Contribution ID: **422** Type: **Presentazione orale**

A Multi-Objective Deep Reinforcement Learning Controller for E-Bike Front-Wheel Anti-Lock Braking

Thursday 3 September 2026 10:30 (15 minutes)

Objective
The primary goal of this work is to develop a Deep Reinforcement Learning (DRL)-based Antilock Braking System (ABS) controller specifically optimized for e-bike front-wheel braking. The system aims to:
- Minimize stopping distance while preserving directional stability and preventing hazardous front-wheel lockups.
- Eliminate the need for explicit, real-time longitudinal slip estimation and reduce reliance on noisy or delayed vehicle speed estimates, making it ideal for resource-constrained, low-cost e-bike sensor suites.

Methods
The proposed framework integrates advanced physical modeling, real-time state estimation, and a robust learning-based control policy:
- System Modeling: A comprehensive, nonlinear e-bike model was developed, capturing longitudinal motion, front-wheel rotational dynamics, vertical suspension, dynamic load transfer, and nonlinear tire-road interactions. This was combined with the unique dynamics of a pump-less ABS hydraulic architecture (hold/release solenoid valves, a check valve, and a volume-limited discharge accumulator) where braking pressure is constrained by rider input and accumulator saturation
- State Estimation: To ensure reliable wheel dynamics under asynchronous real-time constraints, a hybrid Kalman-filter-based estimator was implemented. It blends encoder-based wheel speed measurements, filtered acceleration states, and event-driven tonewheel pulse updates into a computationally lightweight embedded architecture.
- Control Policy & Training: The controller utilizes a discrete-action Proximal Policy Optimization (PPO) agent with an actor-critic architecture. The actor maps a velocity-independent observation vector (vehicle deceleration, estimated wheel speed, estimated wheel acceleration, and caliper pressure) into three discrete actions: pressure release, hold, or reapply. The policy was trained in simulation using a multi-objective reward function balancing braking progress, stopping distance, and safety (penalizing wheel lockup and rollover risks). Longitudinal slip was utilized exclusively during training to guide the agent but omitted during deployment. Domain randomization of vehicle mass, center-of-gravity position, and tire-road friction coefficients was applied during training to guarantee controller robustness.

Results
Simulation testing validated the efficacy of the DRL-based controller, demonstrating:
- Effective pressure modulation and robust wheel-lock prevention across highly diverse riding scenarios.
- Near-optimal friction utilization accomplished entirely without explicit real-time slip feedback.
- High adaptability to varying loading conditions, distinct rider profiles, and fluctuating road surface friction coefficients.

Conclusion
The proposed DRL architecture delivers a sensor-minimal, computationally efficient, and highly adaptive front-wheel ABS solution. By successfully mitigating the hardware limitations inherent to micro-mobility vehicles—such as limited pressure authority and low-cost sensing—this approach establishes a viable, scalable framework for significantly enhancing e-bike safety.

Primary authors: TONOLI, Andrea (Politecnico di Torino); CASTELLANOS MOLINA, Luis Miguel (Politecnico di Torino)

Presenter: CASTELLANOS MOLINA, Luis Miguel (Politecnico di Torino)

Session Classification: Automotive

Track Classification: Automotive