

CHATANI: LANGUAGE-DRIVEN MULTI-ACTOR ANIMATION GENERATION IN STREET SCENES

Anonymous authors

Paper under double-blind review

ABSTRACT

Generating interactive and realistic traffic participant animations from instructions is essential for autonomous driving simulations. Existing methods, however, fail to comprehensively address the diverse participants and their dynamic interactions in street scenes. In this paper, we present ChatAni, the first system capable of generating interactive, realistic, and controllable multi-actor animations based on language instructions. To produce fine-grained, realistic animations, ChatAni introduces two novel animators: PedAnimator, a unified multi-task animator that generates interaction-aware pedestrian animations under varying task plans, and VehAnimator, a kinematics-based policy that generates physically plausible vehicle animations. For precise control through complex language, ChatAni employs a multi-LLM-agent role-playing approach, using natural language to plan the trajectories and behaviors of different participants. Extensive experiments demonstrate that ChatAni can generate realistic street scenes with interacting vehicles and pedestrians, benefiting tasks like prediction and understanding. All related code, data, and checkpoints will be open-sourced.

1 INTRODUCTION

Realistic and diverse street scene animations of traffic participants serve in various fields, including game, movie, especially autonomous driving simulations Caesar et al. (2020); Sun et al. (2020); Xiao et al. (2021); Dosovitskiy et al. (2017). The participants' animation reflect their temporal dynamics and motion. Different behaviors and interactions of traffic participants constitute the temporal events within the scene, which contain crucial temporal information. Perceiving, understanding, and predicting the trajectories and motions of traffic participants Sun et al. (2020) composed of animations is a key step in driving system and making decisions within them. The animation incorporates the participants' motion representation, thereby serving as a prior to drive the rendering engine or diffusion model for the final simulation. Therefore, constructing diverse and high-quality animation is a key component of driving simulators. However, how to generate pedestrian and vehicle animations with interactive attributes remains an unsolved problem.

In driving simulator like Dosovitskiy et al. (2017), animations are manually generated as templates, which is inefficient, not diverse, and especially not realistic. Some existing works adopt learning-based approaches to address physical and dynamic infeasibility or to generate more reasonable behaviors, thereby improving the quality of animations. However, all of these methods overlook the interactivity among participants in street scenes. The generation process and the occurrence of interactions typically exist on two levels: the high-level, which is concerned with trajectory planning, and the low-level, which focuses on detailed

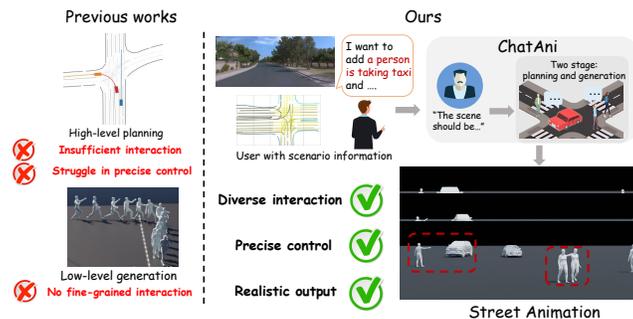


Figure 1: ChatAni generates interactive and realistic multi-actor street animation by language.

054 animations. Existing works have not comprehensively addressed both aspects and the interactive at-
055 tributes inherent in them. At the high level, LCTGen Tan et al. (2023) and CTG++Zhong et al.
056 (2023a) use language to generate vehicle trajectories, but they do not consider different participants,
057 including pedestrians, and thus lack pedestrian-involved high-level interactive attributes. At the low
058 level, PacerRempe et al. (2023) and Pacer+ Wang et al. (2024) generate fine-grained pedestrian ani-
059 mations but still do not account for potential low-level interactions between multiple pedestrians or
060 between pedestrian and vehicle.

061 We introduce ChatAni, the first system achieving interactive and realistic language-driven multi-
062 actor street animation generation, as shown in Fig. 1. ChatAni simultaneously introduces interaction
063 designs to both low-level and high-level aspects, ensuring interactivity, realism, and controllability.

064 In ChatAni, we design specialized low-level animation generators tailored to the distinct charac-
065 teristics of pedestrian and vehicle motions. For interactive and realistic pedestrian animation, we
066 introduce PedAnimator, a unified multi-task framework that generates physically plausible anima-
067 tions via physics-driven control and diverse input signals. PedAnimator is the first method to achieve
068 physical interactions between pedestrians and pedestrian-vehicle through a novel interaction train-
069 ing strategy that integrates general observations and goals, enabling diverse interactions with mini-
070 mal adjustments. It supports control over both trajectory and body motion, with policy unification
071 achieved through a task-masking mechanism that enables generation under a single policy. Further-
072 more, hierarchical control and body-masked adversarial motion priors are incorporated, introducing
073 priors into both action and reward spaces. These not only support multiple control objectives but
074 also enhance the realism and human-likeness of the generated animations.

075 To generate physically feasible vehicle animations, we introduce VehAnimator, a kinematic-based
076 control policy generating physically feasible vehicle animations. It converts planned raw trajectories
077 into kinematically compliant animations through a dynamic model, eliminating unrealistic artifacts
078 like tail swinging/drift. The framework ensures interactivity via obstacle-aware training enabling
079 inherent collision avoidance. Temporal-aware designs improve precision and temporal consistency.

080 To achieve language-driven high-level control and interactive attributes, ChatAni models each traf-
081 fic participant as an agent and employs a multi-LLM-agent role-playing approach. Agents organize
082 their own information based on requirements and communicate with other agents, ultimately invoking
083 tools to output planning, which are then executed by PedAnimator and VehAnimator. Benefit
084 from the powerful understanding capability and prior knowledge of LLM, the design enables precise
085 control over the entire scene animation generation through natural language, while the communica-
086 tion process between the agents facilitates the planning of high-level interactive attributes.

087 ChatAni achieves three core features: *interactivity*, *realism*, and *controllability*. For interactivity,
088 low-level animators incorporate interactive attributes, while multi-agent LLM role-playing handles
089 high-level planning. For realism, PedAnimator employs hierarchical control with body-masked
090 adversarial motion priors to generate human-like motions, while VehAnimator simulates real vehicle
091 physics and kinematics. For controllability, LLM agents adhere strictly to language commands,
092 with animators executing actions with precision. The final output animations are suitable for driving
093 simulators, based on either rendering engines or diffusion models.

094 Our contributions include: (1) the first language-driven system for multi-actor street scene anima-
095 tion generation; (2) PedAnimator enabling interactive pedestrian animations with varied control
096 signals; (3) VehAnimator creating physically plausible animations from trajectories; (4) a multi-
097 LLM-agent framework generating language-controlled interactive high-level plans; (5) com-
098 prehensive experiments validating ChatAni’s command-aligned animations with enhanced interac-
099 tivity/authenticity and benefits to the prediction and understanding tasks.

101 2 RELATED WORKS

102
103 **Human animation generation.** Human animation generation can be broadly divided into kine-
104 matic and physics-based approaches. For kinematics-based generation, transformer-based meth-
105 ods Athanasiou et al. (2022); Guo et al. (2024; 2022) and diffusion models Zhang et al. (2022; 2023)
106 are used to generate the corresponding animation based on language inputs. Recent methods Shafir
107 et al. (2023); Xie et al. (2023); Wan et al. (2023) introduce further control conditions. However,
these approaches do not account for physical constraints. For physics-based generation, existing

work such as Peng et al. (2021; 2022); Won et al. (2022); Luo et al. (2023; 2024) achieves predefined tasks with plausible animations. Tessler et al. (2023); Juravsky et al. (2022); Bae et al. (2023); Xu et al. (2023a;b) extend the animation content with different designs. Rempe et al. (2023); Wang et al. (2024) focus on the animation of pedestrians in the street scene. However, these methods do not involve the interaction behaviors between multi-pedestrians or pedestrian-vehicle. Our PedAnimator considers the interaction behaviors and is trained as a unified policy for multiple scenarios also including following, imitation.

Vehicle traffic generation. In industry, vehicle traffic is usually generated by software tools Chen et al. (2022); Queiroz et al. (2019); Fremont et al. (2019); Jesenski et al. (2019). Some research works Bergamini et al. (2021); Tan et al. (2021); Feng et al. (2023); Rempe et al. (2022) focus on unconditioned generation to simplify the generation process. Recent works on vehicle traffic generation Suo et al. (2021); Zhong et al. (2023b;a); Tan et al. (2023); Lu et al. (2024); Ding et al. (2023) introduce different conditions to the process. However, these works generally do not directly consider physical and kinematic constraints for the traffic and lack sufficient interaction. [Several early studies, such as those by Li et al. \(2017\); Lin et al. \(2016\); Wang et al. \(2018\), consider certain potential interaction behaviors in traffic. However, these works are unable to achieve effective control through means such as linguistic commands, and they also lack sufficiently granular modeling of interactions involving pedestrians, such as hailing a taxi or vehicle-pedestrian collisions.](#) And these language-controlled methods cannot achieve precise control over different vehicles or other participants. Our LLM-agents planning utilizes the language understanding capability of LLMs to achieve precise control of participants considering interactive information, and VehAnimator generates the final physical plausible vehicle animation.

Large language models and agents. Recently, numerous large language models Touvron et al. (2023); Liu et al. (2024); Bai et al. (2023); Achiam et al. (2023) have been proposed and released. Many works have used these models by fine-tuning them Hu et al. (2021); Qiu et al. (2023) or integrating them with relevant tools to build LLM-based agents Liu et al. (2023); Wu et al. (2023). These agents have been applied to a wide range of downstream tasks Hong et al. (2023); Zhou et al. (2023); Li et al. (2024b;a); Leng & Yuan (2023); Shen et al. (2024). In this paper, we explore the application of agents in traffic simulation, using them to interact and perform trajectory and behavior planning within traffic scenarios.

3 METHOD

ChatAni is the first system to enable interactive, realistic, language-driven multi-actor animation generation for street scenes, composed of three components: (1) PedAnimator—a unified multi-task controller that produces physically plausible pedestrian animations by integrating interaction signals, trajectory tracking, and upper-body motion. It incorporates interaction task training, task-masking with embeddings for policy unification, Body-masked Adversarial Motion Priors, and hierarchical control to enhance animation quality; (2) VehAnimator—a vehicle animation generator that translates high-level raw trajectories into physically feasible vehicle motions via kinematic constraints and action-to-state control; and (3) a Multi-LLM-Agent Role-Playing Framework—a language-driven planner leveraging LLMs’ natural language understanding to construct role-specific interactive agents with inter-agent communication, ensuring semantically consistent, language-aligned high-level plans. During execution, traffic participants are initialized as LLM agents that collaboratively establish the scene context through communication. The generated high-level plans are then dispatched to PedAnimator or VehAnimator, which synthesize low-level animations under physical constraints. These outputs are composited into final animations that jointly satisfy physical plausibility and linguistic specifications. The system overview is shown in 2.

3.1 PEDANIMATOR

PedAnimator (Fig. 3) is a unified multi-task controller generating interaction-aware pedestrian animations through diverse control signals (interactions, trajectories, upper-body motions). Multi-pedestrian interactions constitute critical event elements in traffic scenarios yet remain unaddressed by existing methods. PedAnimator employs a dedicated interaction training pipeline with universal observation-action frameworks, enabling physical plausibility across interaction tasks via minor objective adjustments. The system also integrates task-masking mechanisms with embeddings for

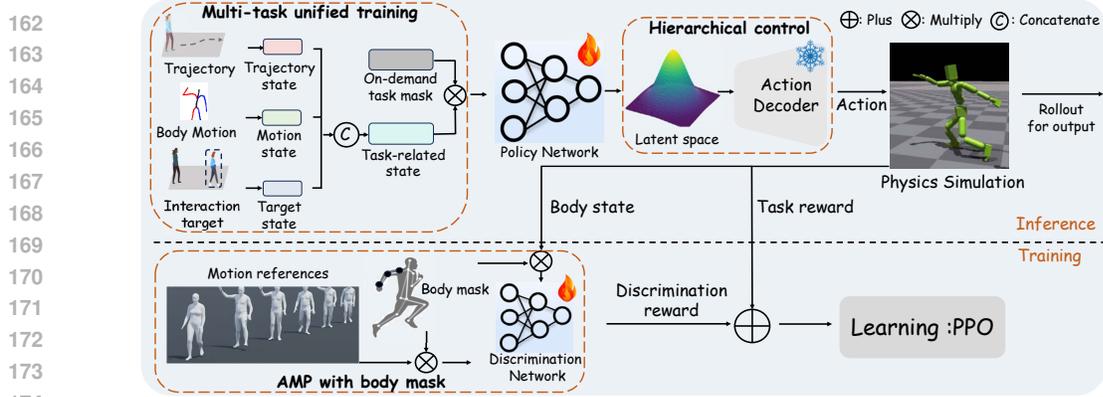


Figure 3: Pedestrian animator (PedAnimator) framework. PedAnimator takes interaction behaviors into account. With multi-task unified training, PedAnimator achieves unified control over various tasks, including following, imitation and interaction. Hierarchical control and AMP with body mask provide prior to action and reward space, improving the realism of animations.

multi-task unification, body-masked Adversarial Motion Priors (AMP), and hierarchical control in reward/action spaces to ensure control fidelity while producing human-like animations that fulfill physical and visual requirements.

The control process is defined by Markov decision process $\mathcal{M}^p = \{S^p, \mathcal{A}^p, \mathcal{T}^p, \mathcal{R}^p, \gamma^p\}$, where the elements represent states, actions, transition process, reward, and the discount factor. Goal-conditioned reinforcement learning (RL) is adopted for training. The transition process is implemented by the physics engine. Others are detailed below.

States and Multi-task Unified Training. Our control policy is designed to handle multiple distinct tasks, requiring specialized processing to achieve a unified policy capable of addressing them. The tasks are categorized into three main groups: trajectory following, single-agent behavior specification, and multi-pedestrian interaction. For single-agent behavior, the LLM generates text descriptions, which are then processed by a Text2Motion model (e.g., Mo-Mask Guo et al. (2024)) to produce upper-body motion for the policy to replicate, specifying single-agent behavior. Multi-agent interaction tasks are trained based on predefined types, with the LLM classifying interaction behaviors, enabling the animator to execute specified interactions.

Task-related states consist of trajectory slices \mathcal{S}_{traj}^p , motion to be imitated \mathcal{S}_{mo}^p , and states of interacting targets \mathcal{S}_{tar}^p . \mathcal{S}_{traj}^p includes the K future steps of the planned trajectory, while $\mathcal{S}_{mo}^p = \text{concat}(\hat{j}_{pos}, \hat{j}_{rot}, \hat{j}_{vel}, \hat{j}_{\omega})$ represents joint position \hat{j}_{pos} , rotation \hat{j}_{rot} , velocity \hat{j}_{vel} , and angular velocity \hat{j}_{ω} , with optional joint masking to focus on relevant parts, such as the upper body. $\mathcal{S}_{tar}^p = \text{concat}(r_{pos}, r_{rot}, r_{vel}, r_{\omega}, r_{bbox}, p_{inter}, e_{inter})$ includes the root position r_{pos} , root rotation r_{rot} , root velocity r_{vel} , angular velocity r_{ω} , bounding box vertices r_{bbox} , interaction contact position $p_{interaction}$, and an interaction embedding e_{inter} as a one-hot vector specifying the interaction characteristic. The states of interacting targets remain consistent across behaviors, with only minor adjustments to control signals, allowing for diverse interactions under a unified framework. The interactive behavior is uniformly modeled. By controlling variables such as the specific interaction points and the intensity of interaction, and by integrating trajectory following and motion imitation,

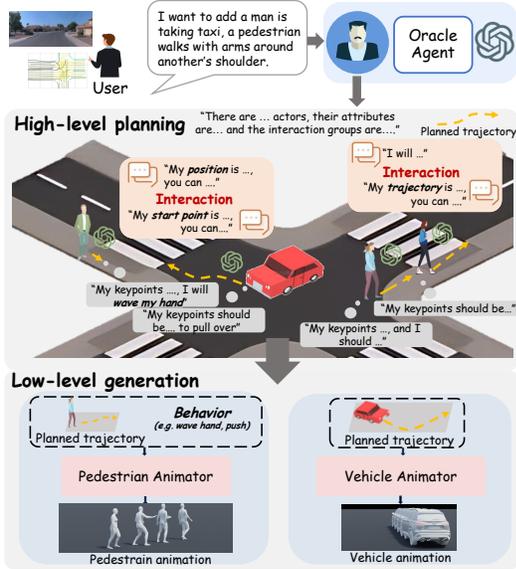


Figure 2: ChatAni adopts multi-LLM-agent role-playing for precise high-level planning. Two specialized animators are designed for low-level animation.

a diverse range of interactive behaviors can be achieved. In addition to task-related states, the final observation includes humanoid proprioception Wang et al. (2024) S_{prop}^p , capturing the internal states of the humanoid. The final state is represented as $S^p = \text{concat}(S_{traj}^p, S_{mo}^p, S_{tar}^p, S_{prop}^p)$. To enable specific parts during training or testing, we introduce a task masking mechanism. During training, tasks for each episode are sampled, and the corresponding task-related states are unmasked while others are masked. This prepares only the relevant environment, and rewards are calculated based on these active tasks. During inference, relevant states are activated while others are masked, ensuring task-specific execution. This approach ensures that the unified policy can handle all tasks without confusion. After sufficient training, the policy performs at the level of individually trained policies while being capable of executing multiple non-conflicting tasks in a single run.

Action hierarchical control. Pedestrian action spaces are typically modeled using a proportional-derivative (PD) controller at each degree of freedom (DoF), but such spaces lack inherent priors, often leading to locally unrealistic actions for completing specific tasks. To address this, we apply a hierarchical action control space from PULSE Luo et al. (2024). The policy network first outputs to a pretrained latent space, obtaining a latent feature f_{action} , which is then decoded by a pretrained decoder into control signals. The pretrained latent space, equipped with the corresponding decoder, ensures that the decoded output distribution closely matches the input data from pretraining, thereby providing the action space with a real-world action prior.

Reward design. The reward consists of two main components. The first is the discrimination reward R_{disc} used to implement AMP Peng et al. (2021), which employs a discriminator to encourage the policy to generate output that aligns with movement patterns observed in a dataset of human-recorded data clips. The second component is the task-related reward R_{task} , designed to motivate the policy to accomplish specific tasks and execute high-level planning. The detailed task reward designs can be found in appendix S2.2. Similar to Luo et al. (2023), we implement early termination if an excessive root distance error or joint distance error occurs, and fail-state recovery task for robustness. The recovery task also plays an important role in low-level interactions between pedestrian and vehicle.

AMP with body mask and warm-up training. We employ Proximal Policy Optimization (PPO) Schulman et al. (2017) for overall training optimization. However, directly optimizing for interaction tasks presents certain challenges: without AMP, the lack of a discrimination reward makes it difficult to achieve human-like animation results. Conversely, proper reference data clips needed for specific tasks are usually complex and not easily available. Using reference data clips from the following or imitation tasks can lead to misalignment with the actual requirements for interaction, resulting in failures during interaction tasks.

We address these challenges through two strategies: (i) A two-phase approach: initial AMP-free training establishes task-completion foundations for subsequent AMP-enhanced optimization that preserves functionality while improving realism; (ii) AMP body masking that excludes interaction-focused joints (e.g., arms) during discriminator calculations, enabling AMP-driven natural motion for non-interactive body parts while maintaining task-specific joint flexibility. These techniques collectively enable visually realistic animations with effective interaction task performance.

3.2 VEHANIMATOR

VehAnimator, as shown in Fig. 4, translates raw vehicle trajectory planning into physically and dynamically feasible animations. Although trajectory data alone could theoretically represent rigid-body vehicle motion, raw trajectories typically lack necessary kinematics constraints, leading to unrealistic phenomena such as tire slippage and drifting when applied directly. As a physics-aware controller, VehAnimator processes trajectory inputs to generate control signals, then applies vehicle animation equations for state

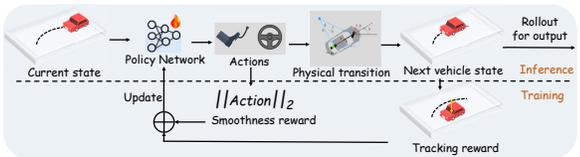


Figure 4: Vehicle animator (VehAnimator) framework. VehAnimator adopts goal-conditioned RL based on physical transition of real vehicle. Combining with history-aware design, VehAnimator generates realistic vehicle animation under planned trajectory. This imple-

270 mentation ensures kinematics validity in the outputs while enhancing control accuracy and temporal
 271 consistency through historical-aware action-observation designs. The control process is also mod-
 272 eled as a Markov decision process defined by $\mathcal{M}^v = \{\mathcal{S}^v, \mathcal{A}^v, \mathcal{T}^v, \mathcal{R}^v, \gamma^v\}$, with goal-conditioned
 273 RL for training.

274 **History-aware states.** The VehAnimator incorporates historical information into its states, along-
 275 side its current state, to improve the temporal consistency of the policy. The vehicle states consist of
 276 the planned trajectory segment $\hat{\mathbf{P}}^v$, temporal velocity \mathbf{V}^v , and dynamic parameters Θ^v . \mathbf{P}^v is a slice
 277 of the planned trajectory in the vehicle’s coordinate system over the current and adjacent τ^v frames.
 278 \mathbf{V}^v represents the vehicle’s centroid velocity over τ^v frames. Θ^v includes inherent vehicle param-
 279 eters such as L (vehicle length), W (vehicle width), l_f (front overhang), and l_r (rear overhang).
 280 These parameters provide prior information that influences the dynamic transition process.

281 **Vehicle actions.** To accurately simulate real vehicle actions and maintain temporal consistency, the
 282 vehicle action space \mathcal{A}^v is defined by the delta steering angle $\Delta\delta$ and scalar acceleration a , which
 283 are the two most direct controls affecting vehicle movement in actual driving.

284 **Vehicle transition function.** We employ the bicycle model to model the vehicle dynamic transition
 285 process. Let x and y denote the vehicle’s coordinates, v represent the scalar velocity, and ϕ indi-
 286 cate the vehicle’s orientation. Utilizing the inherent parameters from the states, the vehicle’s state
 287 transition process can be expressed as:

$$289 \dot{x} = v \cos(\phi + \beta), \dot{y} = v \sin(\phi + \beta), \beta = \arctan\left(\frac{l_r}{l_f + l_r} \tan(\delta)\right), \dot{\phi} = \frac{v}{l_f + l_r} \cos(\beta) \tan(\delta).$$

292 Where β denotes tire slip angle and $\dot{\cdot}$ denotes derivation. The bicycle model effectively simulates the
 293 dynamic state changes of the vehicle determined by vehicle actions.

294 **Reward and training.** The reward focuses on two aspects: following the planned trajectory and
 295 smoothness of the results. Therefore, the reward consists of two components: $R_{pos}^v = -\|\hat{p}_t - p_t\|_2$
 296 for following the planned trajectory and $R_{act}^v = -(c_\delta \|\Delta\delta\|_2 + c_a \|a\|_2)$ for smoothness, with c_δ and
 297 c_a are two coefficients to balance the two different units. Training uses TD3 Fujimoto et al. (2018)
 298 to maximize the accumulated discounted reward. Note that: (i) actions can be further smoothed
 299 with temporal filtering: $\mathcal{A}_t^v = \alpha \mathcal{A}_{t-1}^v + (1 - \alpha) \mathcal{A}_{policy}^v$, where \mathcal{A}_t is the action taken in timestep
 300 t and \mathcal{A}_{policy} is the action directly output from policy network; (ii) obstacles can be considered by
 301 concatenating their positions and radius in the state and adding $R_{obs}^v = \frac{\epsilon}{L_o}$ to the reward if $L_o < D$,
 302 where L_o is the distance to obstacle, ϵ is a coefficient and D is a threshold.

304 3.3 MULTI-LLM-AGENT ROLE-PLAYING SYSTEM DESIGN

305 PedAnimator and VehAnimator facilitate precise low-level control for pedestrian and vehicle ani-
 306 mations. For language-based holistic scene control, we developed a multi-LLM-agent role-playing
 307 framework that generates high-level planning for all participants from user input, offering control
 308 signals for low-level animation. The framework includes agents with LLM-based reasoning and tool
 309 interfaces, divided into two types: (1) Oracle Agent processes user language inputs to create global
 310 scene context and scheduling instructions; (2) Actor Agents represent individual traffic participants,
 311 initialized with Oracle Agent information, and engage in inter-agent communication to collabora-
 312 tively develop high-level plans. This design utilizes the strong language understanding and domain
 313 knowledge of LLM to accurately interpret user input. Actor Agents dynamically adjust planning
 314 based on participant types (e.g., pedestrians, vehicles), meeting diverse behavioral needs through
 315 specific tool functions while ensuring scene-level coordination. See Fig. 2 for details.

316 3.3.1 ORACLE AGENT

318 The Oracle Agent decomposes complex user instructions into participant-specific actions, ensuring
 319 precise interpretation and schedule generation. Its LLM component is prompted with role def-
 320 initions, requirements, and few-shot examples, and outputs initialization information, interaction
 321 groups, and schedules in natural language. We implement a structured output tool that parses these
 322 linguistic outputs into executable predefined data structures. Oracle Agent enables the system to
 323 process composite abstract instructions while enhancing operational clarity and granularity. See
 more details in the appendix S4.

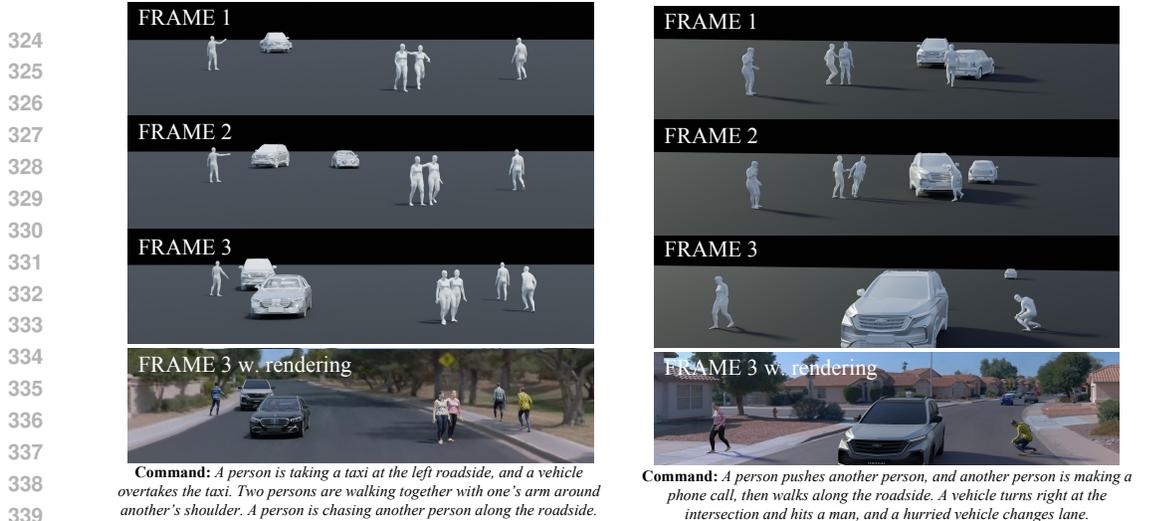


Figure 5: System results under complex and composite commands, with realistic animation output and diverse interaction information between pedestrian-vehicle, multi-pedestrian, multi-vehicle.

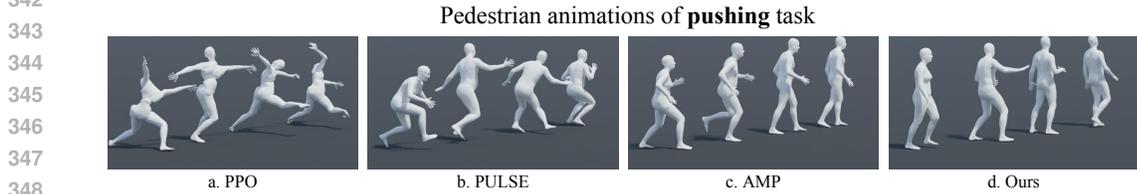


Figure 6: Visualization comparison of pushing task. PedAnimator complete the task with high visual quality.

3.3.2 ACTOR AGENT

351
352
353
354
355
356
357
358
359
360

Each Actor Agent corresponds to a participant in the scene, initialized with information from Oracle Agent, including agent type, trajectory, and behavioral descriptions. These agents communicate according to the Oracle Agent schedule to collect interactive information and formulate final high-level planning. The planning process incorporates map data, where the scene map is represented as a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$. \mathcal{N} denotes lane sections, each containing a point set and metadata such as orientation and driving type (e.g., straight, turn, lane/boundary). Edges \mathcal{E} encode relationships between lane sections (e.g., adjacency, connectivity). The planning adopts a keypoints-based methodology, synthesizing actor interactions to refine the final trajectory.

361
362
363
364
365
366
367
368

Keypoints-based trajectory planning. Trajectories are defined by keypoints interpolated into continuous paths. The Actor Agent’s LLM component determines the required keypoint count and their generation logic based on initialization information. Keypoints originate from two sources: (1) Map retrieval: extracted from the map \mathcal{G} by analyzing attributes (e.g., lane orientation, driving type) and topological relationships (via edges \mathcal{E}) between nodes \mathcal{N} ; (2) Interaction derivation: determined through inter-agent communication, where dependencies on other participants’ states influence keypoint placement. This hybrid approach ensures trajectory coherence with both environmental constraints and multi-agent interactions.

369
370
371
372
373

Inter-agent communication. Communication allows agents to acquire information from others to generate interaction-dependent keypoints. Predefined interaction groups, through the Oracle Agent, guide the process. The LLM identifies the required parameters (e.g., start/end points, trajectories) from interacting agents. Agents exchange requested data, which are processed by tools to generate keypoints. This mechanism supports interactive trajectory generation.

374
375
376
377

Bézier curve interpolation aggregates keypoints to generate complete planning trajectories, with three operational notes: (1) interpolated trajectories serve as physics-agnostic plans for animators; (2) static actors require two orientation-defining keypoints; (3) pedestrian agents have LLM-inferred behavioral instructions (e.g., "wave hand while walking") from initialization data—derived from explicit commands/implicit semantic context—formatted as executable text directives.

4 EXPERIMENTS

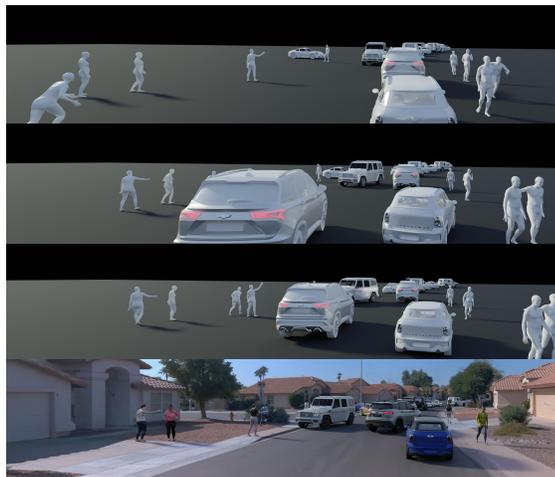
4.1 IMPLEMENTATION DETAILS AND DATASET

PedAnimator uses Isaac Gym Makoviychuk et al. (2021) as physics simulator, AMASS Mahmood et al. (2019) providing reference data, employing a SMPL Loper et al. (2023) model as the simulated entity. LLM components employ GPT-4 Achiam et al. (2023) API. The final rendering adopts pipeline from ChatSim Wei et al. (2024). PedAnimator and VehAnimator are trained independently in separate environments and models. Detailed settings, configurations, more experiments, video results, and part of codes can be found in the appendix and supplementary materials.

4.2 SYSTEM RESULTS

4.2.1 INTERACTIVE, REALISTIC AND CONTROLLABLE

We showcase representative keyframes from two simulated scenarios in Figure 5, each containing various traffic participants generated according to commands. The results demonstrate three critical capabilities: (i) interaction is thoroughly depicted, including pedestrian-vehicle (crashing, taxi hailing, collision avoidance, speed adaptation), vehicle-vehicle (lane switching, overtaking), and pedestrian-pedestrian (pushing, chasing, walking with arm around shoulder), which enrich the scene animations while enhancing the suitability for more diverse and dangerous cases. For instance, pushing is dangerous and corner in the street scenes, which is not easy to collect from real scene and crucial to safety; (ii) realistic low-level animation is achieved by the two physics-based animators, evident in the physically plausible human interaction animation with explicit physical feedback, and robust PedAnimator supports recovery from falling; (iii) precise control is enabled with the multi-LLM-agent framework through decomposition of complex descriptions containing abstract semantics into executable instructions. Combining these features, the generated animations are more realistic and better capture dangerous corner cases that are otherwise difficult to obtain



Command: Make the scene crowded like a traffic jam. On the left, one person has slipped, and another is about to help.

Figure 7: System results under abstract command with large-scale output. ChatAni can handle abstract semantics and scale effectively to manage larger scenarios.

4.2.2 SCALABLE AND ABSTRACT UNDERSTANDING

We demonstrate ChatAni’s ability to understand and decompose abstract semantics in 7, while also supporting more diverse interaction behaviors. At the same time, this result showcases ChatAni’s scalability in large-scale scene traffic generation, as it is capable of handling a large number of traffic participants simultaneously present in the scene, constructing crowded and complex traffic scenarios through the interactions of numerous people, vehicles, and their behaviors.

4.3 COMPONENT RESULTS

4.3.1 PEDANIMATOR

We evaluate PedAnimator on trajectory following and motion imitation tasks, as shown in Table 1. The quality and diversity of the generated results are measured by Frechet Inception Distance (FID) Heusel et al. (2017) and diversity metric (Div.) Wang et al. (2024) at normal speed, and by l-FID and l-Div. at low speed. Imitation accuracy is measured by Mean Per-Joint Position Error (MPJPE), while following accuracy is evaluated using following error (E_f). User preference (UP)

432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485

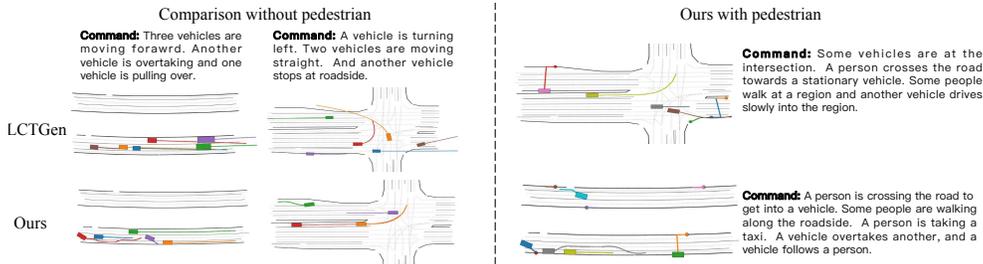


Figure 8: High-level planning results comparison under vehicle-only command, and our planning results with pedestrians involved. The boxes indicate vehicles and the circles indicate pedestrians. Our method significantly and consistently achieves more precise and reasonable plannings.

is assessed by 100 users who evaluate 33 animation segments. Benefit from hierarchical control and its training strategy, PedAnimator generates high-quality animations, demonstrating competitive performance in both following and imitation tasks.

We evaluate interaction task performance across methods, with PedAnimator employing a unified policy versus task-specific training in others (Table 2, Fig. 6). Testing three interaction tasks - pushing, patting, and arm-around-shoulder walking - reveals AMP Peng et al. (2021) produces reference-like motions but lacks task-specific data, yielding superficial mimicry without task completion. While PPO Schulman et al. (2017) and PULSE Luo et al. (2024)

achieve higher success rates, their animations exhibit unnatural movements. PedAnimator’s body-masked AMP and hierarchical control enable both high task success and human-like motion quality through integrated physical constraints and style preservation. More results and videos can be found in S2.6 including ablation and more visualizations or in the supplementary video.

4.3.2 VEHANIMATOR

As shown in 3, we evaluate the error between VehAnimator-generated animations and reference trajectories at different initial velocities to measure animation accuracy, comparing with Pure Pursuit Craig Coulter (1992) and Xu et al. Xu & Yu (2023). VehAnimator consistently outperforms other methods in all cases. Additional experiments including ablation studies and more visualizations can be found in S2.

4.3.3 MULTI-LLM-AGENTS PLANNING

We benchmark LLM-agent trajectory planning against language-based traffic generation methods (LCTGen Tan et al. (2023), ChatSim Wei et al. (2024)). Due to pedestrian scenario incompatibility in other methods, evaluation focuses on vehicle-only commands. Three instruction categories are evaluated: single vehicle, interaction, and composite. Using 5 maps with 26 samples per category, 1000 users assess description matching, road boundary compliance, and preference, as shown in Table 4. ChatAni achieves superior command-aligned planning and user preference. Visual comparisons in Fig. 8 (including pedestrian-involved

Methods	FID↓	Div.↑	l-FID↓	l-Div.↑	E_{mpjpe} ↓	E_f ↓	UP↑
Pacer	7.25	1.24	7.93	1.05	\	0.122	0.171
Pacer+	6.62	1.58	7.76	1.28	82.33	0.128	0.284
Ours	6.21	1.76	7.07	1.49	79.82	0.124	0.545

Table 1: Animation quality and accuracy evaluation for trajectory following and motion imitation.

Methods	Unified policy	Interaction 1	Interaction 2	Interaction 3
PPO	×	0.971	0.934	0.914
AMP	×	0.234	0.179	0.108
Pulse	×	0.975	0.942	0.925
Ours	✓	0.982	0.977	0.971

Table 2: Interaction tasks success rate evaluation.

Methods/Speed	0	5	10	20
PP	0.129/0.103	0.143/0.125	0.162/0.142	0.231/0.208
Xu et al.	0.075/0.054	0.084/0.066	0.095/0.077	0.138/0.114
Ours	0.059/0.037	0.062/0.041	0.077/0.054	0.106/0.088

Table 3: Comparison of position/velocity error.

Methods	Language command category			Within road	User preference
	single	interaction	compound		
LCTGen	91.9	20.7	64.1	59.9	15.4
ChatSim	84.6	5.79	77.1	86.1	5.51
Ours	93.7	87.6	88.4	92.6	79.1

Table 4: High-level planning evaluation for command matching rate, within road rate, and user preferences.

cases) demonstrate ChatAni’s precise fulfillment of complex requirements, contrasting with LCTGen’s limitations in interaction control and complex command execution.

4.4 APPLICATIONS TO AUTONOMOUS DRIVING

4.4.1 AUGMENTATION FOR PREDICTION TASK

We apply ChatAni for data augmentation in two prediction tasks. For traffic prediction (minADE/minFDE), using MTR Shi et al. (2022) on the Waymo Open Dataset Sun et al. (2020), we train on 243k scenes and augment with 20% (48k) via LCTGen Tan et al. (2023) and ChatAni. Both improve minADE/minFDE, but ChatAni’s superior diversity and reasonableness yielded greater gains. For human motion prediction (ADE/FDE), using HumanMAC Chen et al. (2023) on Human3.6M Ionescu et al. (2013), we

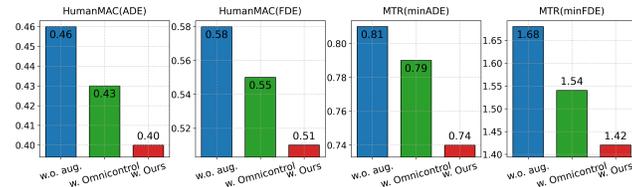


Figure 9: Prediction augmentation for traffic prediction and human keypoints prediction. All prediction tasks benefit from the enhancement provided by ChatAni, and the enhancement effect outperforms those of other baselines.

augment 20% (2k) with Omnicontrol Xie et al. (2023) and PedAnimator. Both methods improve performance, but Omnicontrol’s lower-quality generations constrained gains, while PedAnimator’s realistic, physics-compliant outputs produced substantially better results. We use the default parameters and training settings for both methods.

4.4.2 AUGMENTATION FOR VLM-BASED DRIVING SCENE UNDERSTANDING

To evaluate the impact of hazardous scenario generation on driving situations, we selected 30 scenes from the Waymo Open Dataset Sun et al. (2020) and used DriveLM Sima et al. (2024) to test collision rates 10 times per scene under three conditions: (1) original scenes, (2) ChatAni-edited scenes with intentionally heightened danger, and (3) edited scenes after fine-tuning with

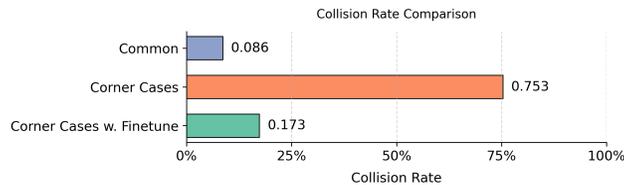


Figure 10: Collision rates of DriveLM Sima et al. (2024) under different scenarios. The data augmentation provided by ChatAni can significantly enhance the safety of VLM.

3000-frame ChatAni-augmented hazardous data. The ChatAni-edited scenes are deliberately modified to create additional hazards, which significantly increase testing collision rates. However, after augmentation and fine-tuning, the model learns to proactively decelerate or stop in these scenarios, thereby reducing collision rates. We choose the default parameters and training setting of DriveLM. See case in section S5 and supplementary video.

5 CONCLUSION

We propose ChatAni, the first system to achieve interactive and realistic language-driven multi-actor animation generation in street scenes. We introduce PedAnimator, a unified control policy for realistic pedestrian animation generation across multiple tasks, enabling fine-grained interactions. PedAnimator uses body-masked AMP to simultaneously improve the realism of action generation and efficiently achieve control, while employing task masking to implement a unified approach to various control methods. We introduce VehAnimator, a kinematics-based vehicle control policy with history-aware design to generate realistic vehicle animation. VehAnimator uses a learning-based approach to optimize the realism of vehicle animations based on a bicycle model, significantly enhancing robustness. ChatAni also utilizes multi-LLM-agent role-playing to enable interaction-aware high-level planning under language command. In the future, we plan to introduce more diverse traffic participants, such as cyclists.

REFERENCES

- 540
541
542 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo
543 Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint*
544 *arXiv:2303.08774*, 2023.
- 545 Nikos Athanasiou, Mathis Petrovich, Michael J Black, and Gül Varol. Teach: Temporal action composition for
546 3d humans. In *2022 International Conference on 3D Vision (3DV)*, pp. 414–423. IEEE, 2022.
- 547 Jinseok Bae, Jungdam Won, Donggeun Lim, Cheol-Hui Min, and Young Min Kim. Pmp: Learning to phys-
548 ically interact with environments using part-wise motion priors. In *ACM SIGGRAPH 2023 Conference*
549 *Proceedings*, pp. 1–10, 2023.
- 550 Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei
551 Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.
- 552 Luca Bergamini, Yawei Ye, Oliver Scheel, Long Chen, Chih Hu, Luca Del Pero, Błażej Osiński, Hugo Grim-
553 met, and Peter Ondruska. Simnet: Learning reactive self-driving simulations from real-world observations.
554 In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5119–5125. IEEE, 2021.
- 555 Michael J Black, Priyanka Patel, Joachim Tesch, and Jinlong Yang. Bedlam: A synthetic dataset of bodies
556 exhibiting detailed lifelike animated motion. In *Proceedings of the IEEE/CVF Conference on Computer*
557 *Vision and Pattern Recognition*, pp. 8726–8737, 2023.
- 558 Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan,
559 Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving.
560 In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11621–11631,
561 2020.
- 562 He Chen, Hongpinng Ren, Rui Li, Guang Yang, and Shanshan Ma. Generating autonomous driving test sce-
563 narios based on openscenario. In *2022 9th International Conference on Dependable Systems and Their*
564 *Applications (DSA)*, pp. 650–658. IEEE, 2022.
- 565 Ling-Hao Chen, Jiawei Zhang, Yewen Li, Yiren Pang, Xiaobo Xia, and Tongliang Liu. Humanmac: Masked
566 motion completion for human motion prediction. In *Proceedings of the IEEE/CVF international conference*
567 *on computer vision*, pp. 9544–9555, 2023.
- 568 R Craig Coulter. Implementation of the pure pursuit path tracking algorithm. *Technical report, DTIC Document*,
569 1992.
- 570 Wenhao Ding, Yulong Cao, Ding Zhao, Chaowei Xiao, and Marco Pavone. Realgen: Retrieval augmented
571 generation for controllable traffic scenarios. *arXiv preprint arXiv:2312.13303*, 2023.
- 572 Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban
573 driving simulator. In *Conference on robot learning*, pp. 1–16. PMLR, 2017.
- 574 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman,
575 Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint*
576 *arXiv:2407.21783*, 2024.
- 577 Lan Feng, Quanyi Li, Zhenghao Peng, Shuhan Tan, and Bolei Zhou. Trafficgen: Learning to generate diverse
578 and realistic traffic scenarios. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*,
579 pp. 3567–3575. IEEE, 2023.
- 580 Daniel J Fremont, Tommaso Dreossi, Shromona Ghosh, Xiangyu Yue, Alberto L Sangiovanni-Vincentelli, and
581 Sanjit A Seshia. Scenic: a language for scenario specification and scene generation. In *Proceedings of the*
582 *40th ACM SIGPLAN conference on programming language design and implementation*, pp. 63–78, 2019.
- 583 Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic meth-
584 ods. In *International conference on machine learning*, pp. 1587–1596. PMLR, 2018.
- 585 Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and
586 natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
587 *Pattern Recognition*, pp. 5152–5161, 2022.
- 588 Chuan Guo, Yuxuan Mu, Muhammad Gohar Javed, Sen Wang, and Li Cheng. Momask: Generative masked
589 modeling of 3d human motions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
590 *Pattern Recognition*, pp. 1900–1910, 2024.
- 591
592
593

- 594 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by
595 a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing*
596 *systems*, 30, 2017.
- 597 Sirui Hong, Xiawu Zheng, Jonathan Chen, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven
598 Ka Shing Yau, Zijuan Lin, Liyang Zhou, et al. Metagtpt: Meta programming for multi-agent collaborative
599 framework. *arXiv preprint arXiv:2308.00352*, 2023.
- 600 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu
601 Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- 602 Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3. 6m: Large scale datasets
603 and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis*
604 *and machine intelligence*, 36(7):1325–1339, 2013.
- 605 Stefan Jesenski, Jan Erik Stellet, Florian Schiegg, and J Marius Zöllner. Generation of scenes in intersections
606 for the validation of highly automated driving functions. In *2019 IEEE Intelligent Vehicles Symposium (IV)*,
607 pp. 502–509. IEEE, 2019.
- 608 Jordan Juravsky, Yunrong Guo, Sanja Fidler, and Xue Bin Peng. Padl: Language-directed physics-based char-
609 acter control. In *SIGGRAPH Asia 2022 Conference Papers*, pp. 1–9, 2022.
- 610 Yan Leng and Yuan Yuan. Do llm agents exhibit social behavior? *arXiv preprint arXiv:2312.15198*, 2023.
- 611 Binxu Li, Tiankai Yan, Yuanting Pan, Jie Luo, Ruiyang Ji, Jiayuan Ding, Zhe Xu, Shilong Liu, Haoyu Dong,
612 Zihao Lin, et al. Mmedagent: Learning to use medical tools with multi-modal agent. *arXiv preprint*
613 *arXiv:2407.02483*, 2024a.
- 614 Junkai Li, Siyu Wang, Meng Zhang, Weitao Li, Yunghwei Lai, Xinhui Kang, Weizhi Ma, and Yang Liu. Agent
615 hospital: A simulacrum of hospital with evolvable medical agents. *arXiv preprint arXiv:2405.02957*, 2024b.
- 616 Weiwei Li and Emanuel Todorov. Iterative linear quadratic regulator design for nonlinear biological movement
617 systems. In *First International Conference on Informatics in Control, Automation and Robotics*, volume 2,
618 pp. 222–229. SciTePress, 2004.
- 619 Weizi Li, David Wolinski, and Ming C. Lin. City-scale traffic animation using statistical learning and
620 metamodel-based optimization. *ACM Trans. Graph.*, 36(6):200:1–200:12, 2017. doi: 10.1145/3130800.
621 3130847. URL <https://doi.org/10.1145/3130800.3130847>.
- 622 Wen-Chieh Lin, Sai-Keung Wong, Cheng-Hsing Li, and Richard Tseng. Generating believable mixed-traffic
623 animation. *IEEE Trans. Intell. Transp. Syst.*, 17(11):3171–3183, 2016. doi: 10.1109/TITS.2016.2542283.
624 URL <https://doi.org/10.1109/TITS.2016.2542283>.
- 625 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural*
626 *information processing systems*, 36, 2024.
- 627 Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men,
628 Kejuan Yang, et al. Agentbench: Evaluating llms as agents. *arXiv preprint arXiv:2308.03688*, 2023.
- 629 Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned
630 multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pp. 851–866,
631 2023.
- 632 Jack Lu, Kelvin Wong, Chris Zhang, Simon Suo, and Raquel Urtasun. Scenecontrol: Diffusion for controllable
633 traffic scene generation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp.
634 16908–16914. IEEE, 2024.
- 635 Zhengyi Luo, Jinkun Cao, Kris Kitani, Weipeng Xu, et al. Perpetual humanoid control for real-time simulated
636 avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10895–10904,
637 2023.
- 638 Zhengyi Luo, Jinkun Cao, Josh Merel, Alexander Winkler, Jing Huang, Kris M. Kitani, and Weipeng Xu. Uni-
639 versal humanoid motion representations for physics-based control. In *The Twelfth International Conference*
640 *on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL
641 <https://openreview.net/forum?id=OrOd8Px002>.
- 642 Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass:
643 Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on*
644 *computer vision*, pp. 5442–5451, 2019.

- 648 Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David
649 Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, et al. Isaac gym: High performance gpu-based physics
650 simulation for robot learning. *arXiv preprint arXiv:2108.10470*, 2021.
- 651 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang,
652 Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with
653 human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- 654 Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors
655 for stylized physics-based character control. *ACM Transactions on Graphics (ToG)*, 40(4):1–20, 2021.
- 656 Xue Bin Peng, Yunrong Guo, Lina Halper, Sergey Levine, and Sanja Fidler. Ase: Large-scale reusable adver-
657 sarial skill embeddings for physically simulated characters. *ACM Transactions On Graphics (TOG)*, 41(4):
658 1–17, 2022.
- 659 Zeju Qiu, Weiyang Liu, Haiwen Feng, Yuxuan Xue, Yao Feng, Zhen Liu, Dan Zhang, Adrian Weller, and
660 Bernhard Schölkopf. Controlling text-to-image diffusion by orthogonal finetuning. In *NeurIPS*, 2023.
- 661 Rodrigo Queiroz, Thorsten Berger, and Krzysztof Czarnecki. Geosenario: An open dsl for autonomous driving
662 scenario representation. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pp. 287–294. IEEE, 2019.
- 663 Davis Rempe, Jonah Philion, Leonidas J Guibas, Sanja Fidler, and Or Litany. Generating useful accident-prone
664 driving scenarios via a learned traffic prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision
665 and Pattern Recognition*, pp. 17305–17315, 2022.
- 666 Davis Rempe, Zhengyi Luo, Xue Bin Peng, Ye Yuan, Kris Kitani, Karsten Kreis, Sanja Fidler, and Or Litany.
667 Trace and pace: Controllable pedestrian animation via guided trajectory diffusion. In *Proceedings of the
668 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13756–13766, 2023.
- 669 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization
670 algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 671 Yonatan Shafir, Guy Tevet, Roy Kapon, and Amit H Bermano. Human motion diffusion as a generative prior.
672 *arXiv preprint arXiv:2303.01418*, 2023.
- 673 Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Hugginggpt: Solving
674 ai tasks with chatgpt and its friends in hugging face. *Advances in Neural Information Processing Systems*,
675 36, 2024.
- 676 Shaoshuai Shi, Li Jiang, Dengxin Dai, and Bernt Schiele. Motion transformer with global intention localization
677 and local movement refinement. *Advances in Neural Information Processing Systems*, 35:6531–6543, 2022.
- 678 Chonghao Sima, Katrin Renz, Kashyap Chitta, Li Chen, Hanxue Zhang, Chengen Xie, Jens Beißwenger, Ping
679 Luo, Andreas Geiger, and Hongyang Li. Drivelm: Driving with graph visual question answering. In *Euro-
680 pean Conference on Computer Vision*, pp. 256–274. Springer, 2024.
- 681 Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo,
682 Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo
683 open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp.
684 2446–2454, 2020.
- 685 Simon Suo, Sebastian Regalado, Sergio Casas, and Raquel Urtasun. Trafficsim: Learning to simulate realis-
686 tic multi-agent behaviors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern
687 Recognition*, pp. 10400–10409, 2021.
- 688 Shuhan Tan, Kelvin Wong, Shenlong Wang, Sivabalan Manivasagam, Mengye Ren, and Raquel Urtasun. Sce-
689 negen: Learning to generate realistic traffic scenes. In *Proceedings of the IEEE/CVF Conference on Com-
690 puter Vision and Pattern Recognition*, pp. 892–901, 2021.
- 691 Shuhan Tan, Boris Ivanovic, Xinshuo Weng, Marco Pavone, and Philipp Kraehenbuehl. Language conditioned
692 traffic generation. *arXiv preprint arXiv:2307.07947*, 2023.
- 693 Chen Tessler, Yoni Kasten, Yunrong Guo, Shie Mannor, Gal Chechik, and Xue Bin Peng. Calm: Conditional
694 adversarial latent models for directable virtual characters. In *ACM SIGGRAPH 2023 Conference Proceed-
695 ings*, pp. 1–9, 2023.
- 696 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix,
697 Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation
698 language models. *arXiv preprint arXiv:2302.13971*, 2023.
- 699
700
701

- 702 Weilin Wan, Zhiyang Dou, Taku Komura, Wenping Wang, Dinesh Jayaraman, and Lingjie Liu. Tlcontrol:
703 Trajectory and language control for human motion synthesis. *arXiv preprint arXiv:2311.17135*, 2023.
704
- 705 Hua Wang, Mingliang Xu, Fubao Zhu, Zhigang Deng, Yafei Li, and Bing Zhou. Shadow traffic: A unified
706 model for abnormal traffic behavior simulation. *Comput. Graph.*, 70:235–241, 2018. doi: 10.1016/J.CAG.
707 2017.07.004. URL <https://doi.org/10.1016/j.cag.2017.07.004>.
- 708 Jingbo Wang, Zhengyi Luo, Ye Yuan, Yixuan Li, and Bo Dai. Pacer+: On-demand pedestrian animation
709 controller in driving scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern
710 Recognition*, pp. 718–728, 2024.
- 711 Yuxi Wei, Zi Wang, Yifan Lu, Chenxin Xu, Changxing Liu, Hao Zhao, Siheng Chen, and Yanfeng Wang.
712 Editable scene simulation for autonomous driving via collaborative llm-agents. In *Proceedings of the
713 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15077–15087, 2024.
- 714 Jungdam Won, Deepak Gopinath, and Jessica Hodgins. Physics-based character controllers using conditional
715 vaes. *ACM Transactions on Graphics (TOG)*, 41(4):1–12, 2022.
- 716 Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xi-
717 aoyun Zhang, and Chi Wang. Autogen: Enabling next-gen llm applications via multi-agent conversation
718 framework. *arXiv preprint arXiv:2308.08155*, 2023.
719
- 720 Pengchuan Xiao, Zhenlei Shao, Steven Hao, Zishuo Zhang, Xiaolin Chai, Judy Jiao, Zesong Li, Jian Wu, Kai
721 Sun, Kun Jiang, et al. Pandaset: Advanced sensor suite dataset for autonomous driving. In *2021 IEEE
722 International Intelligent Transportation Systems Conference (ITSC)*, pp. 3095–3101. IEEE, 2021.
- 723 Yiming Xie, Varun Jampani, Lei Zhong, Deqing Sun, and Huaizu Jiang. Omnicontrol: Control any joint at any
724 time for human motion generation. *arXiv preprint arXiv:2310.08580*, 2023.
- 725 Pei Xu, Xiumin Shang, Victor Zordan, and Ioannis Karamouzas. Composite motion learning with task control.
726 *ACM Transactions on Graphics (TOG)*, 42(4):1–16, 2023a.
- 727 Pei Xu, Kaixiang Xie, Sheldon Andrews, Paul G Kry, Michael Neff, Morgan McGuire, Ioannis Karamouzas,
728 and Victor Zordan. Adaptnet: Policy adaptation for physics-based character control. *ACM Transactions on
729 Graphics (TOG)*, 42(6):1–17, 2023b.
- 730 Yinda Xu and Lidong Yu. Drl-based trajectory tracking for motion-related modules in autonomous driving.
731 *arXiv preprint arXiv:2308.15991*, 2023.
732
- 733 Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Mo-
734 tiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*,
735 2022.
- 736 Mingyuan Zhang, Xinying Guo, Liang Pan, Zhongang Cai, Fangzhou Hong, Huirong Li, Lei Yang, and Zi-
737 wei Liu. Remodiffuse: Retrieval-augmented motion diffusion model. In *Proceedings of the IEEE/CVF
738 International Conference on Computer Vision*, pp. 364–373, 2023.
- 739 Ziyuan Zhong, Davis Rempe, Yuxiao Chen, Boris Ivanovic, Yulong Cao, Danfei Xu, Marco Pavone, and
740 Baishakhi Ray. Language-guided traffic simulation via scene-level diffusion. In *Conference on Robot Learn-
741 ing*, pp. 144–177. PMLR, 2023a.
- 742 Ziyuan Zhong, Davis Rempe, Danfei Xu, Yuxiao Chen, Sushant Veer, Tong Che, Baishakhi Ray, and Marco
743 Pavone. Guided conditional diffusion for controllable traffic simulation. In *2023 IEEE International Con-
744 ference on Robotics and Automation (ICRA)*, pp. 3560–3566. IEEE, 2023b.
- 745 Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou,
746 Yonatan Bisk, Daniel Fried, et al. Webarena: A realistic web environment for building autonomous agents.
747 *arXiv preprint arXiv:2307.13854*, 2023.
748
749
750
751
752
753
754
755

756 S1 ETHIC, REPRODUCIBILITY STATEMENT AND USAGE OF LLM

757
758 All related human participants involved in the user study provided consent, and there are no addi-
759 tional ethical concerns.

760
761 Regarding reproducibility, the experimental setup is described in the main text or appendix, with
762 some relevant code and data included in the supplementary materials. Full code and data will be
763 open-sourced in the future.

764 The LLM in this work is used for writing enhancement purposes.
765
766

767 S2 DETAILS OF PEDANIMATOR

768 S2.1 PEDANIMATOR STATE DETAILS

769
770 The components of the humanoid’s proprioception S_{prop}^p are as follows: joint positions $\mathbf{j} \in \mathbb{R}^{24 \times 3}$,
771 rotations $\mathbf{q} \in \mathbb{R}^{24 \times 6}$, linear velocities $\mathbf{v} \in \mathbb{R}^{24 \times 3}$, and angular velocities $\omega \in \mathbb{R}^{24 \times 3}$ Wang et al.
772 (2024). These components are normalized with respect to the agent’s heading and root position in
773 our simulator. The rotation \mathbf{q} is represented using a 6-degree-of-freedom rotation representation.
774 S_{prop}^p along with the trajectory state S_{traj}^p , the motion state S_{mo}^p to be mimicked, and the target state
775 S_{tar}^p for interaction, collectively forms the complete observation. During the task masking process,
776 the remaining relevant states are masked by multiplying them by 0 based on the task to be executed,
777 and in the motion state, specific joints can also be masked by multiplying them by 0 if necessary.
778 For example, in the experiment, only the upper-body related states were retained.
779
780

781 S2.2 TASK-RELATED REWARD AND TRAINING DETAILS

782
783 **Reward designs.** For trajectory following tasks Rempe et al. (2023), the reward is defined as
784 $R_{\text{trajectory}} = e^{-2\|\hat{p}_t - p_t\|}$, where \hat{p}_t is the position to followed at t , p_t is the current charac-
785 ter position. For motion imitation tasks Luo et al. (2023), $R_{\text{imitation}} = e^{-100\|\hat{j}_{pos}^t - j_{pos}^t\| \odot m^t} +$
786 $e^{-10\|\hat{j}_{rot}^t - j_{rot}^t\| \odot m^t} + e^{-0.1\|\hat{j}_{vel}^t - j_{vel}^t\| \odot m^t} + e^{-0.1\|\hat{j}_{\omega}^t - j_{\omega}^t\| \odot m^t}$, where $\hat{\cdot}$ indicates the motion states
787 to be imitated, m^t is the mask to select the joints for imitation, which are the joints of upper
788 body in our experiments. Different interaction tasks require distinct reward designs Peng et al.
789 (2022); here, we present three sample tasks: for pushing, $R_{\text{pushing}} = 1 - u^{up} \cdot u$, u^{up} is the
790 global up vector u^* is the target’s up vector; for patting, $R_{\text{patting}} = e^{-\|p^{rh} - c\|}$, p^{rh} is the posi-
791 tion of right hand and c is the target contact position; and for walking while bending the shoulder,
792 $R_{\text{walking_bending}} = e^{-2\|\hat{p}_t - p_t\|} + e^{-\|p^{rh} - c\|}$, which combines the trajectory following and target po-
793 sition contacting. The final reward is calculated as $R = 0.5 \cdot R_{\text{disc}} + 0.5 \cdot R_{\text{task}}$.
794
795

796 **Interaction process.** The specific execution process for the three interaction tasks can be understood
797 as follows: i) pushing, where the goal is to push the interaction object over, ii) patting, where the
798 task is to gently tap a specific part of the interaction object (e.g., the shoulder) without knocking it
799 over, and iii) walking with arm around another’s shoulder, where the agent walks along a specific
800 path while keeping the arm in contact with a specific location on the interaction object (e.g., the
801 shoulder). During the training of the interaction tasks, the interaction object is replaced with a box in
802 the physical environment to facilitate more stable training. However, during testing, the interaction
803 object is replaced with another character in the simulation environment, and physical collisions and
804 interactions are present, leading to the final output in the test phase.

805 **Failure recovery and crushing.** To enhance the robustness of PedAnimator, specifically to ensure
806 that it can handle external disturbances without failing due to perturbations caused by physical col-
807 lisions, we incorporate recovery during its training Luo et al. (2023). In this approach, the human
808 pose is initialized in a collapsed or otherwise unstable standing state. Training is then conducted
809 from this initial state, allowing the policy to learn how to recover from failure. This enables the
policy to exhibit robustness when responding to physical collisions and interactions. Without this,
the policy might fail to continue the action after even minor disturbances.

810 The crushing process is implemented for pedestrians in IsaacGym by a vehicle-like shape box. The
811 pedestrian is crushed by the box and related animation can be readout. For the crushing vehicle, the
812 velocity changing due to crushing is already achieved by the planning process.
813

814 S2.3 NETWORK ARCHITECTURE

815

816 All relevant models, including the baselines, adopt the same network architecture, using an MLP as
817 the policy network with hidden layers of 2048 and 1024 units. The final output is directed to either
818 the latent space or the PD control signal, depending on whether hierarchical control is employed.
819 The remaining network components, such as the discriminator, value network, control frequency
820 (30Hz), and hyperparameters used for training, are consistent with those adopted in Pacer+ Wang
821 et al. (2024). All training and testing are conducted on an NVIDIA 4090. The entire training process
822 requires approximately 20 hours to fully converge.
823

824 S2.4 EVALUATION DETAILS

825

826 **Following and imitation.** For the following and imitation tasks, we adopt the same computation
827 methods as those used in Pacer+ Wang et al. (2024). The calculations of FID and diversity are
828 performed using the same manual feature extraction approach as in Pacer+, with 1000 segments
829 selected from the AMASS dataset for FID computation. For the low-speed l-FID and l-diversity, we
830 also follow Pacer+ by testing on instances where the speed is less than 1 m/s.

831 **Interaction tasks.** For the three interaction-related tasks: i) pushing, the success criterion is whether
832 the object is pushed over within the specified timestep (with a tilt along the z-axis greater than 30°),
833 and no part of the body other than the hands is in contact with the target; ii) patting, the success
834 criterion is whether, within the specified timestep, the right hand is within 5 cm of the target’s specific
835 location and remains there for at least 50 timesteps, with no other part of the body in contact with
836 the target except the hands; iii) walking with arm around another’s shoulder, the success criterion
837 is that the maximum deviation from the reference trajectory is no greater than 10 cm, and the right
838 hand is within 5 cm of the target’s specific location for at least 150 timesteps. In the interaction
839 ablation study, for the user study, we recruit 15 participants to test 30 dynamic sequences and select
840 the ones they considered to have the highest quality.

841 S2.5 RENDERING DETAILS

842

843 We utilize the rendering pipeline from ChatSim Wei et al. (2024), employing Blender’s Cycles as the
844 rendering engine. Background rendering and HDRI lighting from ChatSim are applied to the scene.
845 The human model is based on SMPL Loper et al. (2023), with the corresponding mesh initialized,
846 and rendered using skin and clothing textures provided in the Bedlam dataset Black et al. (2023).
847

848 S2.6 SUPPLYMENTARY EXPERIMENTS

849

850 We provide more comprehensive visualization results to compare different aspects of PedAnimato-
851 r’s characteristics. All comparisons, along with animation results, can also be observed in the
852 supplementary video material.

853 **Ablation study** As shown in Table S1, we also conduct ablation studies on interaction tasks to
854 verify the effectiveness of task embedding (TE), hierarchical control (HC) and body masked AMP
855 (M. AMP). The results indicate that, without task embedding, PedAnimator struggles to accurately
856 identify the required interaction task, leading to difficulty in task completion. When body-masked
857 AMP is removed, similar to AMP, the discrimination reward lacks a task completion reference.
858 In this situation, the results can be merely close to the reference data clip without the ability to
859 successfully execute the intended task. Without hierarchical control, the absence of action priors
860 slows the training process and reduces the success rates.

861 **Imitation and following comparison.** As shown in S1, we compared the visual effects of
862 Pacer+ during the imitation and following processes. It is evident that, under the influence of hi-
863 erarchical control, PedAnimator enables smoother and more seamless transitions between different
tasks—transitions that cannot be achieved with the priors provided by AMP.

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

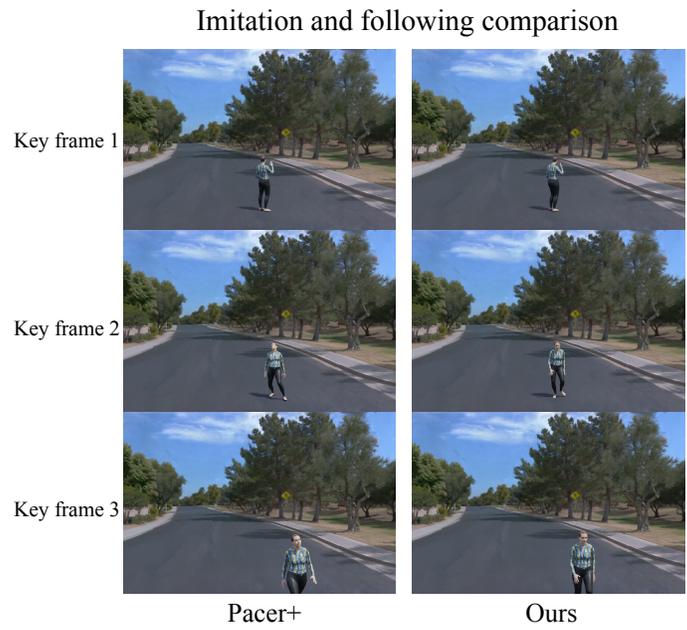


Figure S1: Comparison of imitation and following.

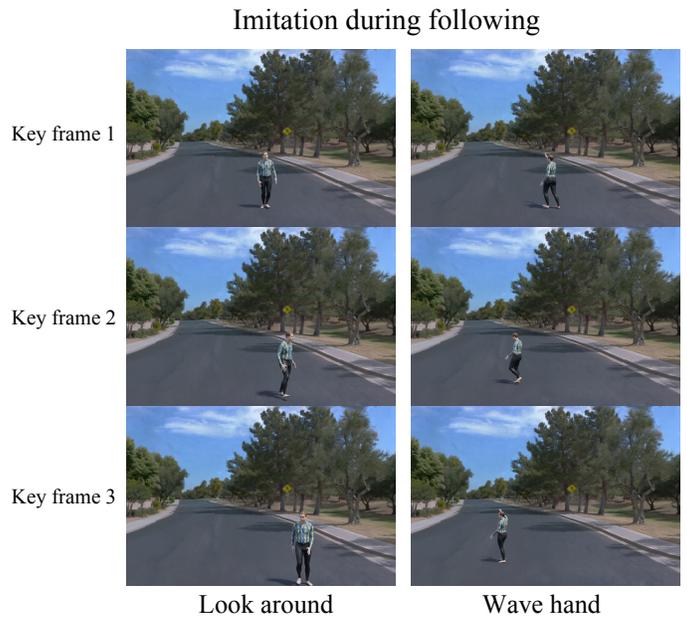


Figure S2: Imitation during following.

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

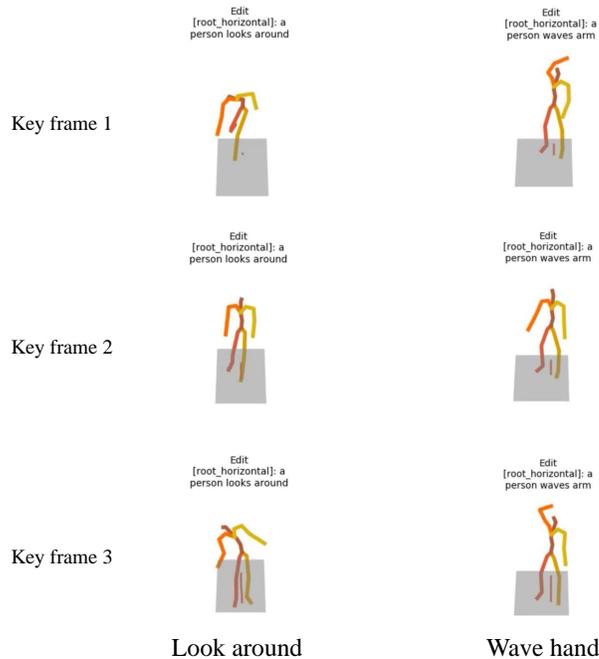


Figure S3: Failure of PriorMDM Shafir et al. (2023) with action specification during following.



Figure S4: Animations generated by PedAnimator with MoCap, T2M, and interaction target control. PedAnimator generates diverse and realistic animations under various control signals.

Results of imitation during following. As shown in S2, we also provide the results of PedAnimator performing both imitation and following simultaneously. In the case of following a specific trajectory, the upper body performs actions such as looking around and waving, showcasing capability of PedAnimator to handle both imitation and following tasks concurrently.

Interaction task comparison. As shown in S6 S7 S8, we compared the performance of AMP, PPO, PULSE, and our method across three interaction tasks. It is still evident that while PPO and PULSE are able to complete the tasks under the given conditions, they fail to produce natural and human-like results. AMP, on the other hand, can only approximate the reference in the discriminator (walking or running), and is completely unable to complete the tasks. In contrast, our PedAnimator successfully completes the tasks while generating natural and human-like results.

Multiple pedestrians visualization. Fig.S4 shows PedAnimator’s generation capabilities under various control modalities, including MoCap, Text2Motion-driven upper body control, trajectory control, and interactive behavior control. The system synthesizes diverse and realistic animations through planned control schemes.

S2.7 FURTHER DISCUSSION OF KINEMATICS METHODS

For kinematics-based methods, some approaches can achieve following and motion specification, but they are completely unable to handle interaction-related tasks. Additionally, for following and motion specification tasks, these methods often suffer from overfitting to the dataset, leading to suboptimal performance. As shown in S3, the results in noticeable sliding steps and unnatural movements.

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

TE	HC	M. AMP	Success rate	User preference
×	✓	✓	0.403	0.120
✓	✓	×	0.186	0.069
✓	×	✓	0.948	0.344
✓	✓	✓	0.977	0.467

Table S1: Ablation study of PedAnimator. Each design contributes to the final results consistently, and the complete PedAnimator demonstrates the best performance.

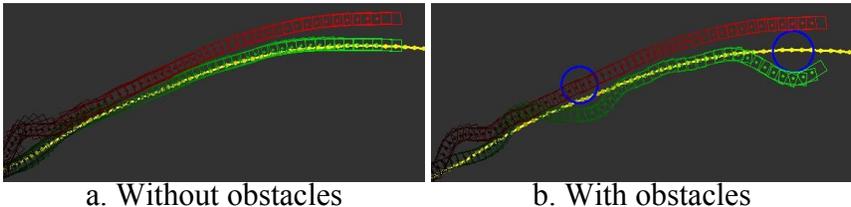


Figure S5: Comparison of vehicle animation generation. Green boxes are our results and red boxes are Xu et al.'s Xu & Yu (2023). Yellow lines are the planned trajectory and blue circles are obstacles. VehAnimator shows better tracking results and avoid the obstacles.

S3 DETAILS OF VEHANIMATOR

S3.1 NETWORK ARCHITECTURE, PARAMETERS OF BICYCLE-MODEL AND TRAINING DETAILS

All networks in VehAnimator are implemented as MLPs. The policy network consists of layers with dimensions 256, 256, 128, 128, 64, and 64, while the value network has layers with dimensions 1024, 512, 256, and 128. During training, the parameters of the bicycle model (L, W, l_f, l_r) are set to two configurations: (2.7, 1.8, 0.9, 0.9) and (6.1, 2.5, 2.3, 2.0), mixed to accommodate vehicles of varying sizes. These parameters can be adjusted as needed based on specific requirements, with the two configurations provided here serving as examples. The training environment for VehAnimator was implemented through our custom codebase, with the full training process requiring approximately 10 hours on a single NVIDIA RTX 4090 GPU.

S3.2 OBSTACLE STATE

The states of obstacles are composed of their orientation and distance relative to the vehicle’s own coordinate system (considering the radius of the obstacles). The state vector is initialized with a maximum number of observable obstacles. The vector is then populated with the specific identifiers of the actual obstacles, and any remaining entries are masked. Obstacles that are too far away are directly excluded, meaning their states are not considered, and they do not contribute to the reward calculation. The distance threshold for exclusion is set to 10 in the experiment.

S3.3 SUPPLYMENTARY EXPERIMENTS

Ablation study. We also conduct ablation study in Table S2, examining the effects of action filtering (Filt.), the composition of the action space (Action spa.) which means it consists of the variations or absolute scalar of velocity and steering or their variations, and the inclusion of historical states (His. state). The ablation results validate the effectiveness of each design.

Visualization. Furthermore, visual experiments in Figure S5 demonstrate that our method achieves more precise tracking and obstacle avoidance capabilities in scenarios with obstacles.

Robustness. We also provide the results using LQR Li & Todorov (2004) in S3. It can be observed that LQR is capable of vehicle animation generation from a planned trajectory to some extent. However, the planned trajectory lacks relevant constraints, which may lead to unreasonable turns or abrupt changes, causing LQR to often produce less than ideal results. Furthermore, we tested the robustness of vehicle animation generation by adding noise with a mean of 0 and variance of σ to

Filt.	Action spa.	His. state	Position error	Velocity error
×	×	×	0.138	0.114
✓	×	×	0.132	0.111
✓	✓	×	0.115	0.092
✓	✓	✓	0.106	0.088

Table S2: VehAnimator ablation study. Each component contributes to the final result for better performance.

Methods/Speed	0	5	10	20
LQR Li & Todorov (2004)	0.074/0.058	0.086/0.070	0.092/0.079	0.125/0.0108

Table S3: Position/velocity error of LQR.

the planned trajectory. As shown in S4, the results show that LQR is highly sensitive to noise, often producing significantly worse outcomes under its influence, while other methods are relatively less affected by the noise.

Visualization for effectiveness. In the supplementary video, we provide results comparing the planned trajectories without using VehAnimator and with VehAnimator. It is evident that, without the involvement of VehAnimator, the animation generated by simply calculating heading between consecutive frames of the directly planned trajectory are highly unnatural, exhibiting noticeable tail swings and abrupt changes. In contrast, the results using VehAnimator are much more realistic and natural. This demonstrates the necessity of VehAnimator, as trajectories without physical constraints are highly unnatural and impractical.

S4 DETAILS OF HIGH-LEVEL PLANNING

S4.1 LLM-AGENT DETAILS

We provide the relevant sample prompts for the LLM agent in S10 and S11. All outputs from the LLM are in JSON format, and corresponding follow-up functions are used to convert the JSON outputs into the required data structures.

Similar to the rendering process, we also use the Waymo Open Dataset Sun et al. (2020) as the planning dataset in all experiments, and the final results are presented based on this dataset.

S4.2 DIFFERENT LLMs

We further validated the impact of different LLMs on the results in S5. We conducted experiments using GPT-3.5 Ouyang et al. (2022) and Llama-3 Dubey et al. (2024) 70B (smaller models struggle to accurately execute the instructions). All other experimental settings remained consistent with those in the main text. It is evident that while other LLMs can handle the task to some extent, GPT-4 Achiam et al. (2023) demonstrates the most accurate understanding and decomposition of the instructions.

S4.3 COLLISION HANDLING

In the high-level planning process, we also designed a collision handling mechanism to avoid unintended collisions. Specifically, during the trajectory generation, collision detection is performed, and when a collision is detected, a velocity adjustment function is applied to modify the speed of one of the agents. This velocity adjustment function uses a nonlinear mapping to combine the original planned result with an interpolated trajectory, leading to a planning result with different speeds. We evaluated the probability of collisions across 50 generated samples, which is calculated by dividing the number of vehicles that experienced a collision by the total number of vehicles. As shown in S6, the collision rate for all methods remains low, and ChatAni also achieves a low collision rate while incorporating collision handling.

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

Pushing



Figure S6: Comparison of pushing.

Patting



Figure S7: Comparison of patting.

σ	PP Craig Coulter (1992)	Xu et al. Xu & Yu (2023)	LQR Li & Todorov (2004)	ours
0.00	0.162/0.142	0.095/0.077	0.092/0.079	0.077/0.054
0.01	0.168/0.147	0.098/0.080	0.322/0.289	0.079/0.058
0.03	0.169/0.151	0.098/0.079	0.568/0.479	0.082/0.060

Table S4: Vehicle animation generation under Gaussian noise. σ indicates the variance of noise.

Methods	Language command category			Within road
	single	interaction	compound	
Ours-Llama3	0.885	0.742	0.812	0.920
Ours-GPT3.5	0.854	0.738	0.834	0.915
Ours-GPT4	0.952	0.883	0.896	0.935

Table S5: High-level planning evaluation for different LLMs.

S5 CORNER CASE FOR VLM

As shown in S9, we demonstrate a corner case for Visual-Language Model (VLM)-related experiments, with the command: "Add a stationary car and have a pedestrian walk out from behind it." In this scenario, the ego vehicle initially fails to detect the pedestrian while moving forward. However, as it approaches the stationary car, it comes dangerously close to the pedestrian, leading to a collision—a classic visibility-induced corner case. Without fine-tuning, collisions occur frequently due to the system's inability to anticipate the pedestrian's presence. After fine-tuning, the vehicle proactively decelerates when significant blind zones are detected in the field of view, reducing collision probability and achieving data augmentation. The related video can be found in the supplementary video.

S6 USER STUDY DETAILS

We conducted user studies via questionnaire format, distributing surveys containing samples from different experimental sources for participants to select the option that best meets requirements or visual preferences. All our studies involve choosing one option from multiple samples without direct rating tasks. Specific implementation details are described in their respective experiment sections, with an interface example provided in Figure S12.

S7 LIMITATIONS

The current framework does not account for complex multi-actor animations involving cyclists, which will be addressed in future work. Additionally, fine-grained pedestrian-vehicle interactions (e.g., collision scenarios) involving intricate physical processes are reserved for subsequent research.

S8 BORDER IMPACTS

This work serves dual purposes: (1) as a street-scene animation generator, it enhances simulation platforms with diverse, realistic pedestrian/vehicle animations for corner-case scenarios while supporting game/film production through efficient animation synthesis; (2) as a generative prior, ChatAni-produced animations can guide video generation models for controlled outputs. Technically, ChatAni pioneers the integration of multi-LLM-agent role-playing in traffic scenarios and addresses multi-pedestrian physical interactions in street environments. Key components—interaction training protocols, task-masking mechanisms, and body-masked Adversarial Motion Priors—are transferable to physics-based human animation and robotic control systems. This work demonstrates no evident potential societal negative impacts or unethical misuse scenarios.

1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200
 1201
 1202
 1203
 1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241

	ChatSim Wei et al. (2024)	LCTGen Tan et al. (2023)	Ours
Collision rate	0.149	0.092	0.067

Table S6: Collision rate of high-level planning

Walking with arm around another’s shoulder



Figure S8: Comparison of walking with arm around another one’s shoulder.



Figure S9: Corner case sample.

1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

Oracle agent prompt.

I have a requirement for analyzing a scenario. I will provide you with a requirement, and I need your help to break it down into four pieces of information: (1) identify all the agents, (2) initialize each agent's state, (3) formulate a text for each agent. You should provide the information in JSON format. Note that your output should only include the JSON format of the information, not the analysis process.

(1) Identify all the agents: This means you need to extract all the agents from the scenario based on the input text. The text may be quite lengthy, so you can locate the agents by identifying the nouns, which may assist you in this task. The agents must be the objects involved in autonomous driving. The format should be as follows: `agent_list = {'0': 'agent_name_0', '1': 'agent_name_1', ...}`. The keys (e.g., '0', '1', etc.) represent unique agent IDs, which you should assign starting from '0'. The agents should only include vehicles like cars, pedestrians, trucks, and buses, and should not include static objects like trees or buildings. Ensure that each agent is given a distinct name. For example, in the sentence 'car a wants to overtake car b', the nouns are 'car a' and 'car b'. Both are objects in autonomous driving scenarios. Therefore, the `agent_list` should be formatted as follows: `agent_list = {'0': 'car_a', '1': 'car_b'}`. Pedestrians can be represented with identifiers like 'ped_0' and may share numbering with entities such as 'car_0'.

(2) Initialize each agent's state. In this task, you need to determine four aspects for each agent: 1. Agent type, 2. Movement, 3. Speed. You can use the results from task (1) to complete this task. Provide the initial state of each agent in a list, formatted in JSON. Most time the speed is bigger than 0. For example, the initial states should be defined as follows: `init-states = ['agent_id': '0', 'agent_type': 'car', 'movement': 'overtaking', 'speed': 60, 'agent_id': '1', 'agent_type': 'car', 'movement': 'straight', 'speed': 30]`. If one car intends to overtake another, it should ideally be at least twice as fast as the car it is overtaking. The initial state must include the agent type, movement, and speed.

Agent type should in [pedestrian,vehicle], action should in ['static', 'straight', 'pull over', 'turn over', 'overtake', 'turn left', 'turn right', 'straight left', 'straight right'] if it's a vehicle, and in ['static', 'crossing', 'straight'] if it is a pedestrian.

(3) Formulate a text for each agent. As an omniscient observer, you should instruct each agent on their actions through a text. The text should contain two pieces of information: 1. The agent type, 2. The agent's intention. You need to provide this in a JSON format. `guide_texts = {'0': 'text1', '1': 'text2', '2': 'text3', ...}` where the key is the agent's ID and the value is the text. The agent's name should be consistent with those in the `agent_list`.

Your answer should be in a JSON format, and must include the three information: `agent_list`, `init-state`, `guide_texts`. And `init-state` must include the agent type, movement and speed.

Figure S10: Oracle agent prompt.

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

Actor agent prompt.

Now, you are an agent in the autonomous driving scenario. I will give you a text, and agent_list, describing who you are and what you need to do. Note that your output should not include your analysis process, only the JSON format of the information you provide. I need you to analyze the text and give me the four information of the ego agent to describe what type of lane the agent should be in: (1) depend (2) speed change (3) keypoints list (4) behavior. Note that you only need to give the information of the ego agent, not the other agents.

(1) depend. You need to determine the depend of the agent according to the agent's intention. And give it like [depend_agent_id, depend_type]. You can get the depend_agent_id from the agent_list. Depend_type should be in ['end', 'start', 'trajectory', 'None']. 'end' represents the ego agent's end point is the depend agent's end point, 'start' represents the ego agent's start point is the depend agent's start point, 'trajectory' represents the ego agent's whole trajectory is the depend agent's trajectory. If it has no depend, you should give it [-1, 'None'].

(2) speed change. You need to determine the speed change of the agent according to the agent's intention. If you want to speed up, the speed change should be 1. If you want to slow down, the speed change should be -1. If you want to keep the speed, the speed change should be 0. For example, if there is someone nearby, you should slow down.

(3) keypoints list. You need to determine the number of keypoints required for your current behavior, the confirmation method for each keypoint, and their respective parameters, then return a list containing this information. Each keypoint can be confirmed in one of three ways: (1) Map-based (type '0'), with parameters: lane position ('left', 'right', or 'front'), lane type ('centerline' or 'boundary'), and driving direction ('turn left', 'turn right', or 'straight'). (2) Lane-relationship-based (type '1'), with the parameter being the relationship to the previous keypoint ('opposite direction adjacent', 'same direction adjacent', 'adjacent straight', 'adjacent left turn', 'adjacent right turn', 'different type adjacent', or 'opposite boundary'). (3) Agent-based (type '2'), with parameters specifying required information type ('point' or 'trajectory'). The result should be a sequential list of dictionaries where the key is the type ('0', '1', or '2') and the value is the corresponding parameters. For example, a left-turning car might return ['0': 'position': 'front', 'lane_type': 'centerline', 'direction': 'turn left', '1': 'relationship': 'adjacent straight'], while a car in the left lane picking up 'ped_a' might return ['0': 'position': 'left', 'lane_type': 'centerline', 'direction': 'straight', '2': 'info_type': 'point']. Ensure compliance with parameters, common sense, and traffic regulations, with the first keypoint typically using type '0'.

(4) behavior. You need to determine your behavior. If you are a vehicle, the behavior is "None." If you are a pedestrian, the behavior corresponds to the description provided. There are two scenarios: if your behavior is a specific action such as calling or waving, simply return the text of that action; if the behavior involves interactive actions such as pushing, patting, or walking with an arm around another person, you must return the exact predefined descriptions for these three types of interactions without modification. For example, if you are calling, your behavior is "calling phone."; if you push ped_1, your behavior is "pushing." Your answer should be in a JSON format, and must include depend, speed change, keypoints list and behavior.

Figure S11: Actor agent prompt.



Figure S12: A sample interface of user study.