

HYUNSU CHO

Box 352350, 185 Stevens Way, Seattle, WA 98195, USA

chohyu01@cs.washington.edu • (206) 453-8718 • <https://github.com/hcho3>

Research Interest

Machine learning systems

Education

Ph.D. in Computer Science and Engineering (in progress), University of Washington, Seattle, WA

- September 2015 – Present
- Advisor: [Carlos Guestrin](#)

B.S. in Computer Science and B.S. in Mathematics, Trinity College, Hartford, CT

- September 2011 – May 2015

Experience

- **Applied Scientist Intern**, AWS Deep Learning, Amazon. June–September 2017.
- **Contributor**, XGBoost project. 2016–Present.
I have submitted a major addition to **XGBoost**, featuring faster training time and flexible tree-growing strategy. XGBoost is a widely adopted machine learning package for scalable **gradient boosting**. Source code available at <https://github.com/dmlc/xgboost/>.
- **Teaching Assistant**, Machine Learning Specialization at Coursera. 2015–2017.
Created interactive programming assignments using **Jupyter notebooks**. The online course series aims to provide an accessible introduction to the field of machine learning. We make extensive use of **GraphLab Create** to support a case study approach. See summary at <https://www.coursera.org/specializations/machine-learning>. The notebooks are available at <https://github.com/learnml/machine-learning-specialization>.
- **Undergraduate Researcher**, Trinity College. 2012–2015.

Awards and Honors

- **Winner of Outstanding Undergraduate Researcher Award**, Computing Research Association, 2015.
- **Recipient of the Goldwater Scholarship**, The United States Congress, 2014.

Peer-Reviewed Publications

- Lin Cheng, Hyunsu Cho, and Peter Yoon. “An Accelerated Procedure for Hypergraph Coarsening on the GPU,” *IEEE High Performance Extreme Computing Conference*, Waltham, MA, September 16, 2015.
- Hyunsu Cho and Peter Yoon. “A Memory-Efficient Algorithm for Large-Scale Symmetric Tridiagonal Eigenvalue Problem on Multi-GPU Systems,” *Proceedings of the 2014 International Conference on Parallel and Distributed Processing Techniques and Applications*, pp. 568-573, Las Vegas, NV, July 24, 2014.
- Lin Cheng, Hyunsu Cho, and Peter Yoon. “GPU Accelerated Vessel Segmentation Using Laplacian Eigenmaps,” *Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Networks*, pp. 177-184, Innsbruck, Austria, February 17, 2014.
- Lin Cheng, Hyunsu Cho, Peter Yoon, and Jiajia Zhao. “An Efficient Out-of-Core Implementation of Block Cholesky Decomposition on a Multi-GPU System,” *Proceedings of the 24th IASTED International Conference on Parallel and Distributed Computing and Systems*, Las Vegas, NV, November 13, 2012. Best Paper Award.