

Application Intention Recognition for Smart Networking

1 Introduction

The increasing diversity of Internet applications has led to highly heterogeneous performance requirements. For example, real-time communication and interactive gaming demand ultra-low latency and jitter, while video streaming requires high throughput with resilience to moderate packet loss. Despite these differences, most applications do not explicitly expose their performance requirements to the network. Instead, their underlying *intention*—the implicit preferences for latency, bandwidth, reliability, and throughput—must be inferred indirectly from recurring traffic patterns.

Inferring application intention is challenging. Modern traffic is encrypted, protocol features evolve rapidly, and application behaviors vary dynamically across sessions. Traditional heuristics such as port-based classification or static thresholds often fail under these conditions. More advanced techniques are needed to extract meaningful requirements from encrypted flows, flow-level statistics, and contextual traffic features.

A robust framework for application intention identification would enable intent-driven networking and path-aware architectures (e.g., SCION) to allocate network resources more efficiently, adapt routing dynamically, and ultimately improve end-user experience.

2 Project Description

The goal of this project is to design and evaluate methods for **inferring application intention from network traffic**, focusing on passive inference mechanisms. The work will be organized into the following research tasks:

Task 1: Traffic Characterization and Feature Selection

Analyze representative traffic traces from diverse applications (e.g., conferencing, streaming, file transfer). Candidate features will include flow-level statistics (duration, packet size distribution, inter-arrival times), NetFlow/IPFIX records, and packet-level representations suitable for machine

learning. The objective is to identify the most informative and lightweight feature set that balances accuracy with computational cost.

Task 2: Application Intention Modeling

Formulate the inference problem at different levels of granularity:

- **Application classification** (categorical output),
- **QoE/QoS labeling** (e.g., “low latency,” “high throughput,” “loss resilience”),
- **Numerical regression** (estimating throughput demand in Mbps, latency tolerance in ms).

Machine learning models, including representation learning approaches, will be explored to capture probabilistic requirements from traffic.

Task 3: Workload Inference under Dynamics

Investigate how application requirements evolve across communication phases (e.g., handshake, steady-state transfer, control channels). Temporal models and adaptive inference mechanisms will be developed to update predictions as session characteristics change, ensuring timely alignment with application intent.

Task 4: System Integration and Evaluation

Develop a prototype that extracts traffic features and produces real-time predictions of application intention. The prototype will be evaluated in controlled experiments with heterogeneous workloads, focusing on accuracy, robustness under encryption, prediction latency, and scalability in high-throughput environments.

Background and Skills

This project sits at the intersection of **computer networks** and **artificial intelligence**. A successful candidate is expected to have a solid understanding of networking concepts such as transport protocols and flow monitoring, combined with familiarity in applying AI or machine learning techniques to real-world data. This interdisciplinary foundation is essential for bridging the gap between traffic-level observations and the inference of higher-level application requirements.

3 Supervisor and Advisor

The project will be carried out under the guidance of both a supervisor and an advisor.

- **Supervisor:** Prof. Adrian Perrig, ETH Zürich.
- **Advisor:** Liwen Xu, ETH Zürich.

4 Organization

The student and the advisor will hold weekly meetings. During each weekly meeting, the student will briefly describe the work completed during the week and outline the work to be completed during the next week. The advisor will, if necessary, assist the student in identifying potential future issues and discuss current issues. Pressing complications arising between two meetings will be promptly discussed. The advisor will assist the student towards completing any agreed-upon milestones.

At the end of the project, the student is required to submit a final project report that summarizes the work, results, and conclusions.

5 Grading Scheme

Grade	Description
6.0	Implementation at production-ready level. The student not just follows the advisor suggestions but brings their own innovative ideas, which add value to the project.
5.5	Project quality significantly exceeds expectations.
5.0	Project meets expectations.
4.5	Project partially meets expectations and has minor deficits.
4.0	Project meets minimum quality requirements but has major deficits and is clearly below expectations.