Using Particle Swarm Optimization to Learn a Lane Change Model for Autonomous Vehicle Merging

Na'Shea Wiesner
Gianforte School of Computing
Montana State University
Bozeman, MT 59717
nashea.wiesner@student.montana.edu

John Sheppard
Gianforte School of Computing
Montana State University
Bozeman, MT 59717
john.sheppard@montana.edu

Brian Haberman
Applied Physics Laboratory
Johns Hopkins University
Laurel, MD 20723
brian.haberman@jhuapl.edu

***Abstract*—This paper presents the results of experiments applying a Particle Swarm Optimization (PSO) approach to lane changing for autonomous vehicles. The lane change model proposed is rule-based, where PSO learns the parameters of the rules. A study was conducted to compare the proposed lane change model to the existing lane change model in the microscopic simulator, SUMO. Experiments performed include simulating vehicles using the Krauss car-following model with the SUMO lane change model, with the proposed PSO lane change model, and with all lane changing decisions turned off. The latter case, where merges are replaced by vehicle reset, serves as a baseline for missed merge opportunities. The objective was to develop an adaptive approach to improve merge efficiency as an example of lane changing behavior. Varying vehicle densities and levels of congestion on the merge lane and through-lane were tested. Empirical results show the proposed lane change model is able to learn merging strategies with minimal collisions and is comparable to the SUMO lane change model in some scenarios. Further investigation is needed to improve performance and safety, but initial results show promise for the proposed PSO-based approach to autonomous lane changing.**

***Index Terms*—Autonomous vehicles, Krauss car-following model, SUMO, lane change model, particle swarm optimization.**

I. Introduction

Autonomous vehicles continue to be developed and tested as a means of improving roadway safety, increasing traffic flows, and reducing fuel consumption. According to Erdmann [1], microscopic driving dynamics are controlled by three main models in SUMO; car-following models, intersection models, and lane change models. Car-following models control the longitudinal motion of a vehicle, which mainly consists of adjusting or maintaining speeds in order to keep a safe distance (or gap) between vehicles. Intersection models control vehicle behavior at various types of intersections based on right of way laws, minimum acceptable gaps, and behavior that avoids intersection blockage. Lane change models determine the lane choice of a vehicle as well as speed adjustments related to changing lanes. However, lane change maneuvers have been empirically shown to be responsible for most traffic perturbations on freeways [2]. This paper focuses primarily on learning a lane change model for safe and efficient merging, and in particular, evaluating a proposed Particle Swarm Optimization Lane Change Model (PSO-LCM). The model is rule-based where PSO will learn the control parameters of the defines rules. PSO is able to explore multidimensional search spaces by maintaining diversity in a population. PSO is therefore suitable for this multi-objective problem of autonomous lane changing since there may be optimal sets of lane changing rules for differing conditions. PSO is also shown to have high speed convergence [3], making it more appealing to real world applications such as autonomous driving.

This work is a continuation of the development of PSO-based approach to autonomous vehicle control [4], where PSO was able to learn three control parameters (speed, inter-vehicle gap, and slow down) while maximizing average speeds and number of merges and minimizing collisions. The merge scenario was simpler, where the merge point was similar to that of a yield intersection. In the present paper, the main contribution is a rule-based lane change model that leverages PSO's balance of exploration vs. exploitation in order to learn the parameters of predefined autonomous control rules. Similar to the previous work, the motivation for developing PSO-LCM is to ensure autonomous vehicles are able to operate safely in a mixed road network with human operated vehicles since road networks will not only consist of autonomous vehicles for many years to come.

In this work, we focus on merging as a specific type of lane changing behavior and evaluate PSO-LCM in a merge scenario, similar to that of a typical on-ramp to a freeway where the merge lane consists of a parallel acceleration lane that eventually ends, forcing the vehicle to change lanes to continue its route. The simulated road network used in this study is closed, meaning the freeway flows in a loop where vehicles can exit at random and merge back on during each lap of the loop, thus allowing us to focus on merging behavior. Vehicles were simulated using Simulation of Urban Mobility (SUMO) [5], an open source microscopic traffic simulator. The Krauss car-following model [6], which is the default car-following model in SUMO, is also used in this study. We compare PSO-LCM to the default SUMO lane change model [1], as well as a set of baseline experiments with no lane change model activated. Overall, this problem is much more difficult for PSO, given the rule-based nature of learning; therefore, it is not comparable to our previous results.

In order to test the effectiveness of the proposed PSO-LCM against the more complex SUMO lane change model in a merge scenario, the following hypotheses are made:

Hypothesis 1 *The use of PSO-LCM on a highway on-ramp can be optimized to allow for a higher merge rate than the default SUMO lane change model.*

For Hypothesis 1, PSO-LCM has the freedom to explore rule parameters that allow for safe and efficient strategies in merging, rather than depending on static rules. This, in turn, will help prevent deadlock at the end of the merge lane while simultaneously increasing overall traffic flow of the closed network.

Hypothesis 2 *PSO-LCM will be able to learn a lane change strategy that minimizes collisions in the merge lane while simultaneously maximizing the number of merges.*

Hypothesis 2 predicts that collisions at or around the merge lane will be minimized. PSO-LCM is unaware of traffic laws stating it must yield to vehicles in the through-lane, which likely will lead to more collisions early in the learning process. However, a merging behavior will emerge such that vehicles will be incentivized to merge only when there are sufficient gaps between vehicles on the through-lane.

The remainder of this paper is organized as follows. As background, a brief discussion of the SUMO lane change model and the Krauss car-following model will be given in Section II. In Section III, we present relevant related work. We then introduce our experimental approach, which includes the simulation set up, PSO-LCM definition, and experimental design in Section IV. The results are presented in Section V, followed in Section VI by a discussion of those results. Finally, in Section VII, we draw our conclusions and discuss potential future work.

II. Background

A. Krauss Car-Following Model

The Krauss car-following model is designed to specify the through-lane behavior of vehicles and is the default car-following model in SUMO [6]. The model maintains a safe gap between a vehicle and its leader by calculating a safe velocity v_{safe}:

$$v_{safe} = v_l(t) + \frac{g(t) - v_l(t)t_r}{\frac{v_l(t)+v_f(t)}{2b} + t_r}$$

where $v_l(t)$ is the speed of the leading vehicle at time t, $g(t)$ is the gap to the leader at time t, t_r is the driver's reaction time (usually 1 s), and b is the maximum deceleration of the vehicle in m/s^2. It is possible that v_{safe} is larger than the speed limit or the maximum acceleration ability of the vehicle. Because of this, the minimum value of the three speeds is used to set the vehicle's desired speed:

$$v_{des} = \min[v_{max}, v + at, v_{safe}].$$

where a is the vehicle acceleration.

B. SUMO Lane Change Model

In general, a lane change model dictates the lateral movement of vehicles between lanes. Lane change models have two main purposes; first, they compute the change decision of a vehicle in a single simulation step based on the route of the vehicle, and second, they compute speed changes for the vehicle itself and any obstructing vehicles to allow for a successful lane change maneuver. For the SUMO lane change model specifically, four motivations for changing lanes are defined [1]:

- Strategic change
- Cooperative change
- Tactical change
- Regulatory change

Strategic lane changes involve changing lanes in order for the vehicle to reach the next edge on a vehicle's route. Cooperative lane changes involves changing lanes and/or speed solely to help other vehicles with lane change maneuvers. Tactical lane changes occur when a vehicle does not wish to follow a slow leader. It balances the speed gains of lane changing against the effort of lane changing. It also has to take into consideration blocking the overtaking lane if the speed gains are minimal, potentially resulting in major impediments to traffic flow. Finally, regulatory lane changes are for maintaining either right hand or left hand driving jurisdictions. If a vehicle is in an overtaking lane and is not currently overtaking another vehicle, it is obligated to change lanes toward the driving lane.

The SUMO lane change model uses a series of decision trees to decide the next appropriate action based on a vehicle's route, speed, current position, and position and speed of neighboring vehicles [1]. These lane change decisions are organized in a hierarchical schema where d is the currently considered direction of lane change ($d = -1$ for change right; $d = 1$ for change left):

1) Urgent strategic change to d needed: change (strategic)
2) Change to d would create an urgent situation: stay (strategic)
3) Vehicle is a blocking follower for another vehicle with urgent strategic change request: change (cooperative)
4) $speedGainProbability$ above threshold and its sign matches d: change (tactical)
5) $keepRightProbability$ above threshold and $d = -1$: change (regulatory)
6) Non-urgent strategic change to d needed: change (strategic)

where $speedGainProbability$ indicates the direction (by its sign) and probability of benefiting (such as an increase in speed) from a lane change, and $keepRightProbability$ is the probability of a lane change to the right occurring and will trigger this lane change after a certain amount of time in order to avoid driving in the overtaking lane. A lane change is considered urgent if the distance to an obstructing leader is less than the required distance needed for a lane change, which is pre-defined in the simulator. The required distance needed for a lane change is based on the vehicle's length, speed, and minimum acceptable gaps between vehicles as determined by the Krauss car-following model.

III. Related Work

The Krauss car-following model is the default car-following model in SUMO and has therefore been studied extensively and compared to real world driving. Ma *et al.* [7] used a modified Krauss model and built a fuzzy control model on top for intelligent vehicles. The behavior of different car-following models at controlled intersections including Krauss was compared, and real world traffic data was used to calibrate the models in order to fit real world trajectories. Results showed the Intelligent Driver Model (IDM) more closely fit real world trajectories [8]. IDM is a time-continuous car-following model [9], where the acceleration a is a function of velocity v_α, gap s_α, and velocity difference to the proceeding vehicle Δv_α:

$$v_\alpha = a^{(\alpha)} \left[1 - \left(\frac{v_\alpha}{v_0^{(\alpha)}} \right)^\delta - \left(\frac{s^*(v_\alpha, \Delta v_\alpha)}{s_\alpha} \right)^2 \right]$$

where δ is the acceleration exponent(usually set to 4), v_0 is the desired velocity, and $s*$ is the desired minimum gap. IDM is considered collision free due to its dependence on the relative velocity.

In recent research focusing on metaheuristic optimization algorithms for autonomous vehicle control, Naranjo *et al.* [10] studied several metaheuristic methods for learning speed parameters of autonomous vehicles. Simulation results showed improvement over previous methods in speed adjustments. Although not directly applied to autonomous vehicles, Olivas *et al.* [11] also used metaheuristics to learn parameters of the gravitational search algorithm (GSA) based on Newton's laws of gravity and acceleration where they proposed dynamic parameter adjustment for GSA using type-2 fuzzy logic.

The cooperative behaviors among communicating vehicles in the SUMO lane change model and MOBIL (Minimizing Overall Braking Induced by Lane Changes) [12] was studied by Khan *et al.* [13]. Vehicle-to-Vehicle (V2V) communication was employed along with distributed learning algorithms to search for local lane change plans to help individual vehicles maintain preferred speeds, as well as the effect a lane change might have on neighboring vehicles. Results showed MOBIL was able to reduce overall braking and allowed vehicles to drive at desired speeds for longer durations. However, this performance boost came at the cost of increasing the frequency of lane changes and burning more fuel, producing more CO_2.

Nie *et al.* proposed a decentralized cooperative lane change model for connected (Vehicle-to-Vehicle communication) autonomous vehicles [14]. Three modules that comprise the model are defined; state prediction, candidate decision generation, and coordination. The state prediction module predicts the future state of vehicles using existing cooperative car-following models. A candidate decision is then generated from an incentive based model. Finally, in the coordination module, an algorithm is proposed that avoids candidate lane changing decisions that may lead to collisions or degraded traffic flow. Empirical results showed the proposed method was consistent with MOBIL, but some limitations were identified for future improvements. These limitations include degradation in traffic efficiency if the lane change execution was not immediate as well as in certain scenarios such as freeways with and without on ramps. This is attributed to their cooperative lane changing model being relatively simple, and therefore plan to introduce more complex cooperative lane change models in the future. PSO-LCM is also a relatively simple model, and will therefore need more elaborate rules to improve efficiency in certain situations.

A novel autonomous lane changing model was proposed by Liu *et al.*, using a support vector machine (SVM) paired with a Bayesian optimization algorithm (BOA) [15]. The proposed model was compared to a rule-based lane changing model using both simulation and real vehicle tests. The authors concluded that while their model showed promise, more research on the feasibility of the model in real world traffic needs to be done.

A convex optimization method was used by Huang *et al.* for solving lane change decisions for autonomous vehicles [16]. The model consists of a path generator and model-predictive-control-based vehicle steering and wheel torque control. A collision-free trajectory is generated when a collision between two vehicles on a two-way path is likely. The authors concluded that based on performance, the proposed method was feasible and avoided collisions in normal driving conditions.

Another approach to autonomous vehicle lane changing, studied by Hu *et al.*, is scheduling [17]. Lane assignment strategies were studied in a highway scenario. These strategies were evaluated based on their effect on overall traffic efficiency and safety. A cooperative lane change maneuver was proposed, known as Politely Change Lane (PCL) and was shown to improve overall efficiency and safety in especially heavy traffic conditions. PCL introduces a politeness index that indicates how much a driver takes into account the vehicle behind them during lane changing. Specifically, it looks whether or not a following vehicle in the target lane is slowing down to let the ego vehicle in, and if not, how much the following vehicle will need to slow down if the lane change occurs. In this way, PCL maneuvers enable a tradeoff between safety and efficiency.

IV. Methods

Our experiments focus on learning the parameters of a set of rules for lane changing behavior. The experiments depend upon a traffic simulator working in conjunction with PSO-LCM. This section describes the overall methods applied in evaluating our approach.

A. The Simulation

The simulation used for experiments was SUMO paired with its Traffic Control Interface (TraCI) [18] in order to enable online updates to the simulation. SUMO is a microscopic simulator, meaning each vehicle is simulated separately with individual routes and movement. The road network used was a

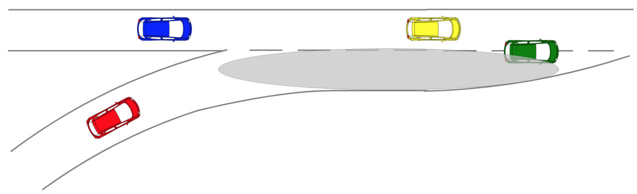

Fig. 1. Notional simulated merge lane.

TABLE I
INITIAL PARAMETER SETTINGS FOR SUMO

Parameter	Value	Unit
Simulation Steps	3,000	steps (s)
Main Loop Speed Limit	30	m/s
Merge Lane Speed Limit	15	m/s
Vehicle Type	Passenger	NA
Vehicle Length	5	m
Acceleration	2.6	m/s^2
Deceleration	4.5	m/s^2
Min Gap	2.5	m
Car Following Model	Krauss	NA
τ	1	s
σ	1	NA
Number of Vehicles	10, 20, 50	vehicles

closed network comprised of a single lane loop with an off/on ramp and merge lane, where vehicles drive counterclockwise around the track. A drawing of the merge lane is shown in Figure 1, where the shaded ellipse denotes where lane detector was located for collecting information on the merging vehicles. Not pictured are re-routers located on each edge to keep vehicles looping around the track or exit off and merging back on the track for the duration of the simulation. The main loop is 1.9km in length and the merge ramp is 0.78km in length from the point where the vehicles exit off of the loop to where the merge lane ends. Figure 1 is not drawn to scale.

Table I shows the initial parameter settings for the simulations run in this study. Each simulation covers 3,000 total simulation steps, where one simulation step is equal to one second. Therefore, each simulation corresponds to 50 minutes of driving. Vehicles are initially placed evenly among the edges throughout the simulation, however, we allow for a 500 step "burn-in" period, where no statistics from the simulation are collected. This allows the simulation to reach equilibrium, where the driving behavior is more realistic and vehicles begin to encounter one another more frequently. Therefore, statistics are only collected for 2,500 simulation steps. The speed limits on the main loop and the exit/merge ramp were set to 30 m/s (i.e., 108 km/hr) and 15 m/s (i.e., 54 km/hr), respectfully. The speed limit increases to 30 m/s once the vehicles reach the merge lane in order to get up to speed with the vehicles on the loop. These speeds were chosen to represent more realistic driving conditions.

All vehicles were passenger vehicles with length 5 m with maximum acceleration and deceleration rates of 2.6 m/s^2 and 4.5 m/s^2 respectfully. The Krauss car-following model was used with a minimum gap ($minGap$) allowed between vehicles on the loop of 2.5 m. These are all default parameter settings in SUMO. The Krauss car-following model has two parameters, $\tau \geq 0$ and $\sigma \in [0, 1]$, that allow for human behavior to be modeled. These parameters were both set to 1. A τ value of 1 indicates a driver's minimum desired headway is 1 s, and a σ value of 1 indicates the most imperfect level of driving modeled. These parameters were chosen to allow for more realistic driving. All other SUMO parameters not listed were set to default values.

Both the Krauss car-following and SUMO lane-change models are considered collision free due to the safety standards they employ such as $minGap$ and v_{safe}. Therefore, collisions within the experiments will only occur if a vehicle using the PSO lane change model violates the minimum space requirements of other vehicles. It should also be noted that collisions are only counted if they occur on the merge lane, adjacent through-lane, or the edge immediately following those lanes. This is to focus only on collisions caused by merging maneuvers and not from car-following behaviors. All vehicles in the simulations will use the Krauss car-following model paired with the either the SUMO lane change model, PSO-LCM, or all lane changing mechanisms turned off.

B. PSO Lane Change Model

PSO-LCM is a rule-based lane change model, where PSO learns the parameters to the predefined rule templates. PSO is a widely studied optimization algorithm that is able to find optimal results in a relatively short time, making it ideal for applications that are highly dynamic such as autonomous vehicle control. The model is designed to exhibit both strategic behavior for the merging vehicles and cooperative behavior for the vehicles on the through-lane. The procedure shown in Algorithm 1 provides a template of the strategic behavior of merging vehicles. The blanks indicate the values PSO-LCM learns during training. Intuitively, if there are ample gaps between a merging vehicle and its left neighbors, the vehicle will merge. Otherwise, the merging vehicle will adjust its speed until ample gaps are created:

The cooperative behavior is defined as shown in Algorithm 2. The gap in this case is known as the polite gap, where the through-lane vehicles will slow down to let merging leaders in front of them. The polite gap will be the same for all vehicle in a simulation due to the autonomous vehicle assumption and does not take into account different driver politeness levels.

The particle representation, along with each value's ranges, are listed in Table II. Each particle has 9 dimensions and is initialized randomly from the listed ranges. There is no value for speeding up after merging since the vehicle will either speed up to the defined speed limit or the speed of its leader while maintaining a safe distance.

The fitness function for PSO-LCM is based on the resulting number of merges and collisions from each simulation:

$$fitness = \ln \#merges - p \ln \#collisions$$

where p indicates how harsh of a penalty to invoke for collisions. For these experiments, p was tuned to 3 since 2

Algorithm 1 PSO Lane Change Process

```
procedure LANECHANGE
    for vehicle in merge lane do
        if vehicle has no leftLeaders & no leftFollowers then
            Merge
            Accel. to speedLimit over ______ seconds
        else if vehicle has leftFollowers then
            if gap > ______ m then
                Merge
                Accel. to speedLimit over ______ seconds
            else gap ≤ ______ m
                Accel. by ______ m/s over ______ seconds
            end if
        else if vehicle has leftLeaders then
            if gap > ______ m then
                Merge
                Accel. to Leader Speed over ______ seconds
            else gap ≤ ______ m
                Decel. by ______ m/s over ______ seconds
            end if
        else vehicle has leftLeader and leftFollower
            Check gaps, adjust speeds, merge
        end if
    end for
end procedure
```

Algorithm 2 Cooperative Behavior

```
procedure COOPERATIVE
    for vehicle in merge lane do
        if vehicle has rightLeader and gap < ______ m then
            Decel. by ______ m/s over ______ seconds
        end if
    end for
end procedure
```

produced many rule sets with collisions, and 4 did not allow for enough exploration of the search space. The natural log is used here to scale the terms across all vehicle densities. The $\#merges$ account for all successful merges occurring within a simulation. A successful merge means the merge did not occur with a collision or a teleportation to the next edge to avoid deadlock. The $\#collisions$ in this case account for all collisions that occur at or immediately following the merge lane.

A standard "gBest" PSO update was used. Inertia (ω) was tuned to 0.5 and both the cognitive and social update parameters (ϕ_1 and ϕ_2 respectively) were tuned to 0.9 to allow for ample exploration of the search space. The swarm size was set to 10 particles.

C. The Experiments

Experiments were conducted to evaluate the effectiveness of the proposed PSO-LCM model against the default SUMO lane change model. Each experiment consisted of 100 iterations where a simulation was run every time the fitness function was evaluated, resulting in (10 particles $\times$ 100 iterations) + 10 initial particles = 1010 simulations per experiment. Metrics collected included the number of successful merges, the number of collisions resulting from merge attempts, the average speed of vehicles on the loop, and the average speed of vehicles on the merge ramp. A range of vehicle densities (10, 20, and 50 vehicles) was tested in order to study different congestion levels. Also, different distributions between merging vehicles and through-lane vehicles (20%, 50%, and 80% of vehicles merging) were tested to specifically study congested merge ramps and through-lanes. PSO-LCM was compared to the SUMO lane change model (SUMO-LCM) as well as no change lane model for a baseline marker.

TABLE II
PSO-LCM PARTICLE REPRESENTATION

Value	Range	Unit
Polite Gap	(0-50)	m
Polite Merge Slow Down	(0-30)	m/s
Polite Merge Slow Down Time	(0-10)	seconds
Merge Gap	(0-50)	m
Speed Up After Merge Time	(0-10)	seconds
Speed Up To Merge	(0-30)	m/s
Speed Up To Merge Time	(0-10)	seconds
Slow Down To Merge	(0-30)	m/s
Slow Down To Merge Time	(0-10)	seconds

V. EXPERIMENTAL RESULTS

This section will present the results from experiments conducted, comparing the proposed PSO-LCM to SUMO-LCM and No-LCM. Number of merges, number of collisions, average loop speed, and average merge speed will be provided.

A. Merges

Figure 2 compares the average number of merges per simulation for PSO-LCM, SUMO-LCM, and no lane change model (No-LCM) for 10, 20 and 50 vehicles in the network. Results are shown for 20%, 50%, and 80% of vehicles merging on and off the loop to test different congestion levels in the through-lane. The merge results reported in Table III have been normalized over the min-max range of all of the experiments. For PSO-LCM this is the average number of successful merges for the swarm after training (1000 simulations). For SUMO-LCM and No-LCM, this is also the average number of merges over 1000 simulations for fair comparison. The merges occurring with No-LCM represents the number of times vehicles had to be moved or teleported to the next edge in order to allow vehicles to continue their route and avoid jamming in the simulation. The No-LCM vehicles did not merge successfully since there was no instruction to the vehicles on how to change lanes to continue their route. The results for No-LCM serve as a baseline for the minimum number of vehicles that need to merge to keep traffic flowing in this specific road network.

Overall, results for merging rate show similar performance between PSO-LCM and SUMO-LCM for 20% and 50% of vehicles merging on and exiting off the loop over the three vehicle densities studied. Comparable results to SUMO-LCM

TABLE III
RANGE NORMALIZED AVERAGE NUMBER OF MERGES ($\#merges$) FOR PSO-LCM, SUMO-LCM, AND NO-LCM

	20% Merging	50% Merging	80% Merging
PSO-LCM-10	0.071 ± 0.004	0.145 ± 0.01	0.246 ± 0.007
SUMO-LCM-10	0.071 ± 0.01	0.169 ± 0.012	0.253 ± 0.01
No-LCM-10	0.066 ± 0.008	0.111 ± 0.006	0.113 ± 0.006
PSO-LCM-20	0.115 ± 0.009	0.276 ± 0.027	0.348 ± 0.017
SUMO-LCM-20	0.142 ± 0.015	0.333 ± 0.017	0.495 ± 0.016
No-LCM-20	0.018 ± 0.008	0.112 ± 0.006	0.112 ± 0.006
PSO-LCM-50	0.183 ± 0.022	0.503 ± 0.051	0.697 ± 0.044
SUMO-LCM-50	0.21 ± 0.018	0.524 ± 0.027	0.858 ± 0.033
No-LCM-50	0.123 ± 0.009	0.114 ± 0.006	0.111 ± 0.006

TABLE IV
AVERAGE NUMBER OF COLLISIONS ($\#collisions$) FOR PSO-LCM

	20% Merging	50% Merging	80% Merging
PSO-LCM-10	1.32 ± 0.55	0.14 ± 0.29	0.09 ± 0.16
PSO-LCM-20	0.35 ± 0.45	3.52 ± 1.15	2.28 ± 0.68
PSO-LCM-50	0.93 ± 0.76	1.53 ± 0.78	11.48 ± 4.92

can also be seen for 10 vehicles with 80% of vehicles merging. For 80% of vehicles merging for 20 and 50 vehicles however, SUMO-LCM allows for more vehicles to merge. This shows a degradation in performance of PSO-LCM when the merge lane is more congested and the through-lane is less congested. This suggests a bottleneck is created in the merge lane when it becomes congested. Further evidence of this will be shown in the collision results.

B. Collisions

The average number of collisions per simulation for PSO-LCM over all vehicle densities and percentages of merging vehicles is shown in Figure 3. SUMO-LCM and No-LCM are collision-free due to the underlying Krauss model influencing safe gaps between vehicles on the merge and through-lanes. The statistical results for average collisions on PSO-LCM are shown in Table IV. We note that PSO-LCM is incentivized to choose solutions that result in the minimum amount of collisions, as shown by PSO-LCM having 0-1.5 average collisions for 10 vehicles and 20 vehicles having no more than an average of 3.8 collisions. There is an interesting behavior with 50 vehicles and 80% of them merging, where there is an average of 11.5 collisions per simulation. This is due to the heavy congestion of vehicles waiting to merge. Currently, the PSO-LCM rules provide little guidance for vehicles waiting to merge in terms of speed and spacing, so rules would need to be added to address this issue in the future.

Figure 4 shows how the swarm average number of collisions evolved over 100 iterations for 10 vehicles. This shows that PSO-LCM is able to learn solutions with no collisions for 10 vehicles, given enough training.

For 50 vehicles, PSO-LCM has a more difficult time finding solutions with minimal collisions due to an increase in congestion in the simulation. Figure 5 shows how the swarm average collisions evolved over 100 iterations for 50 vehicles. Overall, PSO-LCM was able to learn merging strategies that resulted in 0-2 average collisions over the swarm (10 simulations). A similar behavior has been seen for 20 vehicles, where PSO-LCM is able to learn solutions averaging 0-1 collisions per swarm.

C. Loop Speed

Figure 6 shows the average loop speed of vehicles when driving on the loop and merge lane for PSO-LCM, SUMO-LCM and No–LCM. The speed limit on the loop is 30 m/s. This shows both SUMO-LCM and No-LCM result in an average speed of 27 m/s. At 20% of vehicles merging, PSO-LCM also has an average loop speed of 27 m/s, when the through-lane is more congested. However, for 50% and 80% of vehicles merging, the vehicles slow down to just under 26 m/s since the merge lane is more congested. This slowdown is slightly greater for 20 and 50 vehicle densities, as the congestion increases more.

D. Merge Speed

The last metric collected was the average merge speed, which measures the speed of the vehicles on the exit and entrance ramps before vehicles reach the merge lane. The speed limit in these sections is 15 m/s. SUMO-LCM and No-LCM hover around an average of 12.8–13 m/s. PSO-LCM also has an average merge speed of 13 m/s when 20% of vehicles are merging, but as the merge lane becomes more congested, this drop to about 12 m/s. This also shows that as more vehicles are waiting to merge, the more likely they are to jam with reduced speeds. This slowdown of vehicles on the merge ramp is enhanced with 20 and 50 vehicles on the road.

VI. DISCUSSION

Overall, PSO-LCM shows some promise in its ability to find safe merging strategies in many cases, but further improvements need to be made to ensure optimal traffic flow and collision free driving. We need to stress that PSO-LCM is much simpler than the SUMO-LCM model, since neither tactical nor regulatory behavior were built into the model for this scenario. As a result, this study emphasized potential gaps in the rules, such as defining instructions for speed and vehicle spacing on the merge ramp and lane before a merge maneuver.

Hypothesis 1 stated, "the use of PSO-LCM on a highway on-ramp can be optimized to allow for a higher merge rate than the default SUMO lane change model." The initial results do not support this hypothesis directly; however, a final verdict on this remains to be seen for now. Even so, PSO-LCM performed comparably to SUMO-LCM in most cases, particularly for 10 vehicles and 50 vehicles when 20% and 50% are merging. This demonstrates the feasibility of PSO-LCM to learn merging strategies when the through-lane is the most congested. Since the through-lane is not as congested when 80% of vehicles are merging, these results suggest the source of degraded performance is the increase in congestion on the merge lane and ramp, resulting in more collisions on the merge lane. More

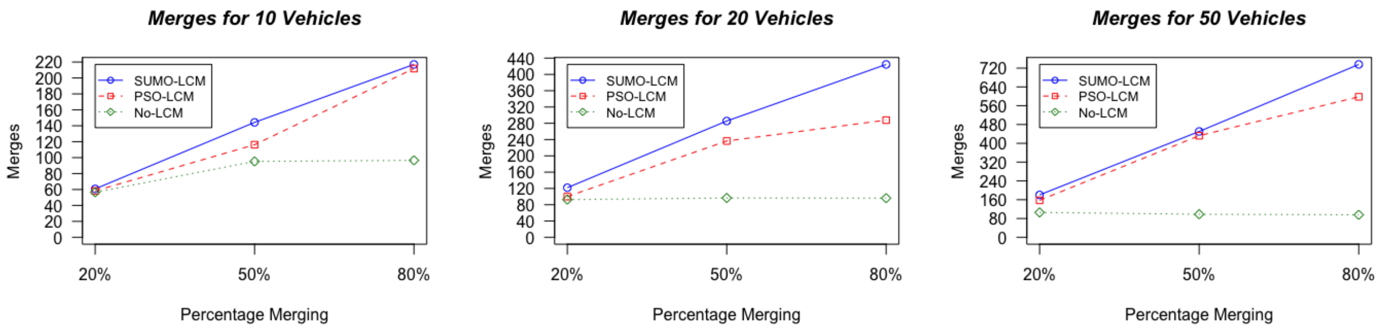

Fig. 2. Average number of merges ($\#merges$) per simulation for 10, 20 and 50 vehicles.

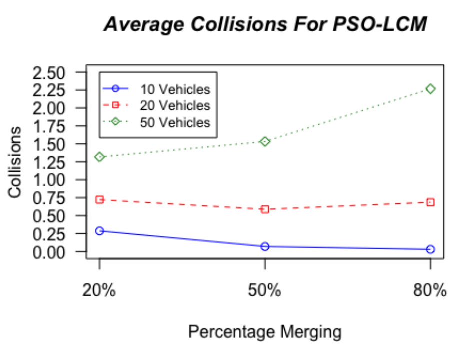

Fig. 3. Average number of collisions ($\#collisions$) per simulation for PSO-LCM with 10, 20, and 50 vehicles

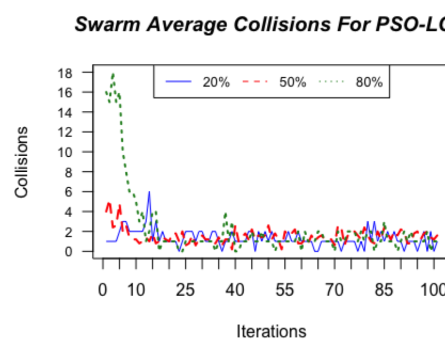

Fig. 5. Swarm average collisions per simulation for PSO-LCM with 50 vehicles.

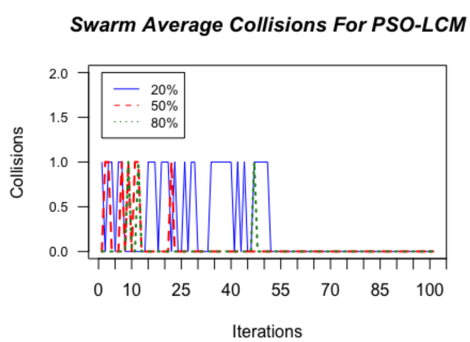

Fig. 4. Swarm average collisions per simulation for PSO-LCM with 10 vehicles

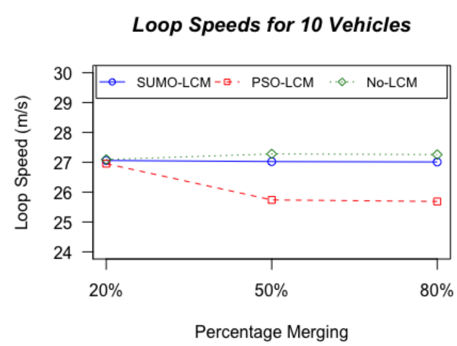

Fig. 6. Average loops speeds for 10 vehicles. Speed limit is $30m/s$

attention will need to be placed on the behavior of vehicles waiting to merge moving forward.

In terms of Hypothesis 2, which predicted PSO-LCM would minimize collisions at or around the merge lane, the results provided partially support this claim. Average collisions are 0 for 10 vehicles, between 0-1 for 20 vehicles, and between 0-2 for 50 vehicles. This suggests that this hypothesis could be verified if more attention is given to collisions occurring on the merge ramp while waiting to merge. The increase in collisions in these cases also coincides with decreased average loop and merge speeds. If collisions are reduced in these cases, average speeds will likely increase, resulting in increased flow and a higher overall successful merge rate.

VII. Conclusions and Future Work

A PSO lane change model was developed and compared to the default SUMO lane change model, as well as to a set of baseline experiments when no lane change model was activated. Three vehicle densities and three different ratios of

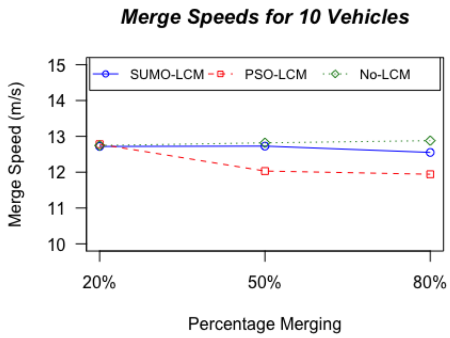

Fig. 7. Average merge speeds for 10 vehicles. Speed limit is $15m/s$

merging to through-lane vehicles were tested. In many cases, PSO-LCM performed comparably to SUMO-LCM in terms of the average number of merges per simulation. The model also produced little to no collisions in most cases, requiring additional rules for behavior on the merge lane.

This work is another step in building a swarm-based vehicle control mechanism for autonomous vehicles in a mixed environment with human operated vehicles. Our previous work [4] focused on more general control mechanisms using PSO and is therefore not comparable to the results provided here. This problem is also more complex due to the rule-based nature, where particles capture the learned parameters to the rules, thus allowing for additional complexity in the underlying learned behaviors.

There are many opportunities to further advance this work. The first would be to incorporate merge ramp rules to reduce collisions further. More rules also need to be added for determining behavior while waiting to merge, which may further reduce the number of collisions on the merge lane. One way of doing this would be to add a car-following rule-based component to the model focused on general behavior throughout the road network. Expanding experiments to test the limitations and scalability of our model would also be important to determine application feasibility. We would also like to compare the results of PSO against other variants of PSO and also study the effects of weather and other environmental factors on our model.

Further work will also focus on increasing problem complexity. For example, a logical next step is to expand the number of lanes on the loop and add rules representing tactical and regulatory behavior. We also intend to incorporate vehicle-to-vehicle communication by pairing ns-3 [19] with TraCI, where control decisions will be based on both predefined rules as well as state information from neighboring vehicles. This would enable more cooperative behavior that has been shown to improve traffic conditions [12]. Finally, we plan to explore other traffic scenarios, such as multiple intersection control, a more heterogeneous traffic makeup with various vehicle types, and lane consolidation (e.g. bridges and tunnels).

The ultimate direction of this work is to develop a distributed, cooperative, co-evolutionary approach. More specifically, we intend to apply a factored evolutionary algorithm FEA [20] where vehicle swarms overlap according to neighbor interactions to further optimize vehicle control and vehicle-to-vehicle communication and interaction.

References

[1] J. Erdmann, *SUMO's Lane-Changing Model*, ser. Modeling Mobility with Open Data. Cham: Springer International Publishing, Mar 12, 2015, pp. 105–123.

[2] Z. Zheng, S. Ahn, D. Chen, and J. Laval, "Freeway traffic oscillations: microscopic analysis of formations and propagations using wavelet transform," *Procedia-Social and Behavioral Sciences*, vol. 17, pp. 702–716, 2011.

[3] R. C. Eberhart, Y. Shi, and J. Kennedy, *Swarm intelligence*. Elsevier, 2001.

[4] N. Wiesner, J. Sheppard, and B. Haberman, "Autonomous vehicle control using particle swarm optimization in a mixed control environment," in *IEEE Symposium Series on Computational Intelligence*, Dec 2020, pp. 2877–2884.

[5] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, "Microscopic traffic simulation using sumo," in *IEEE International Conference on Intelligence Transportation Systems*, 2018, pp. 2575–2582.

[6] S. Krauß, P. Wagner, and C. Gawron, "Metastable states in a microscopic model of traffic flow," *Physical Review E*, vol. 55, no. 5, pp. 5597–5602, 1997.

[7] X. Ma, X. Hu, T. Weber, and D. Schramm, "Traffic simulation of future intelligent vehicles in Duisburg City inner ring," *Applied sciences*, vol. 11, no. 29, p. 29, Dec 1, 2021.

[8] L. Bieker-Walz, M. Behrisch, M. Junghans, and K. Gimm, "Evaluation of car-following-models at controlled intersections," in *European Simulation and Modelling Conference*, Oct 2017.

[9] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical review E*, vol. 62, no. 2, p. 1805, 2000.

[10] J. E. Naranjo, F. Serradilla, and F. Nashashibi, "Speed control optimization for autonomous vehicles with metaheuristics," *Electronics*, vol. 9, no. 4, p. 551, 2020.

[11] F. Olivas, F. Valdez, P. Melin, A. Sombra, and O. Castillo, "Interval type-2 fuzzy logic for dynamic parameter adaptation in a modified gravitational search algorithm," *Information Sciences*, vol. 476, pp. 159–175, 2019.

[12] A. Kesting, M. Treiber, and D. Helbing, "General lane-changing model mobil for car-following models," *Transportation Research Record*, vol. 1999, no. 1, pp. 86–94, 2007.

[13] U. Khan, P. Basaras, L. Schmidt-Thieme, A. Nanopoulos, and D. Katsaros, "Analyzing cooperative lane change models for connected vehicles," in *International Conference on Connected Vehicles and Expo*, Nov 2014, pp. 565–570.

[14] J. Nie, J. Zhang, W. Ding, X. Wan, X. Chen, and B. Ran, "Decentralized cooperative lane-changing decision-making for connected autonomous vehicles," *IEEE access*, vol. 4, pp. 9413–9420, 2016.

[15] Y. Liu, X. Wang, L. Li, S. Cheng, and Z. Chen, "A novel lane change decision-making model of autonomous vehicle based on support vector machine," *IEEE access*, vol. 7, pp. 26 543–26 550, 2019.

[16] C. Huang, F. Naghdy, and H. Du, "Model predictive control-based lane change control system for an autonomous vehicle," in *IEEE Region 10 Conference*, Nov 2016, pp. 3349–3354.

[17] J. Hu, L. Kong, W. Shu, and M.-Y. Wu, "Scheduling of connected autonomous vehicles on highway lanes," in *IEEE Global Communications Conference (GLOBECOM)*, 2012, pp. 5556–5561.

[18] A. Wegener, M. Piórkowski, M. Raya, H. Hellbrück, S. Fischer, and J.-P. Hubaux, "Traci: an interface for coupling road traffic and network simulators," in *11th Communications and Networking Simulation Symposium*, 2008, pp. 155–163.

[19] G. F. Riley and T. R. Henderson, "The ns-3 network simulator," in *Modeling and tools for network simulation*. Springer, 2010, pp. 15–34.

[20] S. Strasser, J. Sheppard, N. Fortier, and R. Goodman, "Factored evolutionary algorithms," *IEEE Transactions on Evolutionary Computation*, vol. 21, no. 2, pp. 281–293, April 2017.