

HYUNSU CHO

chohyu01@cs.washington.edu • (206) 453-8718 • <https://hyunsu-cho.io>

Experience

- **Lead Maintainer** of XGBoost project. 2016–Present.
XGBoost is an efficient, scalable library for gradient boosting. Available at <https://github.com/dmlc/xgboost>. Over the years, I made substantial contribution to the C++/Python portion of XGBoost. I also built and have maintained a continuous integration server for XGBoost. The server is hosted on AWS and runs a variety of tests for all incoming code contributions, ensuring that XGBoost works well with multiple operating systems, data infrastructures (Dask, Apache Spark) and NVIDIA GPUs. In addition, I made multiple renovations for the build and packaging system used in XGBoost, ensuring that XGBoost works well with other packages in Python and C++ ecosystem.
- **Senior Systems Software Engineer**, NVIDIA. April 2020–Present.
I am part of the RAPIDS team, who builds [RAPIDS](#), the cutting-edge software suite for accelerating data science and machine learning on NVIDIA GPUs. My achievements include: build an optimize runtime Triton-FIL for inferencing with decision trees; identify performance bottlenecks in our random forests; build end-to-end examples of deploying RAPIDS and XGBoost on the cloud. I work with some of our corporate clients, delivering new and improved features to meet their needs.
- **Applied Scientist**, Amazon Web Services. January 2018–February 2020.
I was part of initial efforts to launch [Amazon SageMaker Neo service](#) in November 2018. Neo is a service that lets users optimize machine learning models and deploy to the cloud and various edge devices. My intern project, Treelite, (see below) was crucial in expanding the scope of Neo to include decision tree models, where originally it only aimed to support neural nets. I helped a number of internal and external customers to deploy and speed up their machine learning models.
- **Applied Scientist Intern**, Amazon Web Services. June–December 2017.
I created the Treelite project, a model compiler/optimizer for decision tree ensembles.
- **Teaching Assistant**, Machine Learning Specialization at Coursera. 2015–2017.

Skills/Proficiencies: Modern C++, Python, CMake, Docker, Continuous integration and delivery (CI/CD)

Education

- **M.S. in Computer Science and Engineering**, University of Washington. September 2015 – March 2018.
- **B.S. in Computer Science and B.S. in Mathematics**, Trinity College. September 2011 – May 2015.

Awards and Honors

- **Best Short Paper Award**, ACM SIGIR 2019, Paris, France, July 25, 2019.
- **Outstanding Undergraduate Researcher Award**, Computing Research Association, 2015.
- **Recipient of the Goldwater Scholarship**, The United States Congress, 2014.

Selected Peer-Reviewed Publications

- Avinash Barnwal, Hyunsu Cho, and Toby Hocking. “Survival Regression with Accelerated Failure Time Model in XGBoost,” *Journal of Computational and Graphical Statistics*, 31:4, p. 1292-1302, May 24, 2022.
- Theodore Vasiloudis, Hyunsu Cho, and Henrik Boström. “Block-distributed Gradient Boosted Trees,” *ACM SIGIR 2019*, Paris, France, July 25, 2019. Best Short Paper Award.