

# Dong He

▪ donghe@cs.washington.edu ▪ dongheuw.github.io ▪ (206) 295-6340

## EDUCATION

- University of Washington, Ph.D. in Computer Science** USA  
▪ Advisor: Prof. Magdalena Balazinska Sep 2019 – present
- Fudan University, B.Sc. in Computer Science (Honors)** China  
▪ Cumulative GPA: 3.6 / 4.0, School Rank: 6 / 118 Sep 2015 – Jul 2019
- University of Birmingham, Exchange Undergraduate Student** UK  
▪ First Class Honors' Grades Sep 2017 – Dec 2017

## RESEARCH EXPERIENCE

- Accelerating Declarative Top-K Queries for Deep Neural Network Interpretation** UW  
Advisor: Prof. Magdalena Balazinska Sep 2019 – Apr 2021
- Designed, implemented, and evaluated DeepEverest, a system for the efficient execution of *interpretation by example* queries over the activation values of a deep neural network.
  - DeepEverest consists of an efficient indexing technique and a query execution algorithm that has a strong theoretical guarantee, as well as various optimizations.
  - Experiments with our prototype implementation show that DeepEverest, using less than 20% of the storage of full materialization, significantly accelerates individual queries by up to 63× and consistently outperforms other methods on multi-query workloads that simulate DNN interpretation processes.
- The VisualWorld Video Data Management Project** UW  
with the VisualWorld team Sep 2019 – present
- **TASM**: a tile-based storage manager for video data which enables spatial random access into encoded videos. TASM speeds up content retrieval queries by an average of over 50% and up to 94%, and also improves the throughput of the full scan phase of object detection queries by up to 2x.
  - **VFS**: a system that decouples application design from video data's physical layout and compression optimizations. This decoupling allows application and system developers to focus on their relevant functionality, while VFS handles the low-level details associated with video data persistence. VFS also improves read performance by up to 54%, and reduces storage costs by up to 45%.
- FPGA-Based Edge Computing for the Acceleration of Mobile Applications** Peking University  
Advisor: Prof. Chenren Xu Jul 2017 – Aug 2017
- Designed an FPGA-based edge computing model, which can effectively reduce the response time and energy consumption of interactive mobile applications.
  - Implemented a proof-of-concept prototype, and conducted experiments in a case study using 3 computer vision-based interactive applications designed by us.
  - Experimental results showed that our system can reduce the response time and execution time by up to 3×/15× respectively over CPU-based edge/cloud offloading and achieve up to 29.5%/16.2% improvement on energy efficiency on mobile device/edge nodes respectively.
- Improving the Prediction of Real-Time Taxi Demand with External Information** Fudan University  
Advisor: Prof. Yang Chen Sep 2018 – Jan 2019
- Proposed a deep learning-based approach which incorporates user check-in data from a Location-Based Social Network to improve the prediction of the taxi demand in different regions at different times.
  - Integrated the taxi trip data with around 1 million user check-ins collected from the Swarm App. Evaluation on a dataset containing 35 million taxi trip records showed that our method achieves 21.27% lower MAPE and 6.96% lower RMSE compared to existing approaches.

## PUBLICATIONS

- DeepEverest: Accelerating Declarative Top-K Queries for Deep Neural Network Interpretation** [Paper] [Technical Report] [Website] [Code]
- **D. He**, M. Daum, W. Cai, M. Balazinska
  - VLDB 2022

**VSS: A Storage System for Video Analytics** [Paper] [Technical Report] [Website] [Code]

- B. Haynes, M. Daum, **D. He**, A. Mazumdar, M. Balazinska, A. Cheung, L. Ceze
- SIGMOD 2021

**TASM: A Tile-Based Storage Manager for Video Analytics** [Paper] [Preprint] [Website] [Code]

- M. Daum, B. Haynes, **D. He**, A. Mazumdar, M. Balazinska
- ICDE 2021

**Accelerating Mobile Applications at the Network Edge with Software-Programmable FPGAs** [PDF]

- S. Jiang, **D. He**, C. Yang, C. Xu, G. Luo, Y. Chen, Y. Liu, J. Jiang
- INFOCOM 2018

**Incorporating Location Based Social Networks in the Prediction of Real-Time Taxi Demand with Deep Learning** [PDF]

- **D. He**, Y. Chen
- CoNEXT 2018, Poster Session

**WORK  
EXPERIENCE**

**Microsoft, Research Intern**

Remote

- With Jim Gray Systems Lab, mentored by Matteo Interlandi.
- Worked in the intersection of systems and machine learning.

Jun 2021 – Sep 2021

**Goldman Sachs, Technology Summer Analyst**

Hong Kong

- With the Product Accounting and Risk Analysis team.
- Global Winner for Goldman Sachs 2018 Intern Engineering Challenge.
- Re-designed and re-implemented the logic of the True-Up job which reconciles the estimated PnL (profit and loss) with the actual PnL. My enhancements are in production.

Jul 2018 – Sep 2018

**Tencent, Engineering Intern**

Shenzhen

- With YouTu Lab led by Prof. Jiaya Jia and Prof. Yu-Wing Tai.
- Analyzed the liveness and dependencies of the nodes in the neural networks, and reduced the memory consumption of such models in real-world products by memory sharing.
- Developed tools for the collection and annotation of large-scale image data, and collected massive image data for the training of image classification models in real-world products.

Jan 2018 – Feb 2018

**SELECTED  
AWARDS**

- Paul G. Allen Fellowship, Paul G. Allen School of Computer Science & Engineering, UW 2019
- Outstanding Undergraduate Graduates, Shanghai Region 2019
- Wangdao Scholar, Undergraduate Research Opportunities Program, Fudan University 2018
- First Class Scholarship, Fudan University 2016 – 2017
- First Prize, the National Mathematical Contest in Modeling, Shanghai Division 2016
- Silver Medal, the ACM International Collegiate Programming Contest, Asia Regional 2015
- Silver Medal, National Olympiad in Informatics, National Finals 2014
- First Prize, National Olympiad in Informatics, Guangdong Division 2009 – 2014

**PROFESSIONAL  
SKILLS**

- **Programming Languages:** C/C++, Java, Python, ...
- **Deep Learning Libraries:** PyTorch, Tensorflow, Keras, ...
- **Others:** SQL, LaTeX, Git, SVN, Gnuplot, ...