W. Andre Perkins

Research Engineer at the Allen Institute for Artificial Intelligence

What is your current occupation?

I am a Research Engineer at the Allen Institute for Artificial Intelligence (Ai2) working on improving weather and climate model forecasts using machine learning. My main responsibilities on the team have been provisioning and managing the cloud computing infrastructure and exploring ways to replace certain model components (i.e., physics parameterizations) with machine-learned emulators.

What is your educational background?

I got my B.S. from the University of Wisconsin – Madison majoring in both Atmospheric and Oceanic Sciences, and Computer Sciences. I then earned a Ph.D. in Atmospheric Science from the University of Washington. My dissertation was on new ways to reconstruct the last 1000 years of climate (think global maps of temperature, winds, etc.) using data assimilation, climate models, and proxy records (e.g., tree rings, ice cores, and corals).

A key message for students is that the geoscience workforce is dynamic, and boundaries between sectors and occupations are fluid. How has this been true in your career?

I have been working at the interface between my two interests of software and weather/climate since starting my undergrad, so this message resonates with me. Weather was my first love, but I was fortunate to be at a university that had computer sciences under the same college. Since there were no major admission requirements for taking CS classes at the time, I started taking introductory courses right alongside the prerequisite courses for atmospheric sciences. This

early experience ultimately led to my first job in "tech", working part-time for a traffic lab on campus helping to develop a material usage reporting system for the state's snow removal operations. From there, I had a taste for building solutions in a research-oriented domain, which was just in time for my main sequence in the data-intensive atmospheric science. I began working with a professor in the department on building climatological maps of upper-level wind events for a research group, which allowed me to start learning and using the scientific stack of Python. Around this same time, I was admitted to a summer internship program (SOARS) at the National Center for Atmospheric Research in Boulder, CO. I participated in this 10-week research experience for three summers prior to grad school doing work on various topics. The first summer I performed analysis of tree growth in simulations from the NCAR climate model. The following summer I built an orchestration tool for running ensembles of atmospheric chemistry simulations on a compute cluster, and then in a follow-on the next summer, I used that ensemble to test the methane sensing capabilities of a proposed satellite instrument.

For graduate school, I sought a project that would continue to develop my interdisciplinary skillset. I landed at the University of Washington – Seattle working with a professor on novel paleoclimate reconstruction techniques where I helped build a new Python-based framework for the project. This work involved learning basic Bayesian statistics as well as finding a way to efficiently emulate climate forecasts. After receiving my Ph.D., I entered my current position as a Machine Learning Scientist at Vulcan Inc. (now continuing at Ai2). Throughout my career, there's certainly been an aspect of luck timing-wise, but at each step when opportunities presented, I looked to explore new skills and tools while growing as a scientist. I'm grateful for all these experiences, which ultimately culminate in my current position working at the cutting edge of research in ML for climate. It has been uncomfortable at times with how much there is to learn, but it is always interesting.

Where do you see your sector moving in future years? How would you advise students to prepare to be competitive job applicants and successful employees?

There is a lot of interest in ML/AI applications in the climate space, and the government is putting its weight of funding in that direction. It's tough to say where will end up because breakthroughs can provide large leaps forward in utility and understanding, but it can sometimes take a while to achieve them. State and local governments, policy institutes, etc. will require a lot of skill and expertise from the geosciences to apply our knowledge directly into impacts and how to best mitigate these impacts, especially for vulnerable populations. We all need to be well-versed in handling data, but more important is the ability to tie that data into coherent and compelling storylines that can persuade on the urgency of our situation. For my sector specifically, the most important skills to have include experience programming in languages like Python or R (maybe Julia for those on the bleeding edge), knowing how to collaborate with others using tools like git/GitHub, and some practice in determining the right level complexity necessary for learning or modeling your problem at hand.

What is the role of networking in your sector? Do you have advice for a student who is just beginning to build their network? What is the best way for students to get their foot in the door?

Networking has not necessarily been my strong suit, but fortunately the space for machine learning in weather/climate is in its earlier stages. There are a handful of labs/groups focusing on these research topics, so find one that fits your interests and read their work, checkout their publicly released code, and reach out with questions! Often showing interest is a great way to get involved, especially when more often-than-not, there is way more research ideas than people to do the research.

What does a "typical" day of work look like for you?

In the morning I'll check through messages on Slack/e-mail and work through any administrative tasks. We have a short daily stand-up meeting around mid-morning where we check-in with each other on what we accomplished the previous day or ask for advice/feedback on a current task. Then, the rest of the day is a mix of coding, analysis, and code reviews, with the occasional meeting to discuss results or design of new features. Towards the end of the day, I'll wrap up by crossing out finished tasks and updating what I should focus on the next day. I'll note that prior to the pandemic we were in an office for every workday but will likely remain remote/hybrid going forward.

What is the best part of your job?

It's tough to pick just one thing, but I really enjoy getting to work on a technically challenging task with a direct application to problems facing the entire global population. The challenge can be stressful, but it helps to have a super-talented team to work through it all with.

Do you have any other comments or advice for students looking to enter your sector of the geoscience workforce?

I'll preface this with the acknowledgement that I have a non-standard position that is not wholly representative of the non-profit sector. First, don't be discouraged by job ads that have a huge list of "required competencies" of the applicant. You should apply even if you only feel like you meet a few of the criteria for the listed items because the whole list often describes a unicorn. Second, if you can get in touch with someone who works at a company of interest or a recruiter for an informal/informational interview, that will always be better than cold applying. Last, have a back-up plan. I find that often focusing on a singular goal (e.g., finishing a graduate degree) for a long time, it's easy to overlook the fact that your personal factors or feelings about the destination have changed. Having thought through back-ups often puts you in a place where you can take advantage of other opportunities if your current setting isn't working for you.

Connect:

https://www.linkedin.com/in/waperkins/

Learn More:

https://allenai.org/climate-modeling

https://www.geekwire.com/2020/w-andre-perkins/