# Abstract

Suspension control systems in cars are a vital component in modern vehicles tasked with enhancing ride comfort for passengers and protecting occupants from the impact of road disturbances. In order to maximise the performance of these systems, modern cars make use of a variety of different types of suspension. One popular type used in many cars is known as semi-active suspension and allows for the suspension systems to dynamically change to match road disturbances. Specifically, these systems dynamically alter a value known as the damping rate, which is responsible for controlling the dampers in a car. By doing this the vehicle can dissipate energy stored in the suspension system's spring, which in turn improves ride comfort.

However, these systems rely on system controllers to monitor and dictate the damping rate. One of the most common types of controller for this purpose is known as a Proportional Integral Derivative (PID) controller. Whilst these controllers have a number of benefits associated with them there are also a number of drawbacks, including the linear nature of the controllers, the difficulties with tuning the controller's parameters, and the sensitivity of the controllers to changes in the environment.

On the other hand, many of these weaknesses are the strengths of reinforcement learning techniques which are capable of dealing with non-linearities and are self-training and highly adaptable. In recent years, a lot of research has gone into combining reinforcement learning with PID-controllers in order to benefit from each system's strengths.

In this research, a novel approach is taken to do this, by directly combining a PID-controller with a Deep Deterministic Policy Gradient (DDPG) reinforcement learning agent, through a technique known as Residual Policy Reinforcement Learning. This approach involves summing both policies together in order to create a final policy which exhibits the benefits of both systems.

This approach is evaluated on a quarter car suspension system model provided by ZF Friedrichshafen AG. The reinforcement learning agent is trained using the DDPG algorithm and is deployed, with a PID-controller as a base policy, on a variety of road disturbances. The results show that the algorithm is capable of improving the suspension control of a car, enhancing the performance of a PID-controller, and adapting to environmental changes. However, it is also shown that the algorithm's performance is highly reliant on the performance of the base policy.