Internship Offer at GIPSA-lab:
Control-Theoretic Enhancements for Gradient-Based Neural Network Training

Context and Motivation

The intersection of control theory and machine learning has recently emerged as a fertile research area, particularly for online training of neural networks. While classical optimization algorithms such as gradient descent and Nesterov acceleration are widely adopted, they often face limitations in convergence speed, robustness, and sensitivity to hyperparameters. Integrating adaptive control strategies into the learning process has proven to mitigate these issues.

Previous work has highlighted two complementary directions. First, Airimitoaie et al. [2023, 2022] have shown that recursive least squares algorithms with dynamically adjusted adaptation gains can substantially improve convergence and robustness in parameter estimation. Second, Zhao et al. [2019, 2020] demonstrated that feedback-based and event-driven modulation of learning rates can accelerate online neural network training while reducing unnecessary computations. Together, these contributions suggest that control-theoretic principles can be systematically applied to enhance classical gradient-based optimization methods.

Building on these findings, this internship will explore new ways to integrate control mechanisms into gradient descent and its accelerated variants. The ultimate goal is to design algorithms that achieve faster convergence, better stability, and robustness in dynamic, online learning scenarios.

Internship Objectives

The primary objective is to develop and evaluate adaptive control strategies for improving gradient-based optimization. The intern will focus on designing controllers that modulate the learning rate and possibly the update schedule, inspired by the feedback and event-driven approaches observed in prior work. In addition, the intern will:

- Formulate theoretical models to analyze the stability, convergence of the proposed algorithms.
- Implement the algorithms in a deep learning framework such as PyTorch or TensorFlow.
- Evaluate performance on standard benchmark datasets (e.g., CIFAR-10, CIFAR-100, MNIST) and compare against classical optimizers like Adam, RMSProp, and vanilla gradient descent.
- Investigate hybrid strategies combining acceleration techniques with adaptive control for online or streaming learning tasks.

This combination of narrative description and concise bullet points allows the reader to quickly grasp the actionable tasks while keeping the overall flow readable and professional.

Expected Contributions

The intern is expected to make both theoretical and practical contributions. They will design algorithms that integrate control-theoretic insights with classical optimization methods, and evaluate them empirically to quantify improvements in speed, stability, and robustness. The work may lead to the preparation of technical reports and publications in relevant conferences or journals, contributing to the wider research community.

Candidate Profile

The ideal candidate is a motivated Master's student or early-stage PhD in Computer Science, Applied Mathematics, Automatic Control, or a related discipline. They should have:

- Good foundations in optimization, machine learning, and control theory.
- Experience in Python programming and deep learning frameworks.
- Analytical skills to study algorithmic stability and convergence.
- Curiosity and independence, with the ability to design experiments and interpret results.

Prior experience in online learning, adaptive control, or accelerated optimization is a plus, but not strictly required.

Supervision and Environment

The internship will be conducted at GIPSA-lab, Grenoble FR, an excellent research laboratory of Univ. Grenoble Alpes and CNRS. The student will work in a collaborative research group specialized in automatic control and machine learning. The intern will benefit from close supervision and access to high-performance computing resources. The environment combines theoretical modeling, algorithm design, and computational experimentation, ensuring a comprehensive research experience.

Duration

The internship is expected to last four to six months and can be arranged as full-time or part-time to accommodate academic schedules. Students should start around February/March 2026.
Please apply before, January 12, 2026 !

References

Tudor-Bogdan Airimitoaie, Ioan D. Landau, B. Vau, and G. Buche. Can dynamic adaptation gain speed up recursive least squares algorithm? *HAL Preprint*, 2022. URL `https://hal.science/hal-04790798v1`.

Tudor-Bogdan Airimitoaie, Ioan D. Landau, B. Vau, and G. Buche. Accelerating recursive least squares adaptation/learning algorithms using a dynamic adaptation gain. *HAL Preprint*, 2023. URL `https://hal.science/hal-04885000v1`.

Zilong Zhao, Sophie Cerf, Bogdan Robu, and Nicolas Marchand. Feedback control for online training of neural networks. *arXiv preprint*, 2019. URL `https://arxiv.org/abs/1911.07710`.

Zilong Zhao, Sophie Cerf, Bogdan Robu, and Nicolas Marchand. Event-based control for online training of neural networks. *arXiv preprint*, 2020. URL `https://arxiv.org/abs/2003.09503`.

Contact

Interested candidates should submit a CV, a motivation letter, and academic transcripts to:

Bogdan Robu
`bogdan.robu@gipsa-lab.fr`
Univ. Grenoble Alpes, GIPSA-lab

Matteo Tacchi-Bénard
`matteo.tacchi@gipsa-lab.fr`
CNRS, GIPSA-lab