


"Much unhappiness has come into the world because of bewilderment and things left unsaid." – Fyodor Dostoevsky



My experience at Georgia Tech (GT) was, for better or worse, life-changing. I had many hopes and dreams going into the Master's in Computer Science (MScS) program there. Some of these hopes were fulfilled, and others were shattered, leaving behind long-lasting scars. While I'm happy now, I can say with certainty that in most of my time at GT, I was far from being so. With a clear and reflective mind, I look back at that experience to voice my pleasure and my pain. And by doing so, I hope to show you how much of an uphill battle it has been for me to do NLP research.

In this piece, I start by describing my background and how I got into Natural Language Processing (NLP) research. Then I explain the difficulties I faced getting into GT and how unlikely my enrollment into the program was. I move forward to tell you how things started well with my first advisor there but ultimately turned ugly. Parallel to that is a running theme of my living circumstances. Next, I discuss my time with another advisor who carried me

through the rest of the program. Finally, I explain why wrapping up the program was not straightforward.

Who Am I?

I am Taha Merghani. I came from Sudan to study computer engineering at Jackson State University (Go Tigers!) on a full-ride, starting in the fall 2013 semester. I had a senior college friend in Sudan who told me she was working in a field called Human-Computer Interaction (HCI) on her graduation project. I was hooked by the name of the field and decided I wanted to work on that kind of stuff. Since my undergrad college didn't have much research going on, and certainly not in HCI, I pursued my passion for research at bigger universities.

During my freshman year summer, I got accepted into a summer research program for underrepresented minorities at Indiana University-Bloomington. My project there ended up being developing features for an Android app that tracks fitness goals for a community of patients, with a social section where users can interact and support one another in their fitness goals. I found that I was not interested in working on the graphical interface, but rather on one that uses voice input. Specifically, on a spoken language interface that can understand all languages, accents, and dialects.

This is exactly what I did the following summer of 2015, which is the best experience I have ever had. I was accepted into MIT's Summer Research Program (MSRP) and got matched with Dr. Jim Glass' lab at MIT CSAIL. There I worked on automatic speech recognition (ASR) for Egyptian Arabic with then-PhD students Tuka AlHanai and Micheal Price. Both of whom tirelessly answered my questions. This was a time when ASR was done with deep feedforward neural networks using KALDI recipes. By the end of it, I became very interested in language understanding, and not only recognition. But I was awakened by my need to take a machine learning course online since JSU did not have that course. And for my future project, I wanted to see a fully-fledged dialogue system with a huge user base in the real world.

Again, that's exactly what I did the following summer as an intern in the Siri team at Apple, thanks to the inauguration of the Apple HBCU Scholars program in the summer of 2016. The program offered a selected cohort of HBCU students an internship at Apple as well as a scholarship during their studies. I ended up not doing research but it was still a valuable experience. I was determined to get back to working in research and focusing on NLP projects that address understanding language and its variation. And whose research interested me the most? Jacob Eisenstein's at Georgia Tech! His research seemed like the right mix of machine learning, language, and social implications.

Getting into Georgia Tech:

I applied for several Ph.D. programs, not grasping the ridiculous standards set for admissions even then. But I was convinced that if I can have some face-to-face time with potential advisors, my odds would be better. Thanks to yet another community outreach effort, the FOCUS program, which sponsors prospective applicants to visit GT, in January (right after most application deadlines for Fall 2017), I was able to meet Jacob Eisenstein in person.

I talked to Jacob about my MIT project and NLP in general and told him I was not sure if I would make it into any Ph.D. program. Visibly impressed, he asked me why I wasn't confident. I didn't know how to answer that question. I just didn't know of any Black, immigrant Ph.D. students doing the kind of stuff I wanted to do at the top NLP labs. We then chatted about his journey to GT. I remember he mentioned that he was having a lot of fun in grad school while almost everyone around him was depressed.

After FOCUS, Trump was elected 45th president of the U.S. and he immediately announced the Muslim ban, in which my country of origin Sudan was on the list of banned countries. After that Jacob sent me an email saying that I won't get accepted into the Ph.D. program and that GT had not decided if they would implement Trump's restrictions. I asked him if I could be transferred into the master's pool, how to improve my odds of doing a Ph.D. in the future, and what I could do if I can't study in the US. He told me that a master's with publications would be the best intermediate step, and in case I could not study in the US he would be happy to introduce me to NLP researchers abroad.

I then didn't hear from him again for a while. But I did hear from other Ph.D. programs, one rejection after another. And around March 2017, I started shifting gears to looking for industry jobs and summer research internships. I also communicated this frustration to Jacob who finally replied telling me that he was able to re-route my application to the MSCS pool and add a recommendation.

While that brought me hope, I still drove from Jackson, MS to the National Society for Black Engineers (NSBE) annual career fair in Kansas City, MO. On my 11-hr drive to Kansas City, I had my major racial reckoning about the reality of how the police treat Black people in America where I was arrested and spent a night in jail for significant over-speeding. But this is a story for another time. At any rate, I arrived at the career fair just in time for it to culminate. But I wasn't that bummed because on the way there, I received my acceptance letter into the MSCS program at GT.

Funding for MSCS:

Since almost all MS programs are not funded, I had to figure that out. Jacob told me in an email that it is likely that students at GT get funding starting their

second semester as they know professors. He also mentioned he wouldn't have funding for me, but he would help me look. At GT, several funding decisions are made during the first week of classes. Then I thought I can ask for a deferral while I save for tuition by working, but he didn't endorse that idea, saying that it was likely something would work out.

With all these variables to process along with my unhappiness with Jackson, MS where I was not able to be my full self and pursue my research agenda, I packed my bags and moved to the Bay Area to look for a job. Most of the tech job openings were for web, Android, and iOS developers, all of which I had little experience in. In addition, I had to brush up on my Leetcode-style interview questions which are challenging for someone who spent most of their time working on embedded systems and hardware. I applied for a number of jobs, the most promising of which were in big tech companies, since smaller companies were focused on web and app dev stuff. And as new graduate recruitment for Big Tech started in August, I ended up not getting a job. And because of student visa restrictions and my soon-to-expire OPT card, I had to suck it up and enroll in the master's program.

Fortunately, Jacob emailed and suggested I apply for the Graduate Research Assistantship (GRA) at GT's Digital Integrative Liberal Arts Center (DILAC). The center worked on a motley collection of applied data science projects in the domain of literature. DILAC's interviews would be conducted during the first week of the semester, so there were still no guarantees of funding.

Moving to Atlanta:

Burnt out from the unsuccessful job search in the Bay Area, I moved to Atlanta just the weekend before classes started and crashed at a friend's couch who was a first-year Ph.D. student at GT. I scraped the rest of my savings, and my family in Sudan had to sacrifice what they were saving for a house to support me in my first semester in the hope that I would get funding in the subsequent semesters. Being so late left me very few options for roommates and no time to get to know them. I ended up with a terrible roommate setup, which I will get to in a bit. Luckily, my GRA interviews at DILAC went well, and I got the offer! Jacob also had an opening for a TA in his Computational Social Science (CSS) class which he offered to me, but I chose the GRA. Both TAs and RAs waive the tuition fee and offer a meager stipend. The irony wasn't lost on me: I came to GT for Jacob's research, but Jacob couldn't fund me. DILAC could. So my first year would be spent earning my keep in a domain I had no connection to, working on 1800s American literature after a lifetime of Arabic-language education in Sudan, while trying to prove myself worthy of the NLP work I actually came to do.

My first semester at GT (Fall 2017):

My living conditions were abysmal. I had to live a 15-minute drive from campus in a very small apartment in a sketchy neighborhood in Westside Atlanta. I ought to make this choice because the timing was tight and funding was equally as tight. After deductions, I was to be left with about \$750 a month. My share of the rent plus utilities was almost \$500. Moreover, my roommate and I did not get along later on, and I had to do my laundry externally. I could compensate for the lack of funding with my savings, but those ran out quickly. I also couldn't afford on-campus parking, so I had to park far enough where I could find free street parking.

As for school, I took the CSS class, a graduate algorithms course, a research project with Jacob, and another at DILAC. Each of these was worth 3 credit hours to maintain the full-time status. The classes were super demanding but manageable on their own. However, I'm going to talk more about the research projects. At DILAC, I was trying to find an NLP project, but I had no future outlook for it since I was just introduced to Digital Humanities research. It turns out, neither did my advisor who was a recent Ph.D. grad in Literature. I was one of two programmers, and I had less experience with the tools necessary for the project. The setup was that my advisor asked us to implement something and we would do it. But I definitely wasn't having fun, and I felt out of place. On the flip side, my research with Jacob was exciting.

I joined what we named the SocialAttention project with second-year master's student Murali Balusu who was a graduate of IIT with extensive programming experience. Murali was ambitious, driven, indefatigable, take-charge, and well-prepared. Unfortunately, he was not quite the hands-on mentor I desperately wanted for my first project. And unlike me, he was not looking for a Ph.D. program, so he had nothing to lose if the project went awry. The stakes were higher for me because I needed to have at least one first-author paper published at top NLP conferences.

The SocialAttention project goal was to develop a model that uses social network + text to improve part-of-speech (POS) tagging on an annotated Twitter dataset. This was to overcome the drop in accuracy for POS taggers developed on well-edited newswire datasets. Social media text featured more kinds of authors and much fewer restrictions. At first, Murali pretty much left me nothing to do and I was just taking notes during the meetings until Jacob noticed and asked me to implement a baseline feature-rich, text-only POS tagger. I then had a foray into representation learning for graphs and applied LINE and Node2Vec on our Twitter author networks. By the time of the call-for-papers for the first Black in AI workshop at NeurIPS 2017, I had a poster to present as ongoing progress.

The miraculous effort by Timnit Gebru, Rediet Abebe, and others made the Black in AI workshop such a refreshing and enlightening experience. Besides presenting my poster and feeling represented at long last, I was able to meet many AI research superstars like Hal Daumé III, Chris Manning, Ian Goodfellow, and others at the Black in AI dinner. But I was consumed by asking about admissions into Ph.D. programs. What Zachary Lipton said to me at the time stuck with me the most. He told me the backstory behind his well-received solo-authored paper on model interpretability. And he said that if I could show my writing and my code as evidence of my ability to do independent research, that would be best. I returned to campus, suffered through a breakup, and wrapped up the semester. While I succeeded in my classes and the SocialAttention project, I had no deliverables for my GRA at DILAC.

Second Semester (It gets ugly):

I was excited to finally take the class I have been waiting to take my whole life: Jacob's NLP course in Spring 2018. But I had to take an equally demanding class which was Polo Chau's Data and Visual Analytics (DVA) along with the two research projects fulfilling the 12-hr full-time requirement. I also applied to summer internships. And I wanted the Microsoft Research (MSR) internship for non-PhD students the most.

Following the trend in his CSS class but much more evident this time, Jacob's NLP lectures were overly focused on unclear mathematical derivations while making students uncomfortable asking questions about them. I remember once a student asked him an optimization question and he shut him down, saying he wasn't going to answer that question, leaving the classroom in utter silence. The programming assignments were significantly clearer and well designed though, and the questions on Piazza were well answered by the TAs.

As for the research projects, Murali and I were wrapping up the SocialAttention project. We weren't able to make the joint text+network model work, but we found strong evidence that errors in the POS tagger accuracy were correlated with the network structure where 1) errors on user i suggests that errors on user j when (i, j) follow/mention each other are more likely, and 2) training a balanced sample from multiple subcommunities is better than focusing on a narrow community. These results were later accepted at the stylistic variation workshop at NAACL 2018. But as I wanted to own a project when the SocialAttention project was coming to a close, I talked with Jacob about working on dialects. This was around the time I received my rejection from the MSR internship for which I felt my interview had gone well. Stung and desperate for feedback, I sent this email to one of my interviewers reflecting on what I thought could be improved in my interviews.

-Reflection of previous research:

The conclusion I drew from my comparing the results I obtained using the feature rich models I used for part-of-speech tagging was unclear. The purpose was to use that as a baseline. Effectively, I was reproducing the results from the original paper on the topic. Since our goal was to examine whether we can leverage the social network as side information in a deep learning model to reach comparable results. This effort was unsuccessful and we recently abandoned the project.

-Coding ability:

I think I mentioned the largest amount of code I have written was around 1k-2k lines. That was a gross underestimation. I am looking back at the code I wrote last week alone, and it easily exceeds that threshold.

- Paper writing:

I think I focused on somewhat unnecessary details (i.e., obtaining the LaTeX template) and not the abstraction. I wanted to say that I partially write as I go. I first started by describing the more fixed parts, such as the description of the data, the model, and the introduction. I also keep a research log and write up intermediate details as they become clearer. The last portions of the paper to be written are the abstract and conclusion as they depend on the results of the experiments. I tried not to cram all the details to right before the submission deadline.

- Independence:

I was definitely very nervous during the interview and didn't project much confidence. Now, I am leading a new project and working independently. I am defining the research questions (with my advisor's assistance) and taking charge of reproducing previous research.

My interviewer responded that the real reason for my rejection was that there were just a very small number of spots for the program. But he added that he admired my reflections and told me that he viewed periodic self-reflection as a key to success.

On campus, Jacob suggested I work on a project applying the Hilbert-Schmidt Independence Criterion (HSIC) on geotagged French tweets to mine geographically dependent words, and then use concepts like network balance and representativeness to characterize subcommunities that use those mined words. There was a dialectal variation workshop submission deadline around the end of the semester, which we would work towards, but our meetings were to occur only once every two weeks. But at that point, I felt like I wanted to improve my expertise in deep learning, so I was hoping I could steer the project more into methodology rather than application. I had grown out of love for the network stuff then. But I was afraid to bring that up with Jacob as he seemed to have strong views on where the project should go.

I was completely overwhelmed. And my situation at home only became worse as my roommate and I didn't talk. I avoided going back home and opted to sleep in the lab instead. I was also getting some traffic tickets which by the end of the semester amounted to the cost of purchasing a parking pass.

In my project with Jacob, I was first looking into applying the method, but that was not trivial for me. I remember I was sweating hard going through the published code trying to adapt it to my application. I had to email Dong Nguyen (Jacob's co-author of the original HSIC application paper) in the Netherlands to understand how to apply it. Fortunately, she responded by explaining to me in a step-by-step fashion how to compute the HSIC value for a word but warned me that their implementation might be hard to scale. Following her advice, I had some progress to show Jacob, but I could not remember everything during the meeting, so I emailed him 3 questions later. In the meeting that followed, he goes "this kind of thing just does not make me wanna meet with you," shaking his head and pointing his hands at me in extreme frustration. He then answers the questions brusquely. Visibly irritated, he then asked me how things were going with DILAC, suggesting that it was an equivalent experience. I said things got better as I was shifting to work on more statistical work. He nods his head saying assertively, "If you want people to accomplish your goals, you gotta meet their goals first! The smart thing was to combine your project with the DILAC project." But the truth is, I never felt like I was more than a code monkey at DILAC to a supervisor who was not sure where to take the project. I did not identify with the 1800s American literature we worked on there. I mean, I had my entire pre-college education in public schools in Sudan where I was taught in Arabic,

and I had very limited literature training during my undergrad at JSU. And as for the other things going on at DILAC like 3D printing and visualization, I had either no interest or not enough time to develop the necessary experience. Most importantly, I didn't know if Jacob would have been open to collaboration until that point since he could only commit to a biweekly meeting with me.

It turns out I was not the only one who received this kind of treatment from Jacob. I later learned from a student that they once had an implementation issue in one of their projects and Jacob said, "This is not my fucking problem!"

Toward the time for midterms, my roommate suddenly decided to travel and find a subletter who later proved to be much worse. About a week before the NLP midterm, my new roommate didn't pay his share of the rent, and only after the first eviction notice did I notice that as I was overwhelmed by what was going on at school. My landlord only accepted money orders so I had to make a trip and pay double my share so we wouldn't get evicted. Needless to say, I had no time to study for my midterm or work on my research projects. At DILAC, I was humiliated in front of the other fellows for not having done my work. As for the DVA class, I used all my 4 slip days on the second and most challenging assignment. And as for the NLP course, I was in the bottom 5% of the class in the written midterm, but that was slightly offset by my great performance in the programming assignments. Around that time, Jacob received his tenure and announced he was going on sabbatical. Although to be fair, there was a harbinger he might be leaving. During my first semester, he mentioned that he was not taking students in the next Fall of 2018. I later learned from a Ph.D. student that Jacob had told them far in advance he was going on sabbatical the following year. But to Jacob, I was boxed as a master's student.

The rest of the semester followed the same theme of overwork. I sacrificed some commitments to rescue my grade in the NLP class. I studied a lot for the final. Luckily, I moved to the top 10% in class in the final, showing the highest improvement in class performance out of all students. Jacob even congratulated me on that achievement. As for my HSIC project, I was able to retrieve some words that were more meaningfully associated with Lyon than chance. But before he left, we had a conversation. In it, he said I was able to do something with the data but "these research projects are really self-led. We didn't get to where we needed to be [meaning the workshop submission]. You spent a week or two trying to understand the calculus of the HSIC method, and that wasn't so smart." He then moved on to ask me about my plans for the summer. Then I was in talks with Mark Riedl about joining his lab. At that time, Mark's lab was primarily focused on game AI stuff but had been running a big automated story generation project, and that was probably the most NLP thing at GT at the time. Jacob's lab was the main NLP lab, and I was not successful in joining Dhruv Batra and Devi Parikh who were

working on vision+language AI. And to be honest, I was not into the vision side of things. So, I told Jacob that I was joining Mark's lab. He responds, "Good. Just make sure you're worth his time!"

Summer 2018:

Still, in my same living situation with sporadic home visits and sleeping in the lab, I joined Mark's lab and started to work with Ph.D. student Lara Martin on collaborative story generation. The overall idea was to go from sentences to events (using WordNet and VerbNet rules), train an event-to-event model, and finally translate the events back into natural language. I mostly worked on evaluating the changes in the story corpora and VerbNet rules as well as building models and metrics. The process could be viewed as a game, and hence there was a shift towards employing deep reinforcement learning techniques.

The anxiety I developed from working in Jacob's lab didn't quite subside. In my meetings with Lara and Mark, I felt like I wasn't able to share my ideas even though Mark and Lara were both very welcoming. But I truly didn't feel like my tasks were pushing me into developing the mathematical maturity I was looking for. I wasn't sure how long I could extend my time in Mark's lab, but almost everyone there seemed happy unlike my colleagues in Jacob's lab.

With that in mind, I was encouraged by Timnit to apply for MILA internships which were all-year-round and on a rolling admission basis. I thought I was grossly underqualified, but I decided I was not going to lose much. To my surprise, my interviewer was Yoshua Bengio. I was very nervous talking to him about my research, but we had an exciting conversation nonetheless. He offered me the internship right then and there and told me I could work on whatever interests me! I then prepared for the cumbersome Canadian visa requirements.

In the same summer, I learned that some of my friends in Sudan who were just getting into AI were unable to secure funding to go to Indaba, a huge AI conference in Africa. This was also during a time when Sudan was undergoing a major revolution against the 30-year-long dictator Omer Al-Bashir. I thought it was such a good opportunity that they don't deserve to miss out on. I started a Gofundme campaign for them, and we were successful in raising around \$9,000 and sending around 10 students there. The campaign attracted the attention of the new director of the new Google's Accra office Moustapha Cisse who was also starting the African Masters of Machine Intelligence (AMMI) master's program in Kigali, Rwanda. The program featured a wide breadth of topics in ML which were taught by superstar instructors like Kyunghyun Cho and Marc Deisenroth to name a couple. Many of my Sudanese student friends were later accepted into the AMMI program and were able to secure internship or residency programs at Google, Meta AI, MILA, Sony, and others.

Do you remember my HSIC project? Well, my results were accepted with great reviews at the Black in AI workshop of 2018 which was in Montreal. And in December, I remember that Meg Mitchell helped set up my remote presentation.

Fall 2018:

My visa to Canada got rejected and I had to work out a solution with the international studies office at GT because if my internship hours were not satisfied, I would be violating my student visa requirement of staying full-time all semesters. I had to come back to campus and have some sort of emergency late registration as well as funding. Being the hero he is, Mark steps in and fulfills my enrollment requirements and offers me an RA. And this time, I was able to find a good roommate, although still far from campus.

Upon my return, Mark wanted me to continue the summer work assisting Lara and first-year Ph.D. student Prthiviraj "Raj" Ammanabrolu, but I wanted to fork off my own project. Mark and I were able to discuss doing some ideas, but we couldn't hammer down a specific project with a reasonable deadline. And no matter how hard I tried to get myself into story generation, I was fixated on the challenge of automatic evaluation and thought that working on the model was not the most pressing task. But I still dreaded going to meetings, and thankfully Mark was never pushy about anything.

As Ph.D. application season approached, I mustered the courage to gather Mark and Jacob's recommendations and apply for Ph.D. programs in the US at the top NLP labs anyways.

Spring 2019:

I'll start with my living setup. Since I wanted to get closer to campus, I was able to find one in the Homepark (sketchy but close) neighborhood. I had 3 roommates, so I could save about \$100/mo on rent and commute. I got along with 2 roommates but not the 3rd. Nonetheless, I avoided drama as much as I could. Sadly though, this was a time when I hit an intense depressive episode. I'm talking about not being able to get out of bed-, stomachache- type depression. But to fulfill my graduation requirements, I had to take the Deep Learning course with Zsolt Kira and the Machine Learning course with Charles Isbell along with my research project with Mark. Keep in mind I couldn't have had these classes earlier because there was always fierce competition around registration. At that point in the program, taking classes was a chore and a source of stress.

I sought out therapy. I started with the student counseling center at GT. The therapist I met there told me I had too much trauma to unpack and that an

external therapist would be the only way for me. I mean, it took a while to schedule that appointment and the center was understaffed. And during my time at GT there were multiple student deaths which propelled the students to advocate for mental health changes and have the president replaced.

The external therapist was never as good of a match as the GT therapist, and she once suggested I leave the field. How would I leave now that I've built a life around it?

As a small escape, I received my admission offers from the University of Maryland – College Park (UMD) and the University of Massachusetts (UMass) Amherst. I visited them and researched them extensively. UMass had an obvious focus on the computational social sciences. The students were happy and there seemed to be a strong camaraderie between students and advisors. While my direct advisor to be there was Przemek Grabowicz, I was also considering Brendan O'Connor and Mohit Iyyer. As for the CRF godfather Andrew McCallum, our conversation wasn't very successful and our interests didn't seem to converge.

At UMD, I was pretty much almost solely focused on my conversation with Hal Daumé III whom I had met during the Black in AI 2017 dinner and had had a great conversation with about the inequalities in the research community.

During the Spring 2019 semester and the few semesters before it, Hal was residing in NYC working at MSR but planning a return to campus in MD in the Fall. This move to NYC was the only complaint his students had about him, and Hal was open about his uncertainty in his future plans.

But there were also things going on at GT. Without the help of the undergrad Kyle Xiao in Mark's lab, group project leadership by the master's student Vishvak Murahari, and open submission lee-way for me from Dr. Kira, I don't think I would have been able to make it out of the Deep Learning class. As for Isbell's Machine Learning class, I practically failed, but the student counseling center got involved and I was let through. Finally, Mark accepted whatever progress I made on my independent text generation project, and I was clear to graduate.

I deferred my admission offers and graduated in May 2019 having no intention to prolong my suffering in the academy. What happened next is worth an article of its own!

Conclusion:

Coming forward with this piece has not been easy. I could never shake off my memories of that experience, so it felt only right to share it with the world.

Do I regret going to the MSCS at GT? No. But I wish I had known Mark earlier, and that I used whatever money I wasted in the Bay Area on my Atlanta housing. I'm happy that the program opened up a world of opportunities for me, but the price I paid for it in my mental health was far too much. I later learned that Jacob

didn't really want to advise students. That reframing helped. It meant that "just make sure you're worth his time" wasn't about me failing to meet a standard. It was about him resenting the ask itself. I was trying to prove my worth to someone who had already decided he didn't want to give that kind of attention to anyone.

But graduation wasn't the end. Seven months later, I was hospitalized at Grady Memorial for ten days after I stopped sleeping and started hallucinating. I reached out to Jacob during that time, and to his credit, he called. We talked for thirty minutes. It went well. Then I went manic again, and the connection I was grasping for became unsustainable. The silence that followed lasted five years. It made me invisible. To the field, to the people who shaped my trajectory, and for a long time, to myself. What came next tested me in ways the program never did. But this is a story for another time.