Date: May 9, 2022
Speaker: Shaunak Chatterjee
Venue: US Chamber of Commerce deposition, Palo Alto, CA

Why should we care about this?
AI and technology play increasingly important roles in our lives. For instance, our mobile phones are becoming critical to an ever increasing number of tasks in our everyday lives. And AI is powering an ever increasing number of suggestions made by our phones.
With this increasing influence, comes greater need to ensure that AI is being used in a responsible manner.

Why is this important to LinkedIn?
LinkedIn is a platform with 830M+ users globally, who use the platform to build and maintain their professional identity and network, to look for jobs, to hire and to engage with content. All of these user needs are met with significant AI usage.

LinkedIn's mission is to create economic opportunity for every member of the global workforce. In order to achieve this mission, and given the pervasive use of AI, we need to take special care to understand and ensure that the AI technology is being built and used responsibly.

LinkedIn, as a part of Microsoft, is also building towards Microsoft's 6 principles of Responsible AI: Fairness, Reliability and Safety, Inclusiveness, Privacy and Security, Accountability and Transparency.

Fairness @ LinkedIn
For today's conversation, I will focus on Fairness. Fairness, according to the dictionary, means "impartial and just treatment or behavior without favoritism or discrimination."

At LinkedIn, we think of our work on Fairness through an Intent and Impact framework.

During the design and implementation of our products and algorithms, we intentionally attempt to build equitable experiences for all our members. This translates into specific steps like conducting inclusive user surveys, evaluating our training data for representation, checking model performance before deployment and designing specific changes to our products and offerings.

However, having the right intent is not enough. We need to understand the impact of our product offerings on our users. This involves detecting and monitoring how people interact with changes after they are deployed. By focusing on outcomes and the impact on people, we can better ensure that we are delivering on our original intent.

Putting it into practice
Let's use a concrete example to ground some of the ideas I have talked about thus far. The People You May Know or PYMK product is one that LinkedIn pioneered and shows suggestions

to our members on who they can invite to connect with and grow their network. The exact sequence of steps is as follows:

- A user, Daniela, is shown several recommendations of people they may know and may want to connect with.
- Daniela then decides to invite Imani to connect with her. Imani was one of the recommendations shown to her by PYMK.
- Imani receives the invitation to connect from Daniela, and can choose to either accept or decline that invitation.
- If Imani accepts the invitation, then Daniela and Imani become connected to each other and part of each other's network.

A user's network leads to several important downstream effects on LinkedIn. For instance, based on data from 2019-2020, a job applicant was 4x more likely to be hired in a company where they have connections. Also, more than 40% of all connections on LinkedIn were initiated via PYMK.

Assessing fairness in PYMK

In order to better understand how PYMK is being used by different groups of members, we look at the product usage funnel. In PYMK, the different steps of the funnel include:

1. Candidates generated by the PYMK algorithm
2. The impressions of candidates in the PYMK product
3. The number of invitations sent to the impressed candidates
4. The number of connections formed by the invited candidates

We look at the representation number of each group of members (grouped by various member attributes) and use the previous funnel step as a reference distribution to understand how the different groups fared and behaved through different phases of the product. Such detailed analyses helps us identify which portions of the PYMK funnel are leading to larger drop-offs for a particular group of members and accordingly plan for improvements.

Algorithmic fairness in PYMK

Another notion of fairness that we strive for is algorithmic fairness, where we wish to ensure **equal opportunity for equally qualified members**. This translates into the PYMK product in the following way: opportunity is equivalent to exposure or being impressed on PYMK as a candidate. Qualification is the outcome of getting invited and accepting the invitation.

There are various statistical tests in the fairness literature to test for algorithmic fairness given the objective you have in mind. We are looking at three different approaches in our algorithmic fairness work:

1. Pre-training: An example would be adjusting the weights of training data points. The general objective is to improve disparity in model performance for various groups.
2. In-training: An example would be adding appropriate regularization functions. This can also help in model performance improvement and passively improve fairness measures.

3. Post-training: A score booster or reranker would be salient examples here. This is the most direct option to enforce a chosen definition of fairness. The primary drawback of this approach is that it may not improve model performance.

Continuous monitoring

The last avenue of our fairness work, and a key part of our impact focus in the “Intent and Impact” framework is continuous monitoring of our systems, especially changes that we deploy. We use randomized testing (or A/B tests) to understand both the intended and (especially the) unintended impact of our changes on different user cohorts.

It is great to see the awareness and amount of active research and thinking happening in industry and academia to make the world, in general, and our tech products and platforms, in particular, more inclusive and fair. We are excited to play our part in this important mission.