Applicant: Azal Ahmad Khan Web: *azalahmadkhan.github.io* Email: azalahmadkhan@gmail.com

Statement of Purpose

Beginning on a Ph.D. journey, my focus lies on tackling crucial challenges hindering the widespread adoption of **Foundation Models** (FMs). I am interested in two fundamental aspects related to FMs. Firstly, the development of **FMs that align with human preferences and ethical standards**, ensuring their resonance with human values. Secondly, my exploration extends to the practical **integration of FMs in Federated Learning** (FL) settings, recognizing the crucial role of edge-generated data and the necessity to train FMs in these distributed environments.

Research Experience

Federated Learning: I began my research journey during my sophomore year at DMLSys UMN, focusing on Personalized FL (PFL) under the guidance of Dr. Ali Anwar. In observing the efficacy of FL, our attention was drawn to a persistent challenge: *how to motivate clients to share local updates, given the resource-intensive nature of training models in heterogeneous networks?* To tackle these challenges, we introduced the PI-FL framework (1), integrating tiering-based personalization with a token-based incentive mechanism, traditionally treated as separate issues. This not only empowered clients to align models with their unique data but also motivated active participation, ensuring the creation of high-quality personalized models while safeguarding data privacy. Under review at **AAAI 2024**, this work represents my significant step toward democratizing the benefits of FL.

My previous project's success motivated me to delve further into FL. *I tackled issues like stragglers and dropouts to enhance model performance, particularly on resource-constrained devices like IoT and mobile gadgets*. Addressing declining performance in current client selection methods, I introduced FLOAT (2), a framework optimizing both model performance and resource efficiency. FLOAT leverages a multi-objective Reinforcement Learning with Human Feedback (RLHF) agent, ensuring a balanced use of resources without compromising training objectives. I developed the RL framework, implemented compression techniques, and contributed to a paper under review for **ACM EuroSys 2024**.

Furthermore, initiating a response to a critical gap in PFL literature, I led the survey paper (3) currently under review at **IEEE TPDS 2024**. My focus identified and addressed two key issues: the *lack of empirical analyses across PFL algorithms and the literature guiding optimal PFL algorithms in real-world scenarios*. I conducted empirical analyses on 10 PFL algorithms, comparing them against two baselines and extracting valuable insights into metrics such as performance, time, and communication cost. This research stems from my commitment to advancing scholarly discourse on PFL.

Continuing my tenure at UMN, I actively concentrate on optimizing GPU usage for Large Language Model (LLM) training. I am dedicated to overcoming resource limitations for faster LLM inference. My focus lies in the development of efficient systems through techniques such as offloading and compression strategies, particularly in the context of single GPU training.

Human-aligned FMs: Taking on the challenge of a unique project amidst the hype around LLMs and RLHF, I focused my Bachelor's thesis on improving text-to-image generation through prompt engineering. Addressing challenges presented by RL (4) and RLHF frameworks, which often resulted in *misaligned images and computational overhead due to separate training of reward models and RL policies*, I proposed the adoption of Direct Preference Optimization (DPO) (5) to align Language Models (LMs) with human preferences.

Our framework (6), efficiently produces human-aligned and visually appealing images. Motivated by these outcomes, I conducted comprehensive in-domain and out-domain prompt tests, revealing comparable and promising results even with a modest training dataset. Excited by the potential, I am eager to submit our work to ICML 2024.

My experiences have led to questions that broadly align with two research themes I intend to explore during my Ph.D.

Future Directions

Human-aligned FMs: FMs often generate outputs that are not well-aligned with human values or intentions, which can have unintended or negative consequences.

RQ 1: How can we enhance the human alignment of foundation models to apply them in real-world scenarios effectively? My undergraduate thesis (6), focused on aligning FMs with human preferences in text-to-image scenarios, has motivated me to advance human alignment in AI applications. Aspiring to pursue a Ph.D., I seek to extend this exploration to align LMs for applications such as providing factually accurate responses, better reasoning, and improving text summarization.

RQ 2: How can the cost of human-alignment processes for FMs be minimized without compromising performance? Recognizing the significant impact of resource-intensive policy training, I aspire to advance the study and refinement of human-alignment techniques for FMs. This involves a particular focus on exploring alternatives to RLHF, such as DPO (5) and IPO (7), contributing to the ongoing enhancement of training processes and gradually improving their efficiency and performance.

Integrating FMs in FL: The integration of FMs into FL is a promising and anticipated future trend. Data generated in edge devices is crucial for training FMs, but the privacy concerns and communication expenses, highlight the need to apply FMs directly on the devices themselves. However, challenges, such as high computational resource demands in edge devices, prompt important questions (8).

RQ 3: How can we improve the scalability and efficiency of FL systems to handle large-scale and complex datasets? My interest lies in addressing the challenge of leveraging FMs for large-scale datasets on resource-constrained edge devices. I propose exploring parallelism techniques in models and data to distribute computational load and enhance training speed. Simultaneously, I aim to contribute to the field by developing novel FL algorithms, specifically in areas like federated reinforcement learning and federated meta-learning, to handle complex datasets.

RQ 4: How can we efficiently update and distribute FMs across a network while managing and balancing the computational load among devices? I'm eager to tackle the complex issue of managing computational load in FL. My approach involves employing hierarchical or decentralized architectures and organizing devices into clusters based on their capabilities and data characteristics, similar to (1). This strategy aims to balance the load and reduce communication overhead by enabling independent training and updation of FMs in each sub-network. Furthermore, I plan to design FL algorithms that dynamically adapt computational load and communication frequency based on device availability and performance. My previous work (2) serves as motivation for these endeavors.

Why Ph.D. at UMN?

Why Ph.D.? In the past two years, my research journey has been deeply fulfilling, sparking a passion for delving deeper into my field. The realization of this passion motivates my pursuit of a Ph.D., where a graduate school environment offers essential resources, guidance, and a collaborative community. Motivated by a desire to advance the applications of FMs in human alignment and FL settings, my dedication is focused on leveraging their potential for efficiency and fairness in real-world scenarios. This commitment propels my aspiration to make significant research contributions, making a Ph.D. the essential pathway to transform these aspirations into impactful realities.

Why at UMN? I believe the UMN CSE program is a great fit for me. I look forward to working with [redacted] on the applications of LLMs and the human alignment of LLMs. I also look forward to continuing my collaboration with **Prof. Ali Anwar** on the integration of FMs in FL. Lastly, I highly value the collaborative and inclusive environment of the UMN CSE program which I believe will help me expand my research perspective and conduct highly impactful research.

References

[1] Ahmad Faraz Khan, Xinran Wang, Qi Le, **Azal Ahmad Khan**, Haider Ali, Jie Ding, Ali Anwar, Ali Butt. PI-FL: Personalized and Incentivized Federated Learning. arXiv preprint arXiv:2304.07514.

[2] Ahmad Faraz Khan, **Azal Ahmad Khan**, Samuel Fountain, Ahmed M. Abdelmoniem, Ali Butt, Ali Anwar. FLOAT: Federated Learning Optimizations with Automated Tunnings

[3] **Azal Ahmad Khan**, Ahmad Faraz Khan, Haider Ali, Ali Anwar. Personalized Federated Learning Techniques: Empirical Analysis.

[4] Yaru Hao, Zewen Chi, Li Dong, Furu Wei. Optimizing prompts for text-to-image generation. Advances in Neural Information Processing Systems (NeurIPS), 2023.

[5] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model.

[6] **Azal Ahmad Khan**, Debanga Raj Neog. Direct Preference Optimization in Prompt Engineering for Text-to-Image Synthesis.

[7] Tianqi Liu, Yao Zhao, Rishabh Joshi, Misha Khalman, Mohammad Saleh, Peter J. Liu, Jialu Liu. Statistical rejection sampling improves preference optimization. arXiv preprint arXiv:2309.06657.

[8] Weiming Zhuang, Chen Chen, Lingjuan Lyu. When foundation model meets federated learning: Motivations, challenges, and future directions. arXiv preprint arXiv:2306.15546.