



STATISTICS
COLORADO STATE UNIVERSITY

Gradient-Flow Optimization as Dynamic Random-Effects Inference with Applications to Deep Learning

Spring 2026 Departmental Seminar

Thursday, Sept 3, 2026

9:00 am

Engineering Rm E105

Zhonghua Liu

Assistant Professor, Biostatistics
Columbia University



Abstract

Gradient-flow optimization is usually viewed as an algorithmic procedure for minimizing empirical loss, with training duration selected by validation or heuristic early stopping rules. We develop a statistical inference framework for gradient-flow training. We show that whenever fitted values evolve through a time-invariant positive semidefinite training operator, the output at each time is equivalent to the best linear unbiased predictor under a corresponding random-effects model. Training time then becomes a variance-component parameter governing variance reallocation from residual noise to structured signal. This turns two training decisions into inferential problems: whether training is needed becomes a variance-component test for signal beyond initialization, and how long to train becomes restricted maximum likelihood (REML) estimation of the training-time variance component. We show that the REML-guided early stopping rule selects the time at which optimized spectral losses become decorrelated from the training-operator eigenvalues. The asymptotic prediction optimality of the REML-guided early stopping time is established for fixed-design in-sample risk and random-design out-of-sample risk. Deep learning models in fixed-kernel gradient regimes provide canonical instantiations for our results. Numerical experiments and a UK Biobank proteomics application show competitive accuracy of the REML-guided early stopping time with reduced reliance on validation splits and repeated checkpoint evaluation.