Nghia Nguyen-Trung

Research summary · last updated November 2026

Ph.D. Candidate, Statistics and Operations Research, UNC–Chapel Hill

Advisor: Quoc Tran-Dinh · Ph.D. expected May 2027 · On the 2026–27 academic job market

nghiant@unc.edu · nghiant01.github.io · Google Scholar

Fast first-order methods for minimax problems, generalized equations, and fixed-point problems

Many problems in machine learning and operations research ask for an equilibrium rather than a minimum: between a model and an adversary, among players in a game, or among agents that share a network. For many such problems, the optimality or equilibrium conditions can be written in a common form called a *generalized equation*, $0 \in Gx + Tx$. In practice, such problems must be solved from noisy samples, on many processors at once, or across agents that cannot share their local data, and often without the monotonicity that classical theory relies on. My research develops first-order algorithms for these settings. Each method comes with a convergence guarantee and is tested numerically on problems from machine learning, reinforcement learning, games, and energy systems, such as a simulated virtual power plant.

1 Deterministic foundations

Which classical methods still work without monotonicity, and how fast?

- A unified theory of extragradient-type methods (extragradient, past-extragradient/optimistic gradient, Tseng's forward-backward-forward splitting, reflected forward-backward splitting, golden ratio), with sublinear convergence rates under a weak-Minty solution or co-hypomonotonicity. (Tran-Dinh & Nguyen-Trung, COAP 2026)
- Anchored (Halpern) and Nesterov-accelerated extragradient classes with last-iterate $O(1/k)$ and $o(1/k)$ rates for the residual under co-hypomonotonicity. (Tran-Dinh & Nguyen-Trung, COAP, to appear)

2 Stochastic acceleration and variance reduction

Can acceleration survive sampled data?

- Accelerated past-extragradient methods with variance reduction for a class of nonmonotone generalized equations: an $O(1/k^2)$ rate for the expected squared residual, compared with $O(1/k)$ for non-accelerated methods, and $o(1/k^2)$ in expectation and almost surely. For finite sums, the oracle complexity is $O(n\log n + n^{2/3}\varepsilon^{-1})$ with L-SVRG or SAGA and $O(n\log n + n^{1/2}\varepsilon^{-1})$ with L-SARAH. (Tran-Dinh & Nguyen-Trung, under revision at MOR)
- Variance-reduced forward-reflected-backward splitting with unbiased (L-SVRG, SAGA) and, to our knowledge for the first time for inclusions, biased (SARAH, Hybrid SGD, a new Hybrid SVRG) estimators. With L-SVRG or SAGA, the oracle complexity is $O(n^{2/3}\varepsilon^{-2})$ for finite sums and $O(\varepsilon^{-10/3})$ for expectations, compared with $O(\varepsilon^{-4}\ln(\varepsilon^{-1}))$ for increasing mini-batch SGD. Tested on AUC maximization and policy evaluation in reinforcement learning. (Tran-Dinh & Nguyen-Trung, under revision at COAP)

3 Systems-aware methods (first author)

Fast methods when computation is delayed, distributed, or partial

- **Delayed and asynchronous.** Accelerated fixed-point-based methods with delayed inexact oracles: $O(1/k^2)$ and $o(1/k^2)$ rates for the squared residual, with complexity linear in the maximum delay τ. (Nguyen-Trung & Tran-Dinh, arXiv:2512.13547)
- **Decentralized.** To our knowledge, the first accelerated decentralized algorithms for composite monotone inclusions over networks. Agents keep their local operators private, one variant lets them use their own network-independent step sizes, and the methods achieve $O(1/k^2)$ and $o(1/k^2)$ rates for squared residual and consensus quantities. Tested on distributed bilinear games and a simulated virtual power plant. (Nguyen-Trung, Necoara & Tran-Dinh, arXiv:2609.14953)
- **Randomized coordinate.** Optimistic gradient algorithms that combine randomized coordinate updates with variance reduction for composite inclusions. (Nguyen-Trung, Necoara & Tran-Dinh, preprint 2026)

Next 3–5 years

- **Distributed and asynchronous methods beyond current theory.** I will extend fast distributed and asynchronous methods to decentralized problems without monotonicity, to agents that have only noisy local information, and to networks whose agents compute and communicate at different speeds.
- **Decision-making under uncertainty.** Many robust and distributionally robust decision models can be formulated as minimax problems. I plan to study how the first-order methods I have developed can be adapted to solve such models reliably at the scale that data-driven applications require.

Record at a glance

Ph.D.: 2 journal articles (1 published, 1 accepted), 5 preprints (2 under revision).
Earlier: 6 journal articles from HUST research.
Service: ICML 2026 Silver Reviewer Award; reviewer for ICML, NeurIPS, and four journals, including JOTA.

Talks and teaching

INFORMS Annual Meeting 2026: contributed talk on the randomized coordinate methods; chair of a session on first-order, proximal, and interior-point algorithms.
SIAM MDS26: poster on the delayed-oracle methods.
Teaching: instructor of record for four summer courses; two departmental teaching awards (2024, 2025).