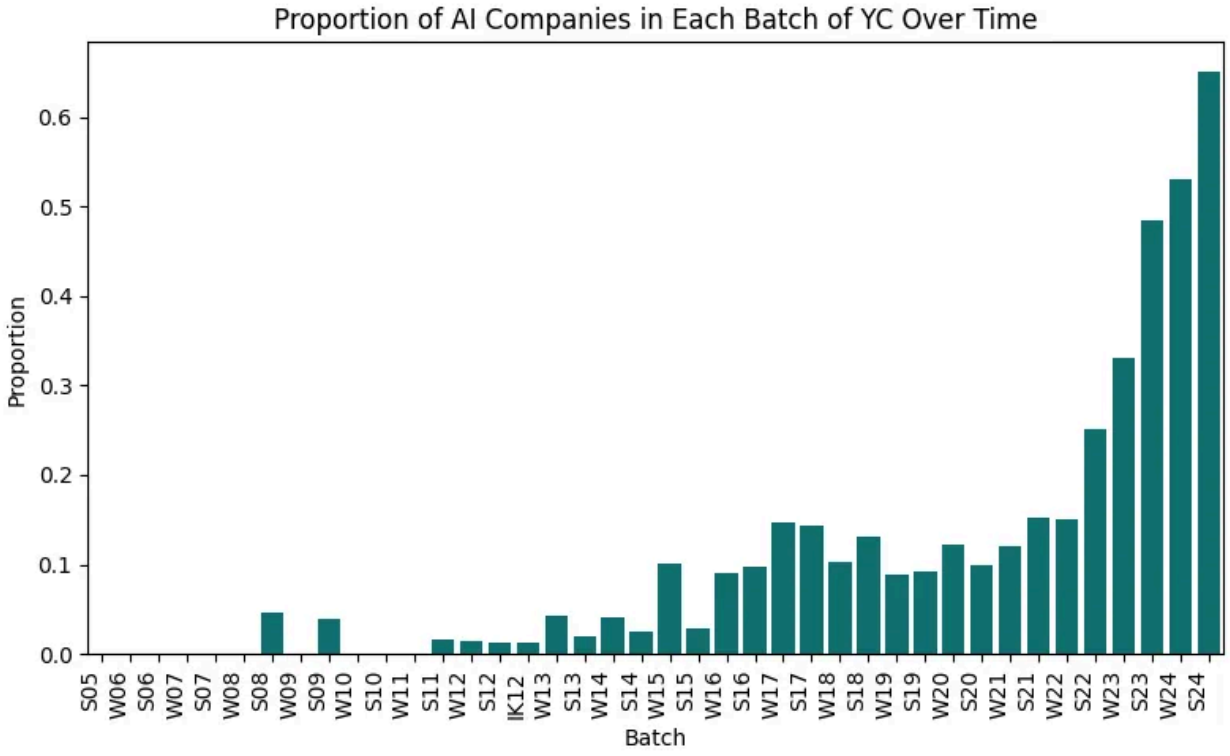

wa

Research and Data Analysis

Data from recent YC cohorts highlights consistent and obvious disparities in representation. Reports have shown that over 45% of founders are graduates from elite institutions such as Stanford, Harvard, MIT, and the University of California-Berkeley (Chung, 2024, p.1).

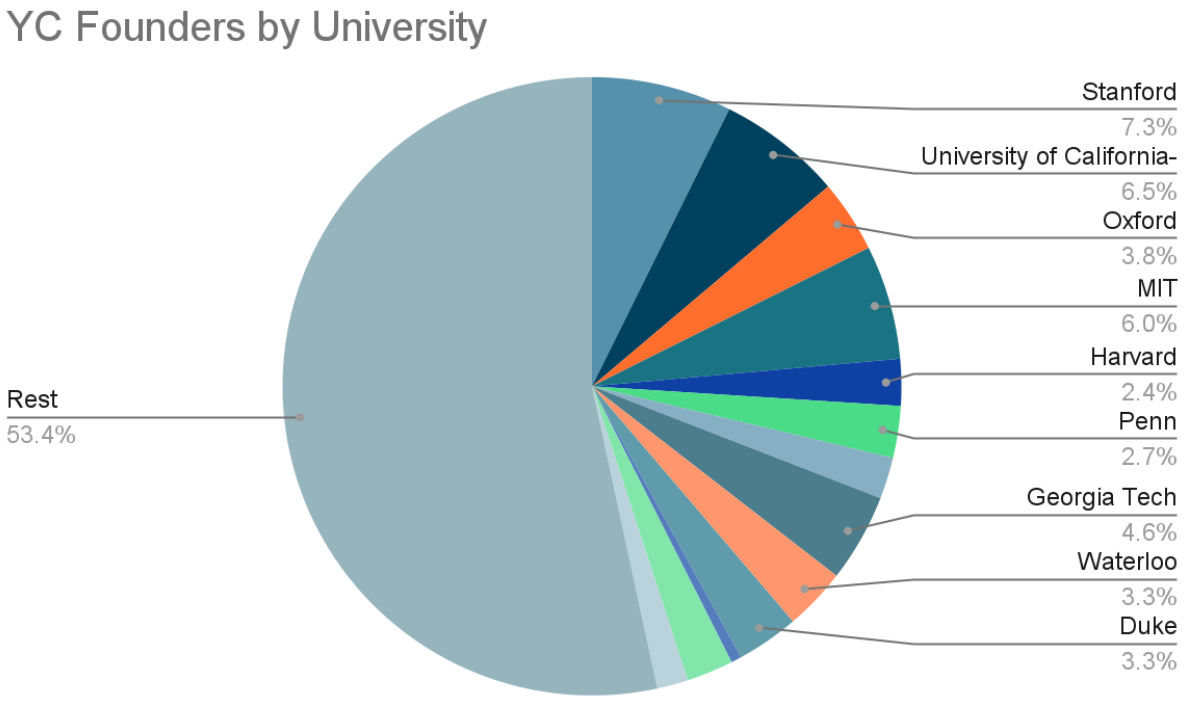

Figure 2: YC Founders Categorized by University (Alma Mater or Attending) (Chung, 2024)

Reports have also shown that less than 11% of founders identify as women (Tan, 2024, p.1). This lack of female voice in YCombinator represents an even more wide spread disparity in venture capital funding for female founded startups. Since 2015, the amount of venture capital funding being allocated to female-only founded startups has fluctuated between two and three percent (Teare, 2024, p.1).

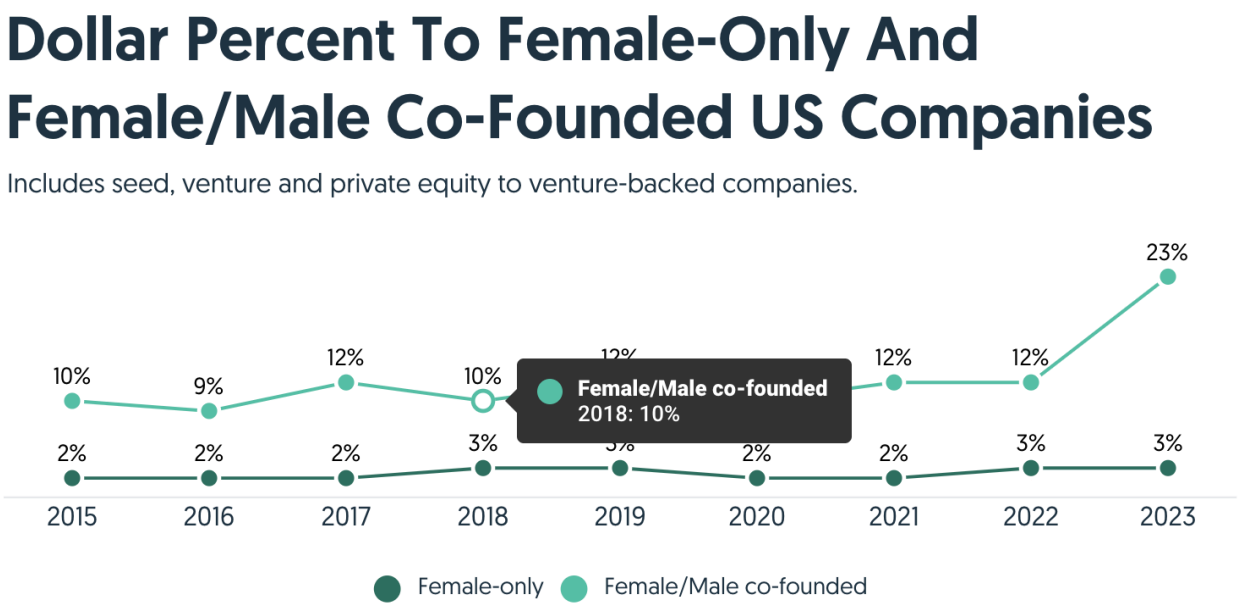

Figure 3: Percentage of Venture-Backed Funding going to Female-Only Founders (Teare, 2024, p.1)

These figures illuminate a structural bias within YC's selection process and in many ways represents structural bias throughout the landscape of venture capital, which favors individuals embedded within not only privileged networks, but individuals who fit the figurative “perfect mold”. By emphasizing credentials and connections from elite institutions, YC effectively narrows the pipeline of entrepreneurs deemed "worthy" of investment.

Exclusionary behaviors are justified by the cultural myths of "meritocracy" and "disruption," which form the foundation of YC's strategy. Meritocracy disregards structural inequalities in access to opportunities, resources, and education in favor of assuming a level playing field. Contrarily, disruption stresses quick change and frequently puts short-term benefits ahead of long-term, community-driven solutions. When combined, these stories marginalize people who don't fit the mold of the "ideal" founder.

The pivot toward AI within YC’s cohorts further highlights its structural biases. AI startups dominate the Winter 2024 batch, reflecting venture capital’s enthusiasm for AI as a

driver of exponential growth (Shah, 2024, p.1). However, this focus often comes at the expense of startups addressing critical societal challenges that fall outside the tech industry. Projects centered on education, environmental sustainability, or community health may struggle to gain traction in an ecosystem hyper-focused on technology's ability to "disrupt" established markets.

The emphasis on disruption is itself a cultural narrative that aligns with YC's ethos. "Disruption" as a value often fetishizes rapid change, sidelining slower, community-driven efforts that might provide more equitable or sustainable solutions. This narrative, while celebrated as meritocratic, effectively excludes founders whose ideas do not align with these norms, further perpetuating inequities.

Meritocracy—the belief that talent and hard work alone determine success—is a central tenet of YC's philosophy. However, the systemic obstacles that minority groups face are hidden by this notion. Regardless of the quality of their ideas, founders who lack connections to Silicon Valley ecosystems, venture capital networks, or prestigious universities are frequently passed over. As a result, YC's cohorts perpetuate the same limited professional and demographic stereotypes year after year, creating a self-reinforcing cycle.

Conclusion and Implications

YCombinator's success in fostering innovation is undeniable. Yet, its model raises critical questions about who benefits from this innovation and at what cost. By shaping its cohorts to reflect a narrow definition of the "ideal" entrepreneur, YC perpetuates structural inequalities that limit access to entrepreneurship for women, individuals from non-elite backgrounds, and those focused on non-technological solutions.

The parallels to Pygmalion's myth are striking. YC's entrepreneurs are molded to align with the expectations of venture capitalists seeking rapid returns, much like Galatea was sculpted

to embody Pygmalion's vision of perfection. This model, while efficient in producing high-growth startups, excludes those who do not conform to its archetype.

To remain a leader in fostering innovation, YC must confront these structural issues. Expanding its selection criteria to prioritize diverse experiences and perspectives could help it foster a broader range of solutions to global challenges. Additionally, reducing its reliance on narratives of meritocracy and disruption would allow it to embrace community-driven and incremental approaches to innovation.

YC's influence on the startup ecosystem means that these biases extend beyond its own cohorts, shaping the trajectory of technological progress in ways that exclude and marginalize.

Who Is Harmed by The AI Revolution?

Artificial intelligence (AI) technologies have emerged as key components of innovation and advancement due to their quick development and application. However, this change has a high price, one that underprivileged populations bear a disproportionate share of. Systemic injustices in resource distribution, environmental damage, and community dislocation are caused by the infrastructure driving the AI revolution, especially the growth of data centers. Northern Virginia, an area essential to the global AI ecosystem, is where these effects are most evident.

With more than 70% of global data traffic residing in Northern Virginia, particularly Loudoun County, the region has become the center of the internet (Rosati, James, Metcalf, 2023, p.10). The processing and storing of data that is essential for training AI models is made possible by this concentration of data centers, but at a huge environmental cost. Data centers put unsustainable demands on local resources by using enormous amounts of water and power. These establishments contribute 0.5% of greenhouse gas emissions in the United States, and their water use exacerbates stress in some places (Siddik, Shehabi, Marston, 2021, p.8). Low-income

households are disproportionately affected by this stress in places like Northern Virginia, where they must contend with growing energy bills without taking advantage of the financial incentives given to tech companies. More precisely, the impact on finances brought on by excessive utility use can lead to difficult situations for low-income families, such as forgoing medical care or other essential needs (Hernandez, 2023, p.1).

Companies profit but local communities bear the price of environmental degradation because corporate tax incentives used to entice data centers frequently come at the expense of public investment in community services. "Demands for cloud services have outpaced sustainable growth, leading to cascading systemic failures, data rationing, and more," emphasizes Monserrate, emphasizing how the energy, water, and land requirements of data centers are given precedence over the welfare of the local communities (Monserrate, 2022, p.18).

Data centers' quick growth has also changed the housing market in Northern Virginia, causing gentrification and displacement. Long-term inhabitants are frequently priced out of their houses when property values rise in proximity to data centers. According to research conducted in the Netherlands, the average transaction price of a home next to a data center is 5.27% more than it was prior to the data center's construction (Tillburg, 2022, p.1). Northern Virginia has seen a 65% increase in median house sale prices over the past ten years, reflecting this trend (Staff, 2024, p.1). Since richer tech workers relocate into the neighborhood, traditionally underprivileged populations suffer disproportionately from this dynamic, since their access to inexpensive housing is reduced.

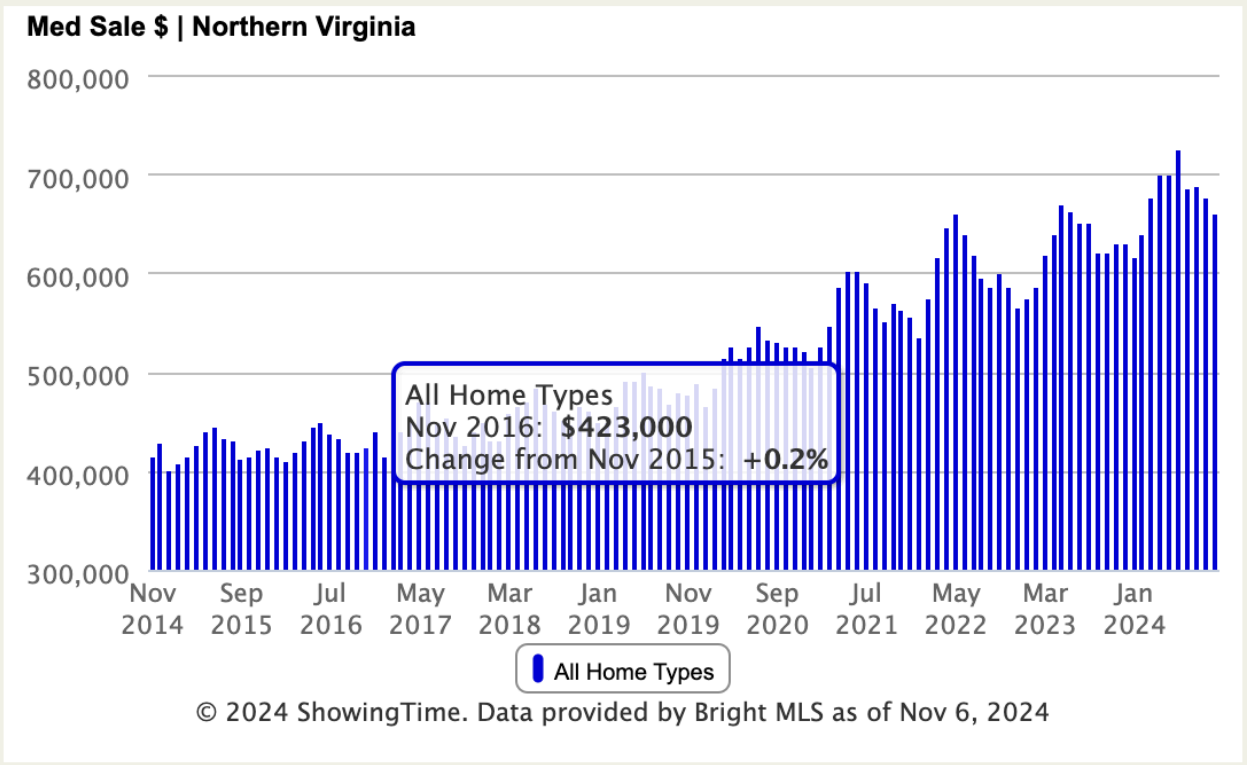

Figure 4: Median Sale Price of Homes in Northern Virginia Over 10 Years (Staff, 2024, p.1)

The socio-economic displacement caused by data centers reflects broader patterns of exploitation in technological development. Rosati compared the growth of data centers to "digital plantations," where elites profit economically from advances in technology but communities bear the costs (Rosati, James, Metcalf, 2023, p.1). Because profits are given precedence over egalitarian development, this framework highlights the injustices ingrained in AI infrastructure.

Beyond their immediate effects, data centers and artificial intelligence might cause deeper systemic injustices. Rich tech elites have a disproportionate amount of influence over municipal laws, influencing choices that put technology development ahead of equity and sustainability. This power dynamic perpetuates a technology-first mindset that exacerbates already-existing imbalances by silencing community voices calling for more balanced development.

There is an urgent need to address the detrimental effects of AI and the infrastructure that supports it. Community-driven solutions must be given priority, and scholars and policymakers must look into how the expansion of data centers perpetuates structural injustices. Addressing these issues requires fair resource distribution, stronger environmental laws, and equitable housing policies.

Designing interventions that strike a balance between social justice and technological advancement requires an understanding of who is impacted by the AI revolution. The ethical rationale for AI's development will continue to be seriously flawed as long as the advantages of the technology are concentrated in the hands of tech elites and its costs are shared by the most vulnerable.

Conclusion

The fast growth of data centers has aided in the development and application of artificial intelligence (AI), positioning it as a revolutionary force with wide-ranging effects. However, this change puts the interests of tech giants and global firms first, leaving the environment and underprivileged communities to bear the costs. Deep structural injustices are revealed by the AI revolution, from Northern Virginia's housing displacement to environmental degradation sufferers worldwide.

The concentrated gains going to powerful organizations are what motivate the unrelenting drive for AI research, not a lack of awareness of its drawbacks. Technology leaders, entrepreneurs, and venture investors influence public opinion and legislation to put profit and innovation ahead of sustainability and equity. These players work within a cultural narrative that elevates technological advancement as intrinsically beneficial while marginalizing critical viewpoints that raise concerns about its wider societal effects.

AI's trajectory will continue to worsen current disparities if its environmental and societal implications are not addressed. The need for more inclusive and accountable approaches to technological advancement is highlighted by the surge in housing and energy prices in areas that host data centers, as well as evidence of corporate incentives that put profit before equitable development.

Future studies must elevate underrepresented perspectives and examine how overlapping power structures affect how the advantages and disadvantages of AI are distributed. To lessen the negative effects of data centers on the environment and society, policymakers must adopt well-informed plans that include energy equity, affordable housing, and enhanced environmental safeguards. These interventions must highlight community needs and values, opposing the technocratic perspective that has dominated AI development to date.

Ultimately, society must confront a fundamental question: What kind of progress do we truly value? If the price of AI is borne by those with the least power to resist its harms, its development cannot be justified without addressing these inequities. The path forward requires a collective reevaluation of priorities—one that places equity, sustainability, and accountability at the forefront. Only by shifting the focus from profit to people can we build a future where technological progress benefits all, rather than a privileged few.

References

Shah, D. (2024, October 15). Data-driven insights into Y-Combinator. Medium. https://medium.com/@devanshshah1309/data-driven-insights-into-y-combinator-f8b695e184bf#:~:text=The%20Rise%20of%20AI&text=The%20above%20graph%20shows%20that,landscape%20and%20where%20it's%20headed.

Teare, G. (2024, June 5). The portion of US VC funding that went to female founders hit a new peak in 2023, thanks to massive AI deals. Crunchbase News. https://news.crunchbase.com/diversity/us-vc-funding-female-founders-peaked-2023-ai-openai-anthropic/

Tan, G. (2024, March 4). Meet the YC winter 2024 batch. Y Combinator. https://www.ycombinator.com/blog/meet-the-yc-winter-2024-batch

Rosati, C., James, A., & Metcalf, K. (2023). Data plantation: Northern Virginia and the territorialization of digital civilization in "the Internet Capital of the World." Online Media and Global Communication, 2(2), 199–227. https://doi.org/10.1515/omgc-2023-0017

Hernandez, D. (2023, June 29). Energy Insecurity And Health: America's Hidden Hardship. Health Affairs. https://www.healthaffairs.org/do/10.1377/hpb20230518.472953/#:~:text=Residential%20patterns%20rooted%20in%20legacies,as%20well%20as%20chronic%20stress.

Monserrate, S. G. (2022). The Cloud Is Material: On the Environmental Impacts of Computation and Data Storage. MIT Case Studies in Social and Ethical Responsibilities of Computing, Winter 2022. https://doi.org/10.21428/2c646de5.031d4553

van Tilburg, T. (2021). Living next to the cloud: The impact of data center development on housing prices. Master's Thesis, University of Groningen.
https://frw.studenttheses.ub.rug.nl/3724/

Staff, U. (2024, May 24). The 10-year path of Northern Virginia home prices in four charts. https://dc.urbanturf.com/articles/blog/the_10-year_path_of_northern_virginia_home_prices_in_four_charts/22350