ManeuverGPT

Agentic Control for Safe Autonomous Stunt Maneuvers

–Shawn Azdam, Pranav Doma, and Aliasghar Arab

NYU

Abstract

LLM-driven multi-agent control architecture (Query Enrichver, Driver, Validator) for executing high-dynamic stunt maneuvers like J-turns in CARLA. Shows 90 % success. With sedans and 70 % with coupes using purely prompt-based agents—no retraining.

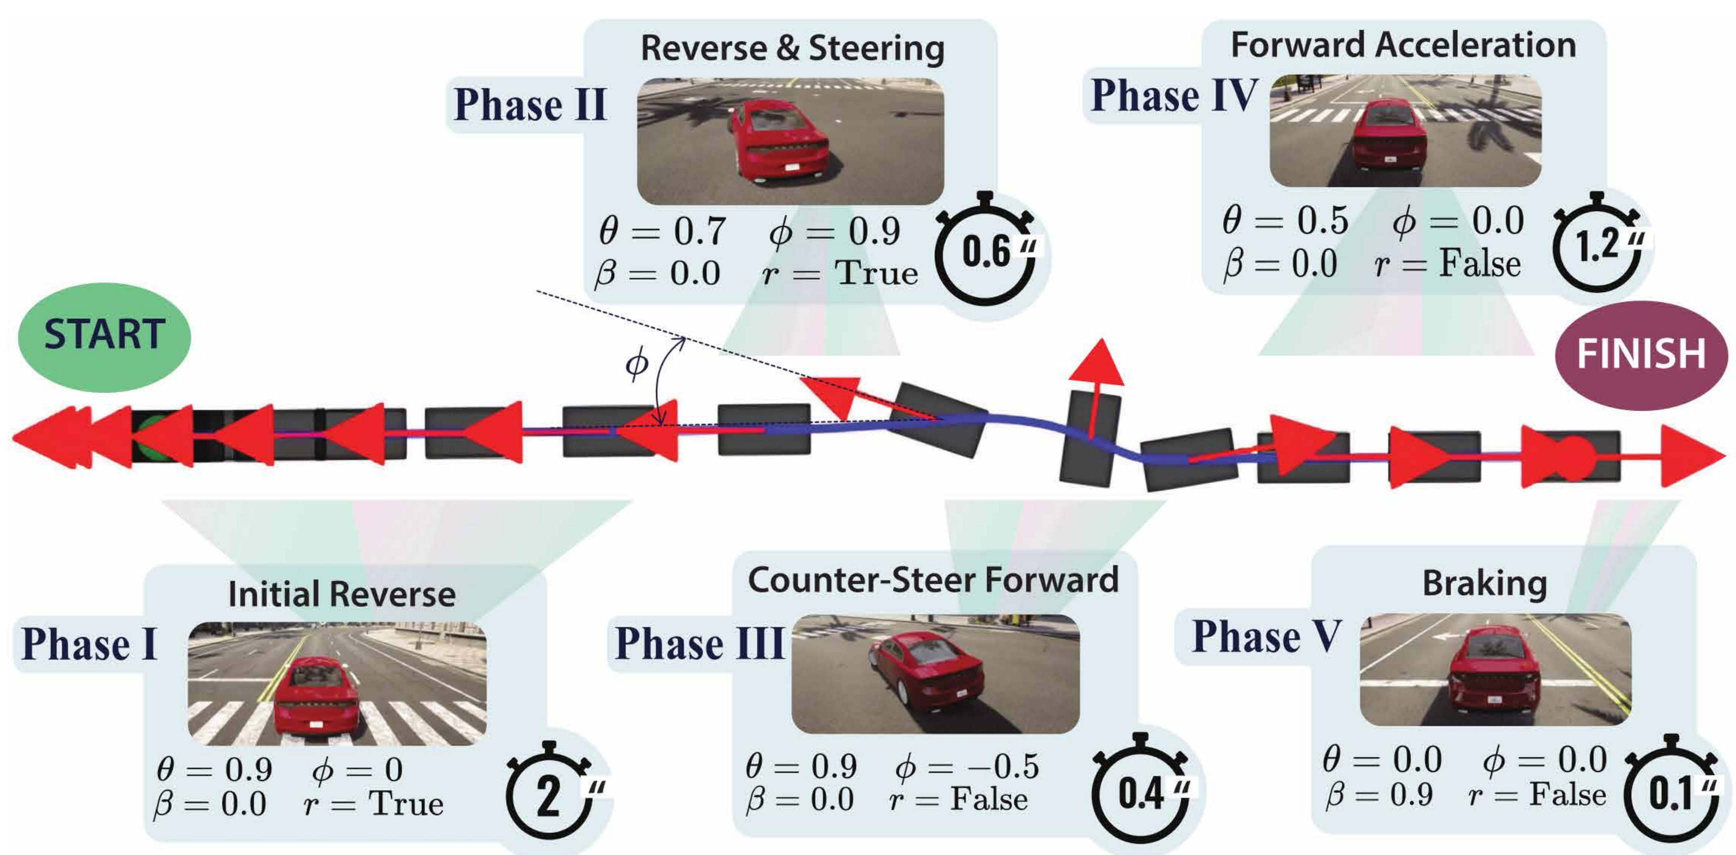

Methodology

- Query Enricher (AE) – adds context
- Driver Agent (AD) – generate parameters
- Parameter Validator (AV) – enforces constraints
- Orchestrator closes the loop

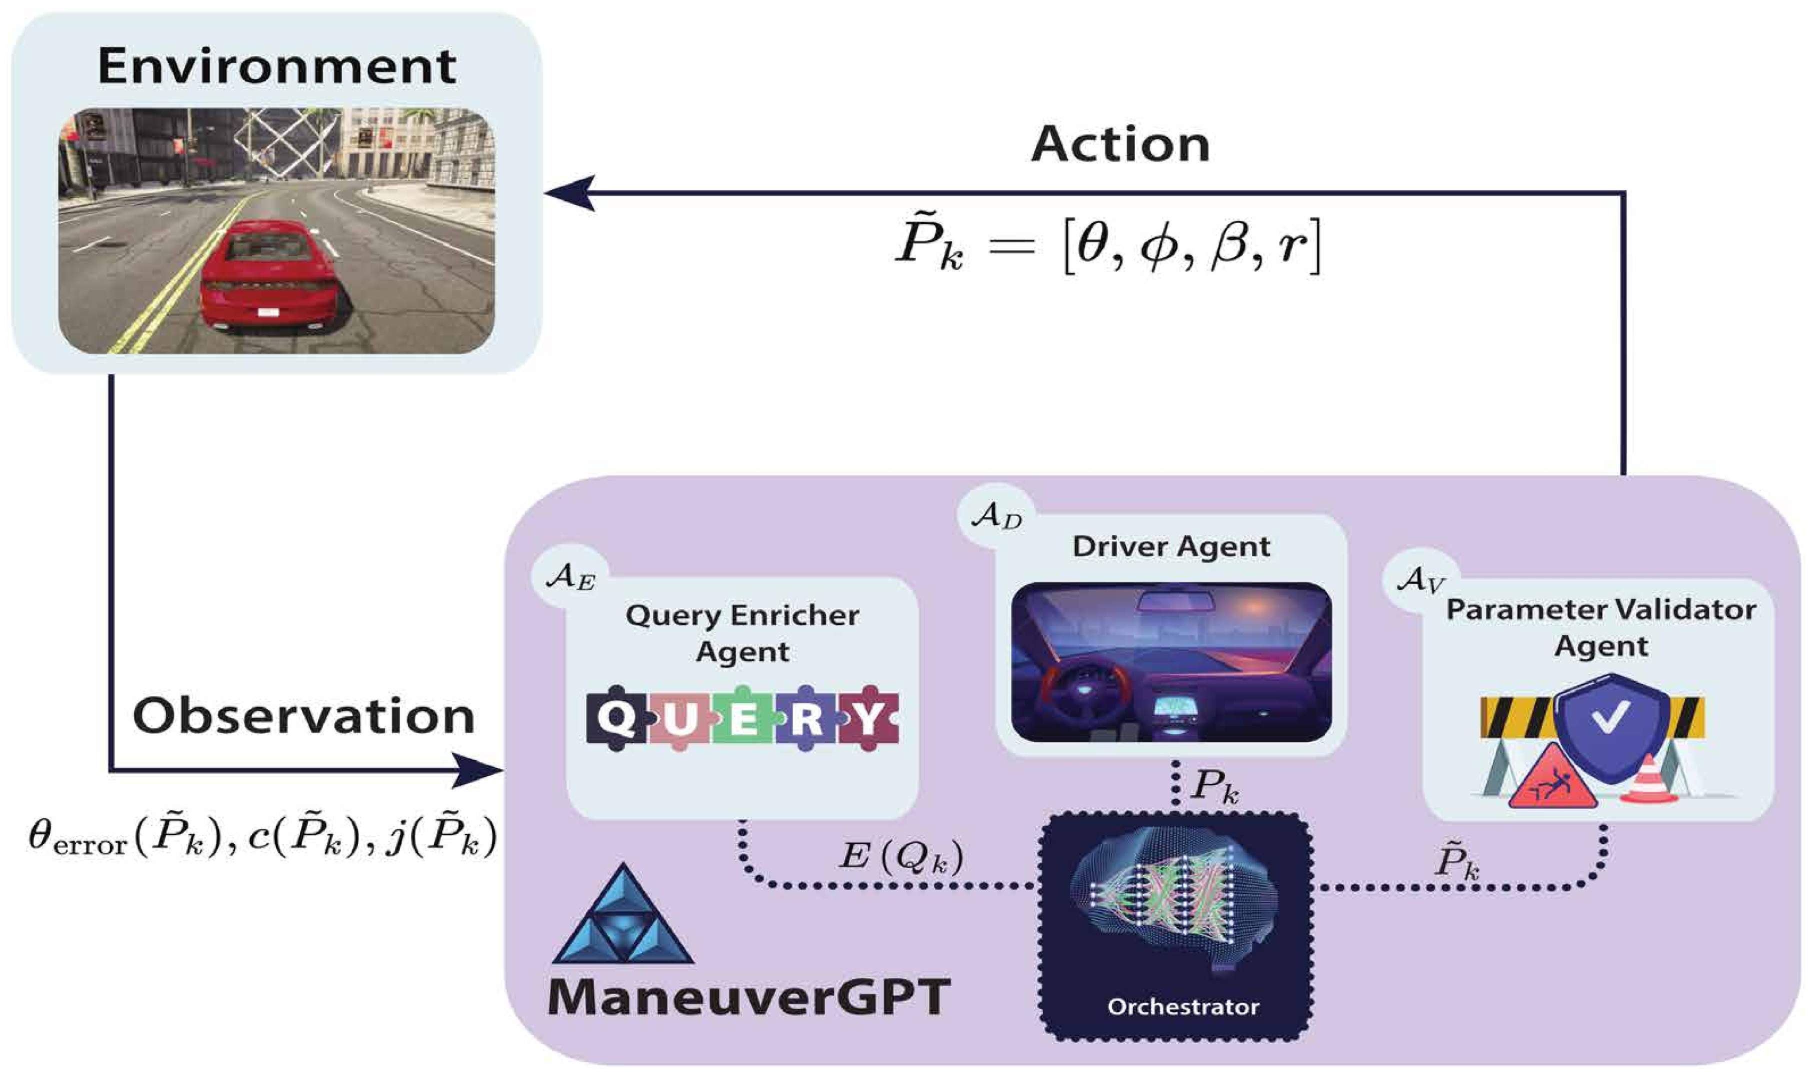

Results

- More stable Multi-agent system vs Single-agent
- Adaptability to different vehicles only via prompting
- Learning across trials from the feedback

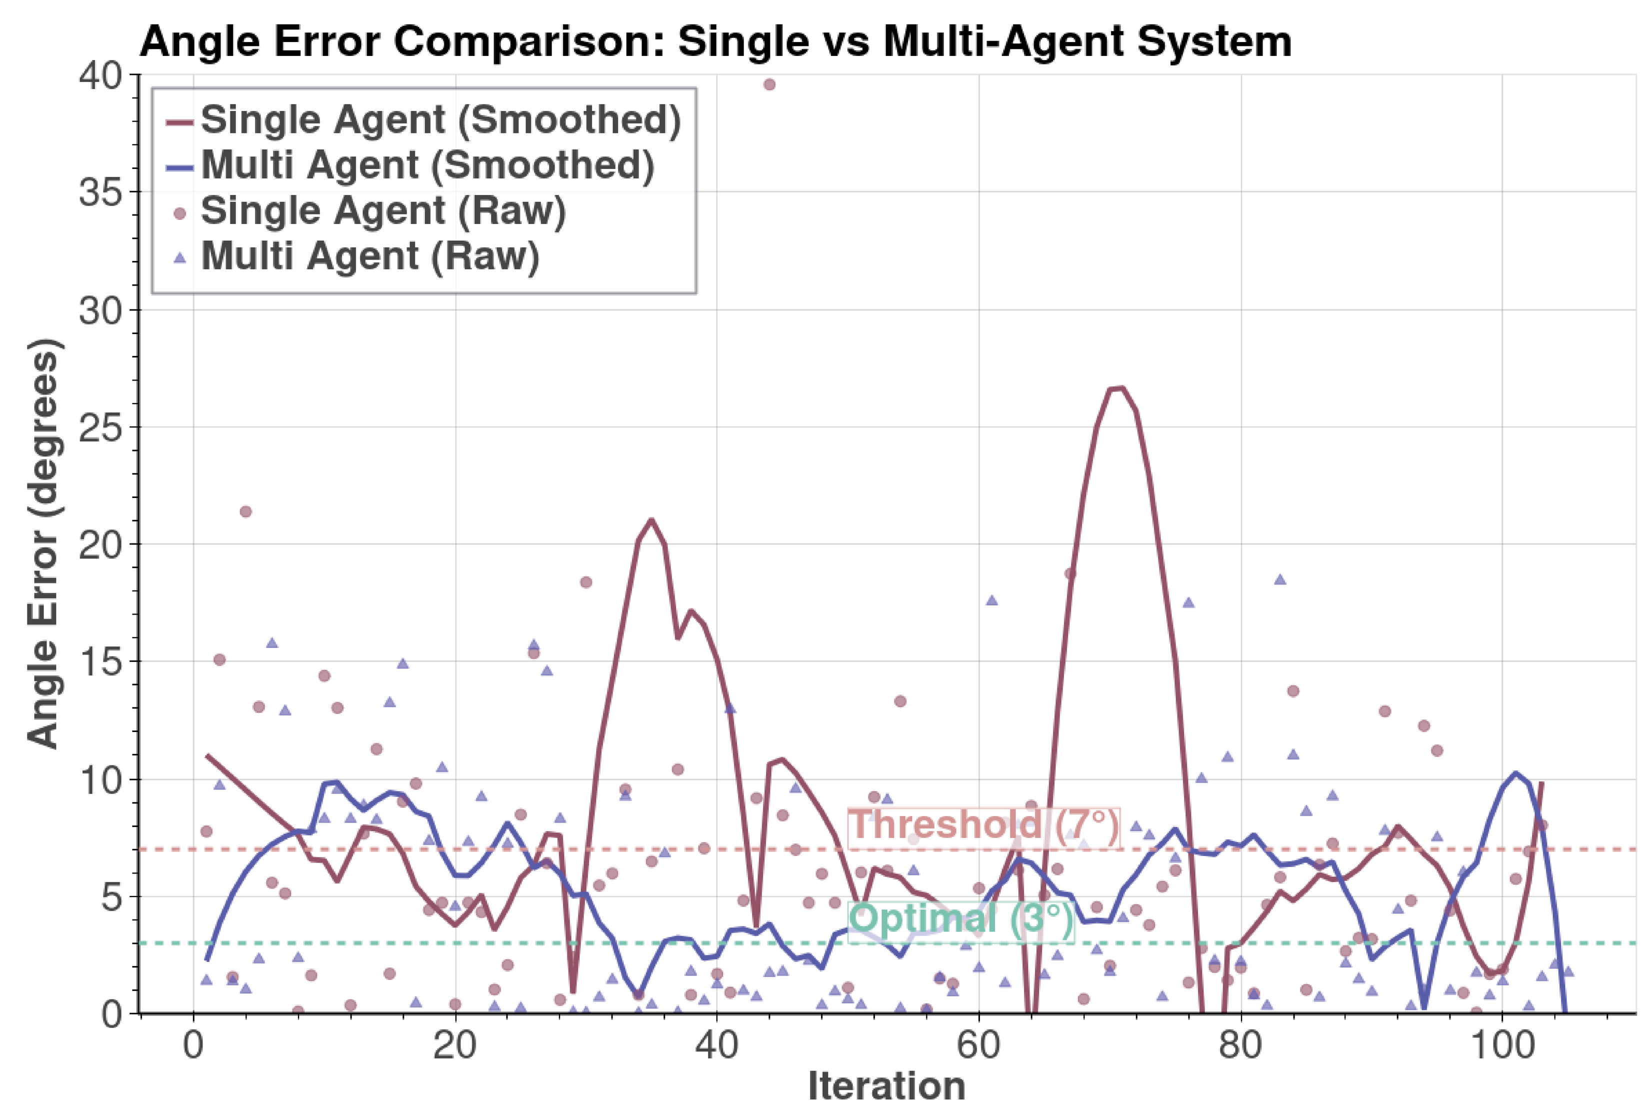

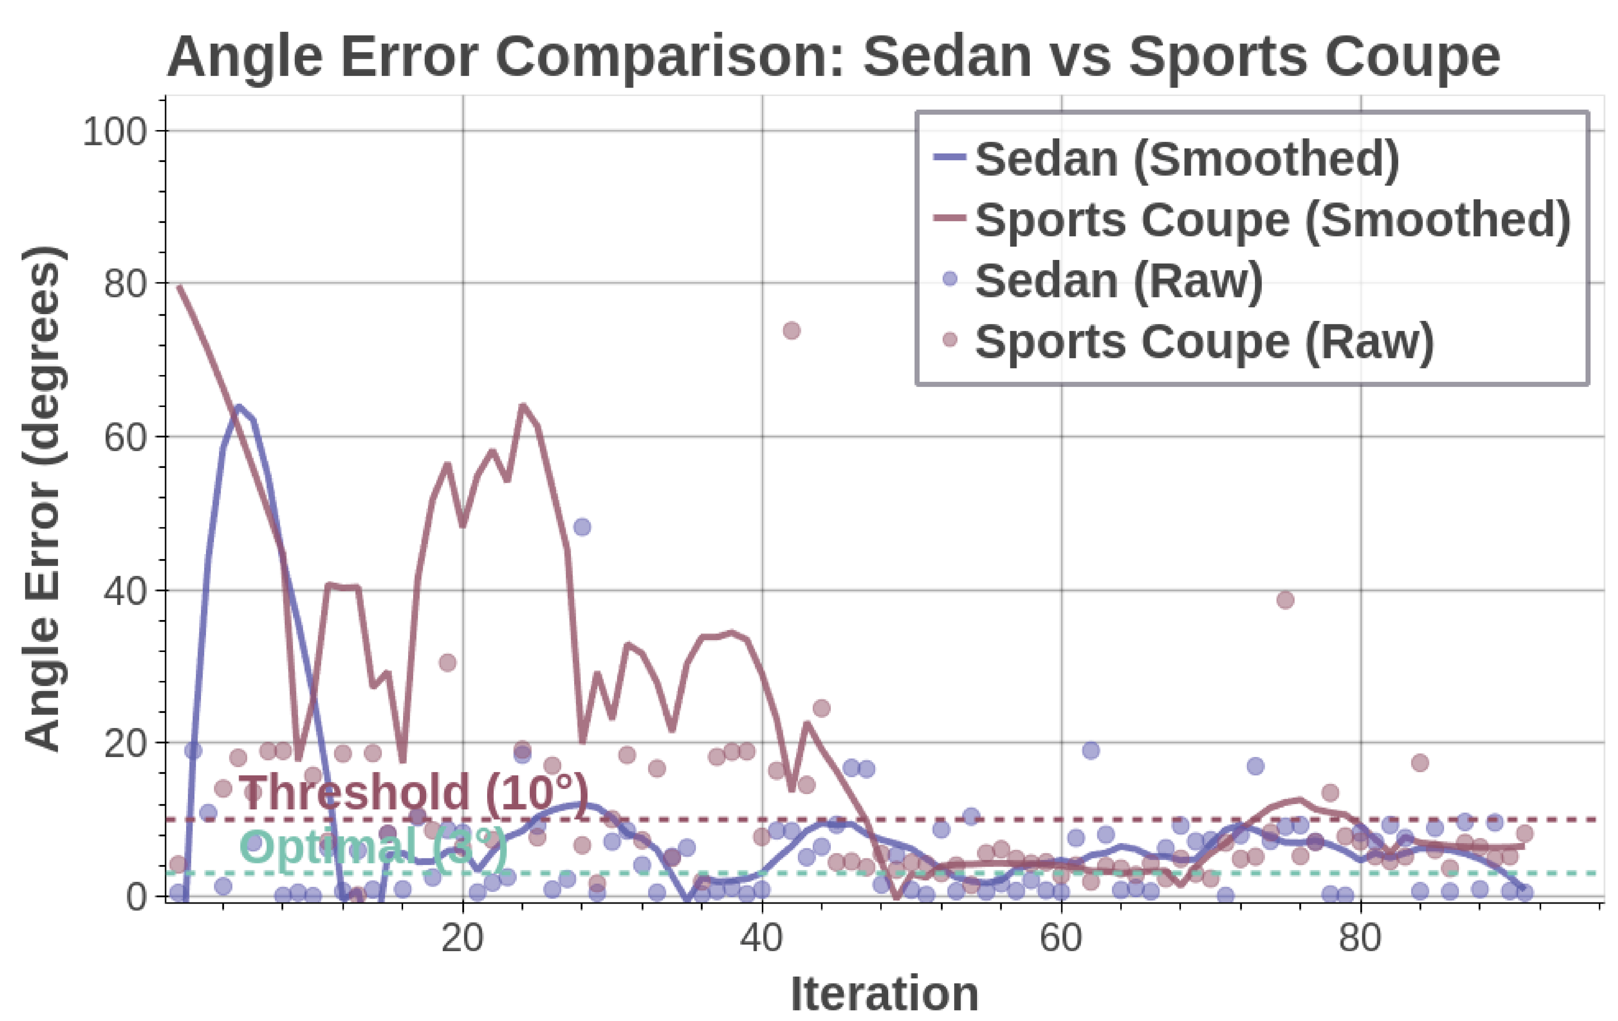

Algorithm 1: ManeuverGPT

Input : - User command Q_1
- Constraints $\mathcal{C}=\{C_s, C_o\}$
- Maximum iterations $k_{\max}$
- Cost threshold ε

Output: Feasible parameter set $\tilde{P}_k$ (or best-effort parameters)

$k \leftarrow 1$
$\tilde{P}_{\text{best}} \leftarrow \emptyset$; $L_{\text{best}} \leftarrow \infty$
while $k \le k_{\max}$ **do**
  $E(Q_k) \leftarrow \mathcal{A}_E(Q_k)$
  $P_k \leftarrow \mathcal{A}_D(E(Q_k))$
  $\tilde{P}_k \leftarrow \mathcal{A}_V(P_k)$
  if $L(\tilde{P}_k) \le L_{best}$ **then**
    $\tilde{P}_{\text{best}} \leftarrow \tilde{P}_k$; $L_{\text{best}} \leftarrow L(\tilde{P}_k)$
  end
  if $L(\tilde{P}_k) \le \varepsilon$ **then**
    return $\tilde{P}_k$ `// Satisfactory solution`
  end
  else
    $Q_{k+1} \leftarrow \text{Feedback}(Q_k, \theta_{\text{error}}(\tilde{P}_k), c(\tilde{P}_k), j(\tilde{P}_k))$
    $k \leftarrow k+1$
  end
end
return $\tilde{P}_{\text{best}}$ `// Best-effort solution`

SCAN ME

IROS HANGZHOU 2025