

# Shehzaad Dhuliawala

---

**BIO**  
I am a PhD student at ETH Zürich. Previously, I spent two years as a Research Engineer at Microsoft Research Montréal. I am interested in building reasoning systems that are explainable, trustable, and robust to distributional shifts. Broadly, I work on machine learning and study its application in various fields of natural language processing.

[!\[\]\(919a2cb85b99741a73c0c31a427236a8\_img.jpg\)](#) [!\[\]\(c9cd5a1c35167a83f09a35036fe5dcbd\_img.jpg\)](#) [!\[\]\(ae1936640fabdea8c18f922ca69733fe\_img.jpg\)](#)

## **EDUCATION** **ETH Zürich**

PhD., Computer Science, 2020-?

**Advisor:** Prof. Mrinmaya Sachan, Prof. Thomas Hofmann

## **University of Massachusetts Amherst**

M.S., Computer Science, 2016-2018

**Advisor:** Prof. Andrew McCallum

**Courses:** *Machine Learning, Deep Learning, Adv. Natural Language Processing, Reinforcement Learning, Adv. Algorithms, Logic, Software Engineering*

## **Visvesvaraya National Institute of Technology,**

B.Tech., Computer Science and Engineering

## **AWARDS** **IBM PhD Fellowship Award 2021**

**RESEARCH & PROFESSIONAL EXPERIENCE**  
**Research Intern** June 2023 - Sept 2023  
**FAIR, Meta AI, New York**  
**Mentor:** Jason Weston

-Investigated how LLM inconsistencies can be used for detecting hallucinations using verification questions

-Developed the [Chain-of-Verification](#) approach where the LLM uses verification questions to detect and mitigate hallucinations in longform generations.

**Research Engineer** June 2018 - Oct 2020

## **Microsoft Research Montréal**

-Worked on annotating and creating question-answering datasets from product manuals

-Researched novel methods to train a system to answer user queries against product manuals using limited amounts of annotated data

**Data Science Research Intern**

June 2017 - Aug 2017

## **@WalmartLabs,**

Sunnyvale, CA

-Experimented with methods to incorporate local weather information to improve item recommendations for Walmart e-commerce

**A Diachronic Perspective on User Trust in AI under Uncertainty**  
(EMNLP 2023)*Shehzaad Dhuliawala, Vile'm Zouhar, Mennatallah El-Assady, and Mrinmaya Sachan.***Variational Classification**

(TMLR 2023)

*Shehzaad Dhuliawala, Mrinmaya Sachan, Carl Allen***Calibration of Machine Reading systems at Scale**

(ACL 2022 findings)

*Shehzaad Dhuliawala, Leonard Adolphs, Rajarshi Das, and Mrinmaya Sachan.***Case-based Reasoning for Better Generalization in Text-Adventure Games**

(ICLR 2022)

*Mattia Atzeni, Shehzaad Dhuliawala, Keerthi Muruganesan, Mrinmaya Sachan***TopiOCQA: Open-domain Conversational Question Answering with Topic Switching**

(TACL 2021)

*Vaibhav Adhlaka, Shehzaad Dhuliawala, Kaheer Suleman, Harm de Vries, Siva Reddy***How to Query Language Models?***Leonard Adolphs, Shehzaad Dhuliawala, and Thomas Hofmann*

Preprint 2021

**A Simple Approach to Case-Based Reasoning in Knowledge Bases**

(Accepted at AKBC 2020, Best paper runner-up)

*Rajarshi Das, Ameya Godbole, Shehzaad Dhuliawala, Manzil Zaheer, Andrew McCallum***Multi-step Retriever-Reader Interaction for Scalable Open-domain Question Answering**

(ICLR 2019)

*Rajarshi Das, Shehzaad Dhuliawala, Manzil Zaheer, Andrew McCallum***Towards Domain Adaptation from Limited Data for Question Answering**  
(ASRU 2019 Life-long Learning Workshop )*T.J Hazen, Shehzaad Dhuliawala, Daniel Boies***Go for a Walk and Arrive at the Answer: Reasoning Over Paths in Knowledge Bases using Reinforcement Learning**

(Accepted at ICLR 2018, Won best paper at AKBC at Neurips 2017)

*Shehzaad Dhuliawala\*, Rajarshi Das\*, Manzil Zaheer, Luke Vilnis, Ishan Durugkar, Akshay Krishnamurthy, Alex Smola, Andrew McCallum*