

Date	Paper	Presenter 1	Presenter 2
14-Sep	Support-Vector Networks	Sarah Kim	Ryan Spring
14-Sep	Random Projection, Margins, Kernels, and Feature-Selection	Adithya Kumar Pediredla	
21-Sep	Very Sparse Random Projections	Chris Metzler	Jeff Lieveuse
21-Sep	Approximate Nearest Neighbors and the Fast Johnson-Lindenstrauss Transform	Jayvee Abella	Jaehee Park
28-Sep	Hash Kernels	David Sullivan	
28-Sep	Hashing Algorithms for Large Scale Learning	Ryan Lear	Betty Chen
5-Oct	Densifying One Permutation Hashing via Rotation for Fast Near Neighbor Search	Runhua Zhang	Derek mikesell
5-Oct	Improved Densification of One Permutation Hashing	Scott Morgan	
26-Oct	Randomized Algorithms and NLP: Using Locality Sensitive Hash Functions for High Speed Noun Clustering	Mohamed "Hafez" AbdElRahman	Chen Luo
26-Oct	Fast, Accurate Detection of 100,000 Object Classes on a Single Machine	Souptik Barua	Xihao Wang
2-Nov	On the difficulty of training Recurrent Neural Networks	Tan Nguyen	Yao-Hsiang Yang
2-Nov	Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift	Ewa Nowara	Yusong Shen
16-Nov	Latent Dirichlet Allocation	Mengchen (Maggie) Tang	
16-Nov	Stochastic Variational Inference	Risa Myers	
23-Nov	An Architecture for Parallel Topic Models	Ankush Mandal	Sadegh Riazi

23-Nov Reducing the Sampling  
Complexity of Topic Models Yulia Baker