

Leadership In Action AI4Good Project

Peter Ma | Aug 26th, 2022

SDG Data Catalog and DEI Storytelling

This report will be split into two major phases, one element focuses on applying current research to solve critical social problems around the world, and secondly, we investigate the storytelling element in working with communities in DEI. We will initially start our focus with the technical project before moving to the community-based project.

Sustainable Development Goals Data Catalog Project

Few-shot learning and Meta-Learning for text classification

Introduction

In the current age, humanity is faced with an increasing number of existential problems one after another. From wars to famine to inequality and climate change, we constantly find ourselves on the edge of the next big problem. These global challenges have been so profound and widespread that the United Nations made a collective effort in addressing these. We call this the Sustainable Development Goals or SDG. There are 17 of them. Addressing any one of these social problems meaningfully would take many lifetimes. From solving world hunger to ending climate change, we, unfortunately, can't do it all.

Does this mean we're doomed? I hope not. I don't think we are. Thankfully technology has advanced so rapidly in the recent decade that what once appeared to be magic, is now in the hands of everyone. One of these pieces of technology that we believe can solve some of the most pressing social problems, is data.

The Impact

The abundance of data is often seen as a double edge sword. On one hand, the surplus in information gave rise to the development of advanced algorithms with highly predictive powers and the ability to discover patterns previously unretrievable by humans. However, at the same time, we see an increase in noise, and we're flooded with a paralyzing amount of information than we have ever had before.

For researchers, engineers, and change-makers, hoping to make advancements in the 17 SDGs, data is crucial in making decisions. This is not only true for policy making but for the engineers and researchers dreaming of tackling these hard problems. Thus in our digital age, the highest leverage product that can meaningfully impact every SDG is an intelligent way of aggregating and curating high-quality data. This is the intention of the SDG Data catalog. We intend to empower decision-makers with the most valuable data to make positive social changes. If we can solve this problem, the clarity we give to policymakers and builders will help accelerate advancements in these problems.

The Project Overview

To accomplish such a task there are a number of steps in bringing the idea to life. Firstly we have web scrapers that first aggregate research papers for data. Currently, we have scraped over 8 million papers of data. To goal is to leverage this accumulation of data and build a

product whereby users can directly access specific data regarding their interests. Effectively we're building a Google search for high-quality SDG data. Currently, classical search engines are insufficient as there's no way to filter for actual datasets related specifically to SDG, hence too much noise. We're filling this gap with our product. However there is a clear problem. how do we label all the accumulated datasets and papers? We have the data, but to make it useful one would need to classify the data for people to filter through. Hence the heart of the problem I lead the solution for.

Responsibilities

There were a few responsibilities in this role. The main task is to engineer and apply the work of modern deep learning techniques to produce a product for researchers and policymakers in the SDG field. This included writing documented code that can be easily interpreted by other rest of the team. This also meant presenting weekly updates on the work done.

The Core Technical Challenge

As mentioned, there exist 8 million undocumented papers. This is useless to us because we can't match the data to the various interests of users. What we want is to give information about these papers and retrieve labels. This would naturally seem almost impossible, there is no means to manually document 8 million papers, and we have no dataset to train a classical machine learning model to automatically classify them either! This is the root cause of the problem, the lack of labels is what makes it difficult. *To solve such tasks we employ modern deep learning techniques in meta-learning to overcome the shortage of data.* Or what I like to call it, black magic!

The Technical Solution

Meta-Learning

What is meta-learning and how did we implement it? Meta-learning lets us use very few examples, maybe a couple of dozen, to train machine learning algorithms that perform as if they were trained on millions of examples. Imagine a child being able to learn the entirety of the English language from just seeing a few examples of each word. This is what we wanted to achieve.

Here is the simplified approach of how we did it.

- 1) Let's start with a few data points. This is simple, we manually read the abstract of research papers and annotate them as to where they belong in any of the 17 SDG categories.
- 2) Now we will randomly sample from the dataset a couple of examples of each.
- 3) We then train the model as usual on a few examples and we repeat!

You might wonder what is so special about this. Well here is the trick. When we update the model during the multiple learning processes, we track how the model's parameters as they're changing. For example, if you see all the artificial neurons get smaller and smaller with each iteration, we can follow this trajectory and extrapolate its movement. It's like extrapolating a line graph in excel except we extrapolate the values of the dynamics of the artificial neurons. This movement doesn't require us to train with more data, it just requires us to estimate the changes in the model. This is called Meta-Learning.

Week By Week Timeline

Here is a breakdown of the week-by-week progress. Note this contains only half of the project.

- **Week 1** Onboarding - review background info
 - This included understanding the algorithm we were implementing.
 - Read papers on meta-learning for few-shot learning
 - Read papers on transformer models
- **Week 2** Developing the data catalog classifier using NLP techniques
 - This involved actually building the first version of the Model
 - Implemented an algorithm called Reptile
 - Results didn't look good and spent time understanding the root causes
 - The training was very unstable. Likely due to learning rate issues
 - Fixed gradient problem with training BERT model
- **Week 3** Improving Models
 - Spent the week fixing the implementation and achieved workable accuracy of 85%
 - Attempted to replicate performance on other datasets but restricted due to computing resources.

Results

The final result of this work produced an MVP algorithm that can accurately predict the labels of papers in a multi-class setting with 85% accuracy. This was only trained with a few hundred data points in total instead of the 8 million. This presents itself as a leading contender in solving the core problem at heart. We now plan to use this technique in the SDG Data Catalog product. However, during this summer we did not get towards fully implementing this solution in a public-facing product. The reason is that the pace at which I developed this work exceeded that of the rest of the team as their timeline was on the order of months while mine was an order of weeks.

Smart Goals

Here I outline the SMART goals I implemented for this specific project and how the result reflected that.

- **Research & Project Management Skills**

This summer I specifically worked on Research in Project management skills. Specifically, the greatest takeaway is learning how to work in a dynamic team with varying speeds of progress. During my experience, I often outpace the other volunteers and this would often lead to distracting amounts of "down" time. Thus whenever my work depended on the answer I often anticipate two or more different scenarios and build out multiple different timelines while I await responses. In practice, this lets me instantly integrate their feedback the moment I get a response and also allows me to explore the problem in much more depth. However this resulted in more work, but this slow down put me back in pace with the rest of the

team and still produce more output! This along with other social skills in coordinating work were some of the PM& Research skills I had picked up.

SMART GOALS

- Specific
 - Develop SDG tools in aiding problem solvers in tackling SDG problems. The specifics of the research are outlined in the start of the report.
- Measurable
 - We can measure the impact by looking at the numerical performance of the model I developed. It achieved ~ 85% accuracy which is suitable for our use case
 - The model was also small in size and thus the execution time is also justifiable. On the order of a 10million parameters which is acceptable.
- Achievable
 - It is achievable that the scope of each project is small but are low-hanging fruits that make massive differences. We also demonstrated how we achieved such a goal.
- Relevant
 - As discussed previously this project directly address the larger problem of SDG!
- Time bounded
 - The schedule is also clearly suitable for my timeframe. However, in hindsight, I should have planned better around the working conditions of my coworkers.
- **Creative & Critical Thinker**

In this department, I learned how to solve problems that were previously never tackled before, much like research. Specifically, I learned how to implement ideas and engineer it in a way that works in the real world. Unlike previous research projects, things could remain theoretical, in this work, we had to merge the theory with reality and this involves creativity and critical thinking skills to accomplish such a task.

- Specific
 - Develop novel problem-solving techniques in implementing the **first** data catalog for SDG problems
 - This
- Measurable
 - Successfully building and launching a product would directly mean that the results are promising. "Promising" is determined by benchmark scores in running application tests. Measuring productivity increase is doable.
- Achievable
 - I have some knowledge of implementing these kinds of projects and believe in my ability to successfully carry this out.
 - This proved important since this gave me a headstart during the onboarding process.

- Relevant
 - Completing this project will help improve the existing pipeline in building these kinds of tools to help address the SDG data problems.

Diversity Equity and Inclusion Scorecard Project

Understanding communities and how to advocate for DEI in an objective way.

Introduction

Despite the social movements in private, public industry, and academia, there is still a long way in improving diversity, equity, and inclusion in the workforce. Although there is a lot of talk amongst executives and directors of companies about improving opportunities for everyone, much of this is still mostly talk. This isn't because of malice or negligence but a lack of easily measurable goals. Could we imagine a world where we can define a metric to gauge how well we align to our DEI goals from first principles? Only then could we make meaningful progress.

The Impact

We know that companies working in competitive environments, like our free market, are extremely efficient in optimizing every quantifiable aspect of their company ranging from KPIs to quarterly earnings to burn rates, etc. Companies are good at fitting their business model to quantifiable goals. If we are successful in making DEI measurable, we'd create yet another quantifiable goal, a metric that directly corresponds to the overall betterment of the company. The implication is that we will effectively leverage the power of free market competition in improving social goals which has rarely been done before in the DEI space.

The Project Overview

To accomplish such a task we want to create a DEI scorecard. A DEI scorecard looks at a number of metrics about a company to then publicly inform people of this score. The DEI Scorecard measures 5 categories of metrics: the talent pipeline, retention, access and advancement, representation, and historical trends. For example, the talent pipeline aims to expose potentially unfair procedures right where it all starts: when employees enter a company. It does so by looking into aspects like new hires by gender, gender bias in job postings, or educational gaps in the hiring process. The data collected serves only as hard numerical data, however, this is often not enough to form suggestions for improvement. Taking the data and formulating constructive criticism requires an element of storytelling to put lived experiences with the data to give compelling suggestions and advice. Hence the project explores the storytelling element of the data we collected.

Responsibilities

There were a few responsibilities in this role. The main task is to perform investigations as to how real people experience DEI problems in their work. This involved reaching out to people and also collecting stories online about what kinds of challenges people faced. Besides the technical requirements, I also had the usual responsibilities of organizing meetings, etc.

Weekly Tasks

- **Phase 2 DEI project**

- **Week 4:** Onboarding with DEI project + Brainstorm Ideas
 - This involved understanding the problem at hand
 - Brainstorming how to craft a story for a public-facing summary of our project
- **Week 5:** Reaching out and Interviewing people of interest that would benefit from their voices being heard
 - Making Surveys
 - Reaching out to the network for suggestions
 - Doing research on other public interviews about the problem
- **Week 6:** Wrap up interviews with reach outs [might extend longer depending on people's availability]
 - Write up the results into a few stories for different audiences.

Results

There were 3 delivered results, first, we crafted a survey for people in our network to fill out, secondly, we interviewed folks about their experiences, and lastly, we wrote stories corresponding to our DEI Scorecard project.

In-depth, our survey looked into asking questions to help voice people's stories, experiences, or concerns surrounding DE&I. The community's input is incredibly valuable to our DEI project, as we depend upon people's lived experiences to be able to communicate to the general public this issue's degree of importance. As developers and researchers, we have the tools to contextualize these results and prove the merit of DEI in the workspace, yet we lack the most important aspect of all which only the participants of our survey can provide. We ask questions like *"Assuming no limitations in money, time, or other resources, what do you wish existed that could help target DEI problems?"*

Secondly, we interviewed people of different backgrounds from people working in HR to people working in academia. We asked questions like:

1. How does one attract talent, such that a diverse group of people apply?
2. How does one then select both for the highest merit but at the same time ensure DEI goals are aligned when picking qualified candidates?
3. As a leader of a community, how do you design a culture that makes it both a comfortable and productive environment for your colleagues?
4. What is the most difficult aspect of DEI that remains unsolved? And why?

The goal is to probe people's expertise in the subject and what pinpoints our scorecard should tackle.

Lastly, we consolidated the work into 2 articles of writing explaining the importance of DEI and how our scorecard can help solve these problems. These are still technically in the drafts as we've yet to fully launch the product during my time working with the group.

More SMART goals

- **Social & Cultural Intelligence & Awareness**

This work really developed my skills in social and cultural awareness. I was working on developing storytelling techniques to promote gender equality and creating messages to directly help improve DEI discussion and this is completely foreign to me. I often interact with experts in the technical sense and rarely do I simply ask how they “feel” about a certain topic. Understanding emotions and how to convey them effectively is something I rarely find myself doing in my professional career and thus this was an excellent experience for me in that sense.

- Specific
 - Reach out to marginalized groups and have discussions about their experience and being vulnerable to the issues in society. Done!
- Measurable
 - Having 3-5 compelling stories to share would be the objective and is definitely measurable. Done!
- Achievable
 - I have reached out to people before and I am well aware of ideal techniques for getting people’s attention in terms of cold outreaches! Done!
- Relevant
 - Interviewing people and asking for people’s stories and input would drastically help improve the discussion and how we build the DEI Scorecard! Done!

Overall Takeaways

This was my first opportunity to work on a climate action project which is a complete 180-degree turn from my previous experiences. Previously I’ve only worked on pure research. Now I am directly thinking of ways to use this to help address SDG problems. The environment is challenging because the work is almost 100% new to me; the only similar aspect is that there will be aspects of coding in building one half of the project. Secondly, I was working on developing storytelling techniques for the DEI scorecard project which helped me explore the emotional and social implications of my work. Lastly, even when working remotely I found myself collaborating with folks overseas and all though I am not in a physically different location, I felt like I was in a new environment working with people of different backgrounds!

Acknowledgments

This project was funded through LAIDLAW FOUNDATION and has supported me financially, and guided me through producing a truly unique experience! Furthermore, I’d like to acknowledge AI4Good for mentoring and supporting the work I’ve done with the group!