

## Education

- **Carnegie Mellon University** Pittsburgh, PA  
*Master of Science in Machine Learning* Sep 2018 – Dec 2019
  - Cumulative GPA: 4.09/4.33
  - Selected coursework: Deep Learning, Advanced ML, Probabilistic Graphical Models, Convex Optimization
- **Yandex Data School** Moscow, Russia  
*Master's equivalent program by Yandex* Sep 2017 – May 2020
  - Cumulative GPA: 5.0/5.0
  - Selected coursework: Algorithms and Data Structures, Parallel and Distributed Computing, Reinforcement Learning
- **Higher School of Economics** Moscow, Russia  
*Bachelor of Science in Mathematics, with distinction* Sep 2014 – Jul 2018
  - Cumulative GPA: 9.7/10. (top 4 in class)

## Work Experience

- **Twitter, Inc.** San Francisco, CA  
*Software Engineering Intern, Twitter Search* June – August 2019
  - Implemented knowledge distillation training for Search ranking models, performed offline and online experiments and analysis, and deployed the end model to production.
  - Used Java for data generation, Scala+scalding for analysis, Python+tensorflow for model implementation.
  - New model brings daily +1.55M tweet clicks and +48K social actions (favorites, retweets, etc.) on search results.
- **Carnegie Mellon University, Machine Learning Department** Pittsburgh, PA  
*Graduate Research Assistant, focusing on Bayesian optimization* Sep 2018 – Present
- **Higher School of Economics, CS department** Moscow, Russia  
*Teaching assistant for Linear and Abstract Algebra course* Sep – Dec 2016

## Selected Projects

- **ChemBO: joint optimization & synthesis of organic molecules** Spring 2019  
*Research project at CMU* [github](#), [preprint](#)
  - Developed an algorithm for molecular optimization with synthesizable recommendations that performs state-of-the art optimization with less queries than alternatives on common benchmarks.
  - Implemented a surrounding framework for custom objectives, exploration strategies and domains, and integrated with existing Bayesian optimization package Dragonfly.
- **Graph neural networks as an inference engine** Spring 2019  
*Course project @ CMU 10-708 Probabilistic Graphical Models* [github](#), [poster](#)
  - Experimented with scaling up inference in probabilistic graphical models with graph-based neural nets.
  - Designed the experiment pipeline and led the team to work on a 1/2 semester time frame.
- **End-to-end speech recognition system** November 2018  
*Homework project @ CMU 11-785 Deep Learning* [gitlab](#)
  - Implemented a system for generating speech transcriptions from audio with attention, input-output masking and beam decoding. Levenshtein distance 15 on WSJ dataset.
- **Some other projects on:** coresets compression for neural networks; VC theory and sample compression schemes; graph neural networks: joint node-edge embeddings, long audio classification with time graph; kaggle competitions.

## Skills

**Experience in:** Deep Learning R&D, Applied ML competitions, Optimization, Probability & Statistics, Algorithms  
**Programming languages:** Python (daily use), C++, Java (learning)  
**Technologies and tools:** Pytorch, Edward, CUDA C, git/github, SQL, L<sup>A</sup>T<sub>E</sub>X

## Awards & Distinctions

**2015, 2016, 2017 Increased academic scholarship** from Higher School of Economics  
**2014 Talented Youth Award** of President of Russia  
**2013-2014 Lomonosov Olympiad in Mathematics** 1<sup>st</sup> Diploma, absolute winner