



THE UNIVERSITY OF
CHICAGO

DEPARTMENT OF STATISTICS

Master's Thesis Presentation

Jose Cruzado

Department of Statistics
The University of Chicago

“Expected Gradient Outer Product Reparameterization in Deep Convolutional Networks”

February 16, 2026, at 2:00 PM
Jones 111, 5747 S. Ellis Avenue

Abstract

Recent work has shown that reparameterizing optimization variables in the eigenbasis of the expected gradient outer product (EGOP) matrix can substantially improve the convergence of adaptive methods commonly used in practice (e.g., Adam, Adagrad, AdamW) by aligning the parameterization with the dominant directions of variation of the objective. In practice, however, the high computational and memory requirements of EGOP-based transformations in high-dimensional settings have restricted their use primarily to linear models and small-scale problems. In this work, we introduce a scalable EGOP reparameterization approach for convolutional neural networks that leverages the structural regularities of convolutional layers through a layerwise construction, in which a single EGOP-derived transformation is shared across all filters within a layer. This design enables EGOP-based reparameterization to extend to deep architectures such as ResNet and to large-scale datasets including ImageNet while substantially reducing both computational and memory overhead. The resulting method integrates seamlessly with standard deep learning pipelines and retains the geometric benefits of EGOP alignment. Experiments on deep convolutional models show that adaptive optimization algorithms converge more rapidly under the proposed reparameterization than in the original parameter space, without incurring prohibitive additional cost, and in some cases outperform strong state-of-the-art optimizers such as SOAP, demonstrating that EGOP reparameterization can be both practical and effective at scale.