
Automated Reward Design for Gran Turismo

Michel Ma*

Mila, University of Montreal
michel.ma@mila.quebec

Takuma Seno

Turing Inc.
takuma.seno@turing-motors.com

Kaushik Subramanian

Sony AI
kaushik.subramanian@sony.com

Peter R. Wurman

Sony AI
peter.wurman@sony.com

Peter Stone

Sony AI, UT Austin
peter.stone@sony.com

Craig Sherstan

Sony AI
craig.sherstan@sony.com

Abstract

When designing reinforcement learning (RL) agents, a designer communicates the desired agent behavior through the definition of reward functions - numerical feedback given to the agent as reward or punishment for its actions. However, mapping desired behaviors to reward functions can be a difficult process, especially in complex environments such as autonomous racing. In this paper, we demonstrate how current foundation models can effectively search over a space of reward functions to produce desirable RL agents for the Gran Turismo 7 racing game, given only text-based instructions. Through a combination of LLM-based reward generation, VLM preference-based evaluation, and human feedback we demonstrate how our system can be used to produce racing agents competitive with GT Sophy, a champion-level RL racing agent, as well as generate novel behaviors, paving the way for practical automated reward design in real world applications.

1 Introduction

Significant algorithmic progress in reinforcement learning has led to impressive results in various domains such as robotics and video games [37, 2, 41]. In particular, GT Sophy [41] used RL to train an agent from scratch to outrace top human competitors in the Gran Turismo video game, marking a significant achievement in applying RL to complex, high-fidelity simulated environments. Achievements like GT Sophy are built on algorithms that maximize a predefined scalar objective known as the reward function. Thus, the reward function is the primary tool for practitioners to use to communicate a desired behavior to an agent. However, even for experts, reward function design remains a difficult process [17] and is a significant barrier to the wider adoption of RL.

When RL is applied to novel tasks, we must first define an appropriate reward function, a process called reward design [33] or reward shaping [25]. Reward design remains largely a manual process of trial and error requiring both RL expertise and specific domain knowledge [3, 17]. Often times, desired behaviors are easily communicated between humans through natural language and yet the same behaviors cannot be readily expressed through numerical functions. For example, designing a reward function for a humanoid to “*perform a back flip*” or for an autonomous racing agent to

*Work done during an internship at Sony AI.

“*win races while maintaining good sportsmanship*” is well understood and recognizable by most individuals, yet difficult to express quantitatively either due to the complexity of encoding the behavior or because of the inherent inexactness of the behavior.

Recently, large-scale foundation models have shown impressive capabilities in code generation, reasoning, image processing, and natural language understanding across a wide range of domains [32, 39, 31, 1]. The increased accessibility and performance of these models has led to a growing body of work leveraging the near expert-level knowledge contained within large language models (LLMs) and video language models (VLMs) for reward design in reinforcement learning [22, 42, 12, 15, 8, 35, 21, 5]. One promising class of approaches to automate the conversion from text-specified behaviors to fully trained agents is to leverage LLMs to turn goals such as “*win races while maintaining good sportsmanship*” into appropriate reward functions as code amenable for any out-of-the-box RL algorithm to optimize.

Of particular interest is Eureka [22], an LLM-based approach that uses an evolutionary algorithm to generate and train multiple reward functions in parallel; an approach which would otherwise be very labor intensive for any one human to perform. However, their method presupposes access to a ground-truth *fitness metric* that provides a numerical signal for assessing the success of a given policy. Not only can it be computationally expensive to train agents for each generated reward function, but fitness metrics are not always available, and designing this metric is not necessarily simpler than directly designing the reward function itself. Alternatively, one could rely exclusively on human feedback as a replacement for the fitness metric [42, 12]. Unfortunately, human feedback inherently requires a certain amount of manual effort that scales unfavorably with the number of reward functions generated, largely defeating the purpose of an automated search system.

In this paper, we set out to develop a general purpose LLM-based system that can efficiently and reliably produce RL-trained agents for Gran Turismo given only text-based inputs from a user. To this end, we introduce a scalable iterative framework for automated reward design without the need for a *fitness metric*. We hypothesize that current foundation models can accurately distinguish good from bad behavior just as well as humans can. Therefore, after each iteration, we elicit preferences from a VLM over a population of learned policies to select the best-performing reward function, serving as a replacement to the *fitness metric*. The resulting set of preferences is used to automatically prune misaligned candidate reward functions for future iterations via a trajectory alignment coefficient [24].

We use GT Sophy, a superhuman agent trained with expert-designed reward functions, as a useful point of reference in the Gran Turismo 7 environment for our experiments. GT 7 is a video game built around a complex racing simulator grounded in reality for which it has historically been difficult to design appropriate reward functions [41]. We show that our framework can replace countless hours of expert reward design, generating reward functions capable of producing a racing agent similar in performance to current state-of-the-art agents with minimal human intervention. We also conduct experiments to better understand the contributions of the VLM’s preferences towards the final results, and examine the novelty of the generated reward functions. Finally, our framework is used end-to-end to produce novel behaviors in GT 7 directly from text, eliminating the need for manual reward design.

2 Related Works

RL for Autonomous Racing. In the first instances of using RL for autonomous racing [9, 29], the agent’s goal, to minimize lap times, was straight-forward to implement for RL, because no other car were on the track. However, as progress was made towards developing agents in a multi-car race setting [34, 41, 20], reward functions became more complex to account for the subtleties of motor-sport racing rules [23]. These challenges are not unlike those in the field of autonomous driving [6], where additional considerations must be dealt with to ensure reasonable behaviors within the broader context of other humans [17, 12]. To the best of our knowledge, our work is the first of its kind to push the edge of superhuman performance in an autonomous racing environment with an automated reward design system.

Coding Reward Functions with LLMs. A number of existing works have used LLMs to code reward functions for RL agents [22, 21, 45, 36, 44, 42, 12, 11] in robotics tasks [22, 36, 44, 42], video games [21] and autonomous driving [11, 12]. The reward generation process in these works largely follow the same blueprint: the environment and task description are fed as text-based context to the LLM, and it must then produce reward code either from scratch or following an existing template to be used for an RL algorithm. Where they differ is mostly in their applications and their automated

evaluation process. Notably, several of these works [22, 36, 44, 21, 45] assume access to some oracle sparse reward function or fitness metric to automatically provide feedback about a reward function’s performance. Others [42, 12, 11] without such access rely heavily on a human’s feedback for each generated agent for evaluation. In contrast, our work attempts to generate superhuman agents with automated LLM reward design in Gran Turismo without access to a fitness metric or oracle reward, while limiting the amount of human feedback necessary.

LLMs and VLMs as the Reward Function. Alternatively, others have suggested using LLMs [15, 16, 18] and VLMs [8, 13, 43, 35, 30] more directly either as a reward function, or as a supplementary shaping reward to the existing rewards of the environment instead of going through the intermediate step of coding a reward function. In all these cases, some reward function is extracted from the foundation models through preference elicitation [15, 16, 18] or by using the embedding space of the foundation model to numerically quantify behaviors [8, 13, 43, 35, 30]. In either case, such methods do not generate interpretable reward functions, and necessitate a white-box foundation model in the latter case in order to access a model’s embedding space. While a part of our system elicits preferences from a foundation model similarly to [15] and uses a VLM to identify behaviors [35, 13], our method uses these components as a supporting module for a larger reward coding framework.

3 Problem Setting

The Reward Generation Problem. Defined as an extension of the *reward design problem* [33], the *reward generation problem (RGP)* [22] can be expressed by the tuple $P = \langle M, \mathcal{R}, \pi, F, g \rangle$, where $M = (S, A, T)$ contains the state space S , action space A and transition function T of the environment. \mathcal{R} is the space of all reward functions $R \in \mathcal{R}$ in the resulting Markov Decision Process (M, R) , and $F : \Pi \rightarrow \mathbb{R}$ is a fitness function that quantifies the performance of a policy. Suppose we have access to an algorithm $\mathcal{A} : \mathcal{R} \rightarrow \Pi$ to find a policy that optimizes any given reward function. The objective of this problem is to find a reward function R such that $\mathcal{A}(R) = \pi^*$ maximizes F . The defining feature of an RGP that separates it from the typical reward design problem is the text-based nature of the problem. Specifically, the task description, or the agent’s goal, is given as the string g , and the components (S, A, T) are assumed to be fully expressible as code.

In this paper, we make the assumption that the fitness function F is not accessible while searching for a reward function. In practice, this can be due to a multitude of reasons. Designing a fitness metric for a task can pose similar challenges to designing reward functions, where the definition of *success* is either too vague or too complex to clearly define numerically. Other times, calls to the fitness metric can simply be too expensive to use during training, such as the case where F is based on human expert evaluations.

Iterative LLM-based Reward Design. One approach to addressing the *reward generation problem* is to use an LLM to iteratively generate reward function code towards maximizing the fitness metric. Such an approach, which we name *iterative LLM-based Reward Design*, follows the general structure in Figure 1. First, the environment and task descriptions are fed into some LLM-powered system that outputs a trained agent. Then, feedback is optionally collected from a human in the form of text, and both the agent and feedback are fed back into the main LLM system for improvement on the following iteration. This general framework encompasses several prior works [22, 42, 12, 21, 36, 11] in LLM-based reward design, where the core systems for reward generation, evaluation and training can vary from method to method, but all follow an iterative principle to automatic reward design. As an additional benefit to these systems, the iterative process of designing reward functions can be

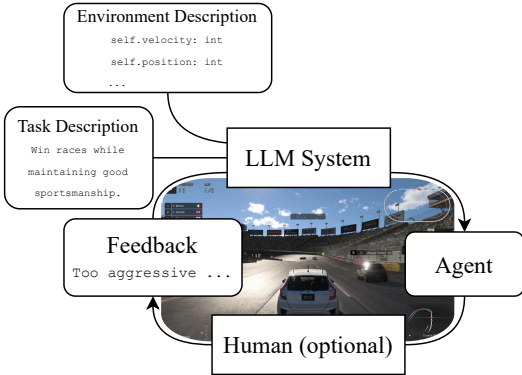


Figure 1: Overview of an iterative LLM-based reward design pipeline. The environment and task description are fed into a black-box system to produce compliant agents, and the agent is fed back into the system for improvement, sometimes accompanied by some human feedback.

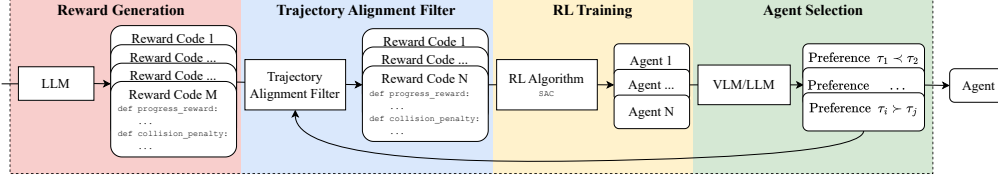


Figure 2: Expansion of the *system* block in Figure 1 for our framework. First, reward code is generated from scratch by an LLM. The rewards are then filtered through a trajectory alignment filter to avoid training misaligned rewards. The rewards are all trained separately by an RL algorithm. Finally, a VLM or LLM selects the best agent for further iterations.

useful for human users as well. In the cases where F is held implicitly in the mind of a user, g is often imperfect due to the vagueness of language, and it is only by observing agent behaviors that a user can iteratively refine its own expectation of an optimal agent.

4 Method

Our automated reward design system follows the same outer loop for *iterative LLM-based Reward Design* described in Section 3 and Figure 1. For the environment description, other methods [22, 42, 44, 36] provide the full code that describes the environment as part of the LLM prompt. We argue that having environment code neatly contained as a single file or class is more an exception than the rule. Many complex environments, even those that are simulations, such as GT, can rarely be captured so concisely. To this end, we provide only the Python dataclass API of the state and action space of the Gran Turismo 7 training environment, which contains variable names, comments, and data types for each field.² Instead of providing code for the dynamics, our system relies on the LLM’s internal world model of physics and racing to reason about the dynamics of our environment. For the remainder of this section, we focus on the inner workings of our LLM system shown in Figure 2. Subscripts are used to denote anything related to the i th generated reward function.

Reward Generation. For the first iteration, the environment context and task description g are combined into a single prompt that tasks an LLM to design reward functions for an RL algorithm. For future iterations, the prompt also contains information about the previous iteration’s best agent (see **Agent Selection**): text-based feedback about the agent’s behavior, the reward function R_{best} , and training diagnostics η_{best} expressed as the reward function’s value at different points during training. We ask the LLM to separate the reward function into multiple simpler components that are then combined by a weighted sum to obtain the final reward. The LLM is also tasked with setting the weights for each component. Reward functions are validated as “runnable” by calling them on randomly sampled states and actions. Generated reward functions that fail this test are iteratively re-generated by the LLM with the error trace included in the prompt until the generated functions pass the test [27, 19]. M different reward functions are independently sampled during this step.

Trajectory Alignment Filter. Training RL agents can be quite expensive, especially in a complex environment such as Gran Turismo. To reduce the number of RL agents trained at every iteration, we use the trajectory alignment coefficient (TAC) [24] based on a Kendall’s Tau-a correlation [14] to filter out potentially misaligned reward functions. To do so, assume access to a dataset of previously collected preferences, $(p_{ij}, \tau_i, \tau_j) \sim \mathcal{D}_{\text{pref}}$, where τ is a sub-trajectory of length L of states and actions $(s_1 a_1, s_2 a_2, \dots, s_L a_L)$, and $p_{ij} \in \{0, 1\}$ indicates a preference of $\tau_i \succ \tau_j$ if $p_{ij} = 0$ and $\tau_i \prec \tau_j$ otherwise. Let $G(\tau, R_m) = 1/L \sum_{t=1}^L R_m(s_t, a_t)$ be the average reward³ over a sub-trajectory τ for any given reward function R_m and $\tilde{p}_{ij}(R_m) \in \{0, 1\}$ is the preference induced by comparing $G(\tau_i, R_m)$ and $G(\tau_j, R_m)$. The trajectory alignment coefficient for a reward function R_m is then calculated as:

$$\sigma(R_m, \mathcal{D}_{\text{pref}}) = \frac{2P - |\mathcal{D}_{\text{pref}}|}{|\mathcal{D}_{\text{pref}}|} \text{ where } P := \sum_{(p_{ij}, \tau_i, \tau_j) \in \mathcal{D}_{\text{pref}}} p_{ij} \tilde{p}_{ij}(R_m) + (1 - p_{ij})(1 - \tilde{p}_{ij}(R_m)) .$$

²We acknowledge that effort was put into ensuring that these objects were properly named and commented before being given to the LLM.

³We use an average reward for the TAC to account for varying lengths of sub-trajectories and due to the very long horizon nature of the task which resembles a continual task in practice.

The TAC ranges from $[-1, 1]$. Intuitively, when $\sigma_m = -1$, the reward function R_m disagrees with all preference labels in $\mathcal{D}_{\text{pref}}$; when $\sigma_m = 1$, the preferences induced by R_m are in total alignment; when $\sigma_m = 0$, the reward function is no better than one that generates random values. Before training, only the top N best aligned reward functions, as indicated by σ_m , are kept for the following training step, where N is a hyperparameter smaller than the number of rewards sampled in the previous step M . In doing so, we hope to reduce computational costs associated with training, while searching over an equally large number of reward functions.

Training Any black-box RL algorithm can be used in this step so long as the following three outputs are obtained: a trained policy π^* , a sample trajectory of the trained policy τ , and training diagnostics η . The diagnostics are represented as a string capturing the reward function’s value at different points during training. For the first iteration, each of the N reward functions are trained from scratch. In subsequent iterations, all models are trained from scratch again, except the replay buffer from the previous iteration’s best agent (see **Agent Selection**) is used as a secondary replay buffer to speed up the learning process [38]. The additional experience from the secondary replay buffer can help reduce reinforcement learning training times [38, 26, 7], which can pose significant computational challenges in an iterative approach where RL is part of the inner loop. An ablation is done on the impact of the secondary seed buffer in Appendix A.

Agent Selection After training each of the N reward functions, our system is tasked to select the best performing policy for human evaluation. We elect to use a pairwise preference-based method due to its robustness [28, 40]. For each trained policy, a sufficiently long trajectory τ is sampled along with its corresponding video footage at ten frames per second. Then, a VLM is prompted to give a preference $p \in \{0, 1\}$ over all $\binom{N}{2}$ pairs of agents in the current iteration given both trajectory and video data for each agent (Appendix C). A Bradley-Terry model [4] is used to rank the agents, and the best policy is kept for the next iteration. Finally, all preferences p_{ij} and corresponding trajectories τ_i, τ_j are saved in the dataset $\mathcal{D}_{\text{pref}}$. The complete pipeline of our framework is shown in Algorithm 1.

5 Experiments

All experiments are done in the Gran Turismo 7 video game racing environment in a single car-track scenario: Lake Maggiore - Full Course⁴ with a Mercedes-AMG GT3 ’16. For all code generation, we use the latest available reasoning model available to us at the time provided by OpenAI, gpt-4o. Similarly, all preferences collected during agent selection are done using the latest vision-enabled model provided by OpenAI, gpt-4. 1. All human feedback is collected from one of the authors.

Training of the RL agent uses the approach described in [41] – a modified version of the soft actor-critic algorithm [10], QR-SAC. On average, a production-level baseline GT Sophy agent is trained

⁴Lake Maggiore is a popular track in GT 7 that features a diversity of interesting racing scenarios such as tight-harpins, both high and low speed corners, and a straight long enough to take advantage of slip stream.

Algorithm 1 Automated Reward Design for GT

Input: Empty dataset of preferences $\mathcal{D}_{\text{pref}}$, RL algorithm \mathcal{A}_E , initial task prompt prompt .

```

1: for  $i = 1$  to  $K$  do
2:   // Generate M reward functions
3:    $R_1, R_2, \dots, R_M \sim \text{LLM}(\text{prompt})$ 
4:   // Compute the alignment
   // coefficient for each reward
5:    $\sigma_m \leftarrow \text{Alignment}(R_m, \mathcal{D}_{\text{pref}})$ 
6:   // Keep only the top N rewards
7:    $R_1, R_2, \dots, R_N \leftarrow \text{top}_n(R_m, \sigma_m)$ 
8:   // Train an agent for each reward,
   // and get a policy, trajectory, and
   // training diagnostics
9:    $\pi_n^*, \tau_n, \eta_n \leftarrow \mathcal{A}(R_n)$ 
10:  // Get preference labels over the
   // newly trained policies
11:   $p_{ij} \leftarrow \text{VLM}(\tau_i, \tau_j) \quad \forall i, j \in \{N\}$ 
12:  // Get Bradley-Terry scores from
   // VLM labels
13:   $b_{1:N} \leftarrow \text{bradley}(\{p_{ij} \mid \forall i, j \in \{N\}\})$ 
14:  // Select best performing agent
15:   $best \leftarrow \text{argmax } b_{1:N}$ 
16:  // Get human feedback for best
   // policy
17:   $feedback \leftarrow \text{human}(\pi_{best})$ 
18:  // Construct new prompt
19:   $\text{prompt} \leftarrow \text{prompt} : feedback : R_{best} : \eta_{best}$ 
20:  // Add new preferences to dataset
21:   $\mathcal{D}_{\text{pref}} \leftarrow \mathcal{D}_{\text{pref}} + \{p_{ij} \mid \forall i, j \in \{N\}\}$ 
22: end for
23: return  $R_{best}$ 

```

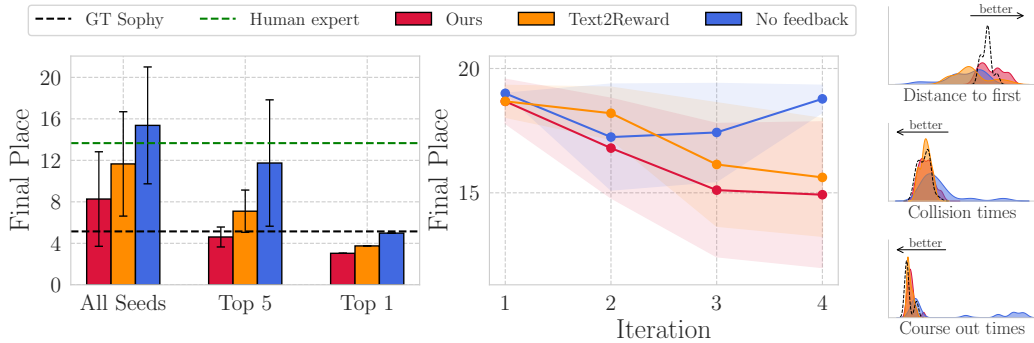


Figure 3: **Left:** Average final placements (\pm std) of final policies over all ten seeds, top five seeds only, and the best seed respectively. **Middle:** Average final placements (\pm std) of intermediate policies over all reward functions trained for each iteration of every seed. **Right:** Distribution of various racing metrics evaluated with the final agents from the top five seeds of each method.

for over 1500 epochs⁵ from scratch, which takes approximately ten days to train on a single A100 GPU and a dozen Playstation 4 systems. For our experiments, we reduced the training time of each iteration to 300 epochs, and only train the policy from the final iteration for 1500 epochs without a secondary replay buffer (to compare with GT Sophy) for evaluation. From now on, we refer to policies trained for 300 epochs as *intermediate* policies or agents, and final policies trained for 1500 epochs without a secondary replay buffer as *final* policies or agents. Three seeds of the final policy are trained for all experiments. Note that intermediate policies are unlikely to have converged even with the use of a secondary replay buffer, but are usually sufficient for preference elicitation in the agent selection step.

5.1 Outracing Human Experts

The goal of our main experiment is to produce an agent competitive with GT Sophy: a superhuman agent that complies with all standard motorsport racing rules. Our GT Sophy baseline was trained using the current state-of-the-art reward function meticulously designed over time with the help of both RL experts and game designers to achieve superhuman performance. The reward function for this agent is detailed in Appendix F. We also compare our agents to the performance of a human expert with over 25 years of experience in the Gran Turismo series who regularly ranks in the top 5% of players in the world. For the RGP, we purposefully choose a vague task description g to limit any prior bias we may have about what constitutes a good reward function in this task: "*win races while obeying standard motorsport racing rules and maintaining good sportsmanship*".

We run three different variations of our framework for four iterations ($K = 4$): (i) **Ours** generates ten different rewards ($N = 10$) and trains only the top five rewards ($M = 5$) from the trajectory alignment filter per iteration; **Text2Reward** generates only a single reward ($N = 1, M = 1$) per iteration, therefore skipping both the trajectory alignment filter and the agent selection steps which essentially recovers the Text2Reward framework [42]; (iii) **No feedback** also generates and trains ten and five reward functions respectively ($N = 10, M = 5$) per iteration, but instead replaces the human feedback by an LLM’s feedback of the agent’s behavior (Appendix D), making it a closed-loop system. All other hyperparameters follow the same parameters used to train GT Sophy.

5.1.1 Evaluation

For the same reasons that make reward design difficult for this task, evaluating agents for *good sportsmanship* or *rule abiding behavior* can be difficult, and is typically done using human expert judges for production-level agents. In order to systemically evaluate our agents in this work, we rely on two different racing statistics to provide a proxy for good sportsmanship and rule abiding behavior: *car collision time*, the total amount of time spent in contact with other cars, and *course out time*, the total amount of time spent off the legal racing area. Other works in autonomous racing also use collision times to assess proper sportsmanship [20, 9, 34]. We found it necessary to also consider

⁵One epoch is equivalent to 6000 training steps.

the time spent off course, since some agents were observed to gain an unfair advantage by taking illegal shortcuts during the race.

Policies are evaluated by starting a 4-lap race against 19 built-in AIs, an agent deployed by the game developers, with the agent to be evaluated beginning the race in last place. A hundred different races are collected to account for the variability of behaviors in the built-in AIs, and only trajectories in the 50% interquartile range of performance are kept. For evaluation, an agent’s final place defaults to last place if it exceeds a predetermined threshold for *car collision time* and *course out time*. Following the procedure in [20], these thresholds were chosen by training five seeds of a GT Sophy agent and selecting the highest *car collision time* and *course out time* among the five agents. In other words, agents who do not achieve a similar or better *car collision time* and *course out time* than GT Sophy are automatically disqualified and set to last place. By definition, all GT Sophy agents pass this threshold, therefore biasing our evaluations in favor of GT Sophy. All placements follow a zero-based indexing scheme where the best final placement is 0, and the worst is 19.

5.1.2 Results

The results of our experiments shown in Figure 3 and Table 1 tell us that LLMs can in fact produce agents which are not only competitive with GT Sophy, but can sometimes surpass the latter’s performance. While our system does not achieve this performance every time, the added ability to automatically search over multiple reward functions per iteration clearly improves both the consistency and the ceiling of our method compared to the Text2Reward baseline, without incurring any additional manual labor.

Interestingly, while the closed-loop system which relies solely on intrinsic feedback from the LLM is capable of producing expert-level reward functions, its average performance is markedly worse than a system incorporating human feedback; most likely due to its tendency to exceed the required thresholds for collision times and time spent off course. The progression of the generated reward functions’ performances shown in the middle of Figure 3 provides some additional insight: reward functions tend to diverge in performance without human grounding.

	Ours	Text2Reward	No feedback	GT Sophy	Human Expert
Valid agents	10/10	8/10	3/10	–	–
Better than Sophy	3/10	1/10	1/10	–	–
Best dist. to first	243.0 ± 9.1	269.1 ± 24.3	320.2 ± 3.6	317.8 ± 17.2	548.5 ± 50.3

Table 1: *Valid agents* counts the number of times the final policy for each seed of our framework passes the thresholds for good sportsmanship and rule-abiding behavior. *Better than Sophy* counts the number of agents where the average final placement during evaluation is *higher* than GT Sophy. *Best dist. to first* reports the **best** final agent’s (out of 10 seeds) average distance (± std) to the first car in meters during evaluation.

5.2 Replacing Human Expertise with VLMs

The preferences generated by the VLM during agent selection plays a crucial role in our framework. Aside from the direct impact they have towards replacing the fitness metric, the trajectory alignment filter also relies on its accuracy to prune reward functions. In the worst case, these preferences may be no better than random guessing, and the system would effectively be equivalent to generating a single reward function ($N = 1, M = 1$) per iteration.

In order to evaluate the accuracy of the preferences collected by the VLM, we first randomly select 10 iterations from those generated by "Ours" in the previous section. Each iteration contains five intermediate policies ($M = 5$) and five corresponding trajectories from which we can create ten possible pairs of trajectories. Then, ten human subjects familiar with the Gran Turismo 7 environment were asked to blindly label




Figure 4: **The VLM tends to select agents that are also favored by humans.** Frequency of a VLM’s preferences, using different forms of inputs, to select the best, second best, third best, fourth best, and fifth best agent in a set of five agents, as decided by human experts.

	Human to humans	VLM to humans	VLM- τ to humans	VLM- I to humans
Agree rate	78.54%	74.14%	71.48%	62.96%

Table 2: Agreement rate of different preference sources with human experts. VLM, VLM- τ , VLM- I indicate different input modalities. Specifically, using both trajectory data and video data, trajectories only, and videos only respectively.

a subset (twenty to thirty) of these pairs given the same task description g defined above to create a small dataset of ground truth pairwise preferences, ensuring that each pair is labeled by at least two different individuals.

To evaluate the VLM’s image processing capabilities, we collect preferences from three different input modalities: (i) **VLM** collects preferences using both trajectory data and video data, which is the version used for the experiments in Section 5.1; (ii) **VLM- τ** uses trajectory data only (no images); (iii) **VLM- I** collects preferences using only video data (images only). From the agreement rates shown in Table 2, we observe that the preferences generated by the VLM have an agreement rate nearly as high as the percentage of times humans agree with each other, the latter of which can be seen as an upper bound for the VLM. The accuracy of the VLM also appears to drop significantly when given only access to videos of the agents’ behaviors, indicating a gap in performance in today’s vision foundation models, since humans rely only on videos to judge performance.

In practice, Figure 4 shows how often the selected agents in the main experiment corresponded to the humans’ top ranked agent within an iteration, the second best agent, the third best agent, and so on. Overall, the preferences from the agent selection step serve as an effective replacement for human expertise, promising an alternative to manual fitness metric design for automated reward function evaluation.

5.3 The Trajectory Alignment Filter Relies on Data

After collecting a dataset of reliable preferences, the TAC can be computed for a new reward function to measure its alignment with respect to the dataset and consequently the data generating process. Although it has not been theoretically proven, the intuition behind the trajectory alignment filter is that reward functions which are better aligned with the preference generator also produce agents that are ultimately preferred by the same generator. To verify this empirically, we reuse all preferences generated during the experiments in Section 5.1 to randomly create varying sizes of preference datasets, denoted \mathcal{D}^x where x is the size of the dataset, all originating from the same preference generator, the VLM. Let $\mathcal{D}^{\text{full}}$ be the complete dataset of 400 preferences. For different values of x , we compute how accurate the TAC for a reward function is at predicting post-training preferences with the following equation:

$$\text{Acc}(\mathcal{D}^x) := \frac{1}{|\mathcal{D}^{\text{full}}|} \sum_{(p_{ij}, \tau_i, \tau_j) \sim \mathcal{D}^{\text{full}}} p_{ij} \hat{p}_{ij} + (1 - p_{ij})(1 - \hat{p}_{ij}) ,$$

$$\hat{p}_{ij} = 1 \text{ if } \sigma(R_i, \mathcal{D}^x) > \sigma(R_j, \mathcal{D}^x) \text{ else } 0 .$$

The results of this experiment shown in Figure 5a) lead to a concerning conclusion: the trajectory alignment coefficient is only a reliable indicator for success after a certain dataset size is reached. In our case, the predictive accuracy of the TAC is only comfortably better than random guessing after collecting around 100 preferences. Unfortunately, the experiments in Section 5.1 only collected 10 preferences per iteration, and reached a maximum of four iterations. Therefore, the trajectory alignment filter was unlikely to be very beneficial in our results. However, the TAC can still be a very cheap method to quickly evaluate new reward functions provided a large enough dataset of preferences is available, either by bootstrapping an initial set of preferences or by scaling M and/or K towards larger values, which we leave for future work.

5.4 On the Novelty of Generated Reward Functions

To measure the *novelty* of the generated rewards by our system, we compute the correlation of all the generated rewards from "**Ours**" against the rewards used to train GT Sophy. From a diverse set of trajectories (Appendix H.1), rewards are computed for *all* LLM-generated reward functions and compared against a reference reward function to measure their correlations with the reference rewards.

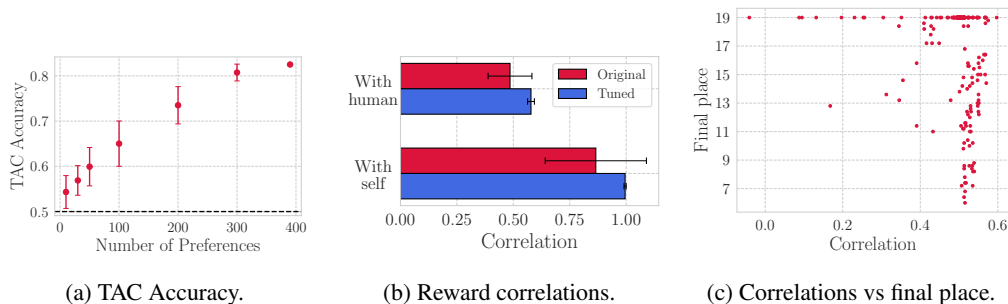


Figure 5: **a)** The trajectory alignment coefficient’s accuracy at predicting a reward function’s post-training performance for varying sizes of preference datasets. The dashed line represents the baseline accuracy of a random guesser. **b)** The average correlation (\pm std) of all generated reward functions with respect to rewards from GT Sophy (with human) and the best LLM-based agent (with self) before (original) and after (tuned) optimizing the reward component weights towards the reference policy. **c)** The relationship between the correlation of a reward function with respect to the rewards from GT Sophy, and its intermediate policy’s performance.

Two different reference reward functions are used to evaluate two different forms of novelty. First, the correlations are computed against the human designed rewards used to train GT Sophy (‘With human’) to see if the generated rewards are novel with respect to humans. Second, the correlations are computed against a reward function generated by the LLM which achieved better performance than GT Sophy (‘With self’) to see if the generated rewards are novel with respect to itself. The resulting correlations are shown in red in Figure 5b).

Recall that the generated reward functions are composed of multiple smaller components, each with a weight coefficient also generated by the LLM. To evaluate the qualitative novelty of reward functions, and not just potential differences in weights, we also compute the correlations with the *tuned* reward functions with respect to the reference rewards, where the weight coefficients are tuned to maximize the correlation with the reference rewards (Appendix H.2). With this change, the blue bars in Figure 5b) sees improved correlations, but the rewards are still relatively novel compared to the ones used to train GT Sophy.

Further, notice that the average tuned correlation of the LLM’s reward functions computed against itself is nearly 1.0, suggesting a very low degree of qualitative diversity within the generated reward functions. In addition, since the reference rewards in this case achieved better performance than the baseline GT Sophy agent, one potential solution to consistently generate high performing agents is to improve the weight tuning process. In our framework, no special care was taken towards tuning these weights: the LLM handled the weight selection concurrently with generating the reward function code. This remains an open problem to be solved in future research.

Finally, the plot in Figure 5c) shows the relationship between a reward function’s correlation with GT Sophy and the final placement achieved by its corresponding intermediate policy. The Pearson correlation between the x and y axis on this plot is -0.177 ; there is no strong relationship between rewards that are closer to the human designed ones and their performance, which further illustrates that rewards designed by humans are not necessarily optimal.

5.5 Generating Novel Behaviors

In our final set of experiments, we use our system to generate two novel behaviors not yet attempted with RL in GT 7. Using the same hyperparameters as before, the system is run on the following task descriptions: "*race as fast as possible in reverse at all times while otherwise obeying standard motorsport rules*" and "*race normally except the agent should drift as much as possible while otherwise obeying standard motorsport rules*". We found that a single seed was sufficient to achieve the desired behaviors in these cases based on human evaluations, although we cannot comment on their optimality. Videos of the trained policies for both novel behaviors can be found here.

6 Conclusion and Limitations

Modern reinforcement learning algorithms present a powerful tool set to produce impressive agents through proper reward specification. In an environment like Gran Turismo, this led to the development of GT Sophy, a superhuman autonomous racing agent. While GT Sophy was trained without explicit human demonstrations, experts had to guide the agent’s behavior through reward functions. These reward functions are an unnatural form of communication for most, especially those unfamiliar with RL. In this work, we demonstrated how today’s foundation models can automate the reward design process for GT 7, creating a direct pipeline from text to policy by leveraging the large-scale domain expertise contained within a foundation model as a useful prior for learning. Our framework was shown to produce superhuman agents more consistently than prior work, as well as to generate novel reward functions and behaviors in a previously unexplored environment for automated reward design. Even though our system was only evaluated on the GT 7 environment, none of its components are specific to this environment. We leave the exploration of its broader applicability in other environments for future work.

Our method shares several limitations with prior approaches to LLM-coding-based automated reward design. These approaches can be very computationally expensive to run, since the outer loop trains multiple RL agents in the inner loop and the foundation models themselves must be quite large to write functional reward code. In addition, our system still relies on a human in the loop to achieve consistent results. Although we take steps to alleviate these limitations, for example with a secondary replay buffer and automatic VLM evaluations, further work can be done in this space towards cheaper and more reliable automated reward design.

References

- [1] M. Awais, M. Naseer, S. Khan, R. M. Anwer, H. Cholakkal, M. Shah, M.-H. Yang, and F. S. Khan. Foundation models defining a new era in vision: a survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.
- [2] A. P. Badia, B. Piot, S. Kapturowski, P. Sprechmann, A. Vitvitskiy, D. Guo, and C. Blundell. Agent57: Outperforming the atari human benchmark, 2020.
- [3] S. Booth, W. B. Knox, J. Shah, S. Niekum, P. Stone, and A. Allievi. The perils of trial-and-error reward design: misdesign through overfitting and invalid task specifications. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 5920–5929, 2023.
- [4] R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- [5] T. Carta, P.-Y. Oudeyer, O. Sigaud, and S. Lamprier. Eager: Asking and answering questions for automatic reward shaping in language-guided rl. *Advances in neural information processing systems*, 35:12478–12490, 2022.
- [6] L. Chen, P. Wu, K. Chitta, B. Jaeger, A. Geiger, and H. Li. End-to-end autonomous driving: Challenges and frontiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [7] P. D’Oro, M. Schwarzer, E. Nikishin, P.-L. Bacon, M. G. Bellemare, and A. Courville. Sample-efficient reinforcement learning by breaking the replay ratio barrier. In *Deep Reinforcement Learning Workshop NeurIPS 2022*, 2022.
- [8] A. Escontrela, A. Adeniji, W. Yan, A. Jain, X. B. Peng, K. Goldberg, Y. Lee, D. Hafner, and P. Abbeel. Video prediction models as rewards for reinforcement learning. *Advances in Neural Information Processing Systems*, 36:68760–68783, 2023.
- [9] F. Fuchs, Y. Song, E. Kaufmann, D. Scaramuzza, and P. Dür. Super-human performance in gran turismo sport using deep reinforcement learning. *IEEE Robotics and Automation Letters*, 6(3):4257–4264, 2021.
- [10] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. Pmlr, 2018.
- [11] X. Han, Q. Yang, X. Chen, Z. Cai, X. Chu, and M. Zhu. Autoreward: Closed-loop reward design with large language models for autonomous driving. *IEEE Transactions on Intelligent Vehicles*, 2024.

- [12] R. Hazra, A. Sygkounas, A. Persson, A. Loutfi, and P. Z. D. Martires. Revolve: Reward evolution with large language models using human feedback. *arXiv preprint arXiv:2406.01309*, 2024.
- [13] Z. Huang, Z. Sheng, Y. Qu, J. You, and S. Chen. Vlm-rl: A unified vision language models and reinforcement learning framework for safe autonomous driving. *arXiv preprint arXiv:2412.15544*, 2024.
- [14] M. G. Kendall. The treatment of ties in ranking problems. *Biometrika*, 33(3):239–251, 1945.
- [15] M. Klissarov, P. D’Oro, S. Sodhani, R. Raileanu, P.-L. Bacon, P. Vincent, A. Zhang, and M. Henaff. Motif: Intrinsic motivation from artificial intelligence feedback. In *NeurIPS 2023 Foundation Models for Decision Making Workshop*, 2023.
- [16] M. Klissarov, M. Henaff, R. Raileanu, S. Sodhani, P. Vincent, A. Zhang, P.-L. Bacon, D. Precup, M. C. Machado, and P. D’Oro. Maestromotif: Skill design from artificial intelligence feedback. *arXiv preprint arXiv:2412.08542*, 2024.
- [17] W. B. Knox, A. Allievi, H. Banzhaf, F. Schmitt, and P. Stone. Reward (mis) design for autonomous driving. *Artificial Intelligence*, 316:103829, 2023.
- [18] M. Kwon, S. M. Xie, K. Bullard, and D. Sadigh. Reward design with language models. *arXiv preprint arXiv:2303.00001*, 2023.
- [19] H. Le, Y. Wang, A. D. Gotmare, S. Savarese, and S. C. H. Hoi. Coderl: Mastering code generation through pretrained models and deep reinforcement learning. *Advances in Neural Information Processing Systems*, 35:21314–21328, 2022.
- [20] H. Lee, T. Seno, J. J. Tai, K. Subramanian, K. Kawamoto, P. Stone, and P. R. Wurman. A champion-level vision-based reinforcement learning agent for competitive racing in gran turismo 7. *IEEE Robotics and Automation Letters*, 2025.
- [21] H. Li, X. Yang, Z. Wang, X. Zhu, J. Zhou, Y. Qiao, X. Wang, H. Li, L. Lu, and J. Dai. Auto mc-reward: Automated dense reward design with large language models for minecraft. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16426–16435, 2024.
- [22] Y. J. Ma, W. Liang, G. Wang, D.-A. Huang, O. Bastani, D. Jayaraman, Y. Zhu, L. Fan, and A. Anandkumar. Eureka: Human-level reward design via coding large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- [23] J. MacGlashan, A. Devlic, and P. MacAlpine. Don’t cross that line! how our ai agent learned sportsmanship. <https://ai.sony/blog/Dont-Cross-That-Line!-How-Our-AI-Agent-Learned-Sportsmanship/>, 2022. Accessed: 2025-08-21.
- [24] C. Muslimani, K. Johnstonbaugh, S. Chandramouli, S. Booth, W. B. Knox, and M. E. Taylor. Towards improving reward design in rl: A reward alignment metric for rl practitioners. *arXiv preprint arXiv:2503.05996*, 2025.
- [25] A. Y. Ng, D. Harada, and S. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *Icml*, volume 99, pages 278–287. Citeseer, 1999.
- [26] E. Nikishin, M. Schwarzer, P. D’Oro, P.-L. Bacon, and A. Courville. The primacy bias in deep reinforcement learning. In *International conference on machine learning*, pages 16828–16847. PMLR, 2022.
- [27] T. X. Olausson, J. P. Inala, C. Wang, J. Gao, and A. Solar-Lezama. Demystifying gpt self-repair for code generation. *CoRR*, 2023.
- [28] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- [29] A. Remonda, S. Krebs, E. Veas, G. Luzhnica, and R. Kern. Formula rl: Deep reinforcement learning for autonomous racing using telemetry data. *arXiv preprint arXiv:2104.11106*, 2021.
- [30] J. Rocamonde, V. Montesinos, E. Nava, E. Perez, and D. Lindner. Vision-language models are zero-shot reward models for reinforcement learning. *arXiv preprint arXiv:2310.12921*, 2023.
- [31] B. Roziere, J. Gehring, F. Gloeckle, S. Sootla, I. Gat, X. E. Tan, Y. Adi, J. Liu, R. Sauvestre, T. Remez, et al. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*, 2023.

- [32] I. Singh, V. Blukis, A. Mousavian, A. Goyal, D. Xu, J. Tremblay, D. Fox, J. Thomason, and A. Garg. Progprompt: Generating situated robot task plans using large language models. *arXiv preprint arXiv:2209.11302*, 2022.
- [33] S. Singh, R. L. Lewis, and A. G. Barto. Where do rewards come from. In *Proceedings of the annual conference of the cognitive science society*, pages 2601–2606. Cognitive Science Society, 2009.
- [34] Y. Song, H. Lin, E. Kaufmann, P. Dürr, and D. Scaramuzza. Autonomous overtaking in gran turismo sport using curriculum reinforcement learning. In *2021 IEEE international conference on robotics and automation (ICRA)*, pages 9403–9409. IEEE, 2021.
- [35] S. Sontakke, J. Zhang, S. Arnold, K. Pertsch, E. Bıyık, D. Sadigh, C. Finn, and L. Itti. Roboclip: One demonstration is enough to learn robot policies. *Advances in Neural Information Processing Systems*, 36:55681–55693, 2023.
- [36] S. Sun, R. Liu, J. Lyu, J. Yang, L. Zhang, and X. Li. A large language model-driven reward design framework via dynamic feedback for reinforcement learning. *CoRR*, 2024.
- [37] C. Tang, B. Abbatematteo, J. Hu, R. Chandra, R. Martin-martin, and P. Stone. Deep reinforcement learning for robotics: A survey of real-world successes. *arXiv preprint arXiv:2408.03539*, 2024.
- [38] D. Tirumala, T. Lampe, J. E. Chen, T. Haarnoja, S. Huang, G. Lever, B. Moran, T. Hertweck, L. Hasenclever, M. Riedmiller, et al. Replay across experiments: A natural extension of off-policy rl. In *The Twelfth International Conference on Learning Representations*, 2023.
- [39] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- [40] C. Wirth, R. Akrou, G. Neumann, and J. Fürnkranz. A survey of preference-based reinforcement learning methods. *Journal of Machine Learning Research*, 18(136):1–46, 2017.
- [41] P. R. Wurman, S. Barrett, K. Kawamoto, J. MacGlashan, K. Subramanian, T. J. Walsh, R. Capobianco, A. Devlic, F. Eckert, F. Fuchs, et al. Outracing champion gran turismo drivers with deep reinforcement learning. *Nature*, 602(7896):223–228, 2022.
- [42] T. Xie, S. Zhao, C. H. Wu, Y. Liu, Q. Luo, V. Zhong, Y. Yang, and T. Yu. Text2reward: Reward shaping with language models for reinforcement learning. In *The Twelfth International Conference on Learning Representations*, 2024.
- [43] Y. Yang, M. Chen, Q. Qiu, J. Wu, W. Wang, B. Lin, Z. Guan, and X. He. Adapt2reward: Adapting video-language models to generalizable robotic rewards via failure prompts. In *European Conference on Computer Vision*, pages 163–180. Springer, 2024.
- [44] W. Yu, N. Gileadi, C. Fu, S. Kirmani, K.-H. Lee, M. G. Arenas, H.-T. L. Chiang, T. Erez, L. Hasenclever, J. Humplik, et al. Language to rewards for robotic skill synthesis. In *Conference on Robot Learning*, pages 374–404. PMLR, 2023.
- [45] C. B. C. Zhang, Z.-W. Hong, A. Pacchiano, and P. Agrawal. Orso: Accelerating reward design via online reward selection and policy optimization. In *ICML 2024 Workshop: Aligning Reinforcement Learning Experimentalists and Theorists*, 2024.

A Ablation on Secondary Replay Buffer

We conduct an additional ablation to see the impact of the secondary replay buffer on training. On five randomly chosen final agents from **Ours**, we train the agent for the full 1500 model clocks with and without a secondary seed buffer in Figure 6. We observe that while the secondary replay buffer can be quite beneficial for earlier epochs, the policies ultimately converge to a similar behavior in either case.

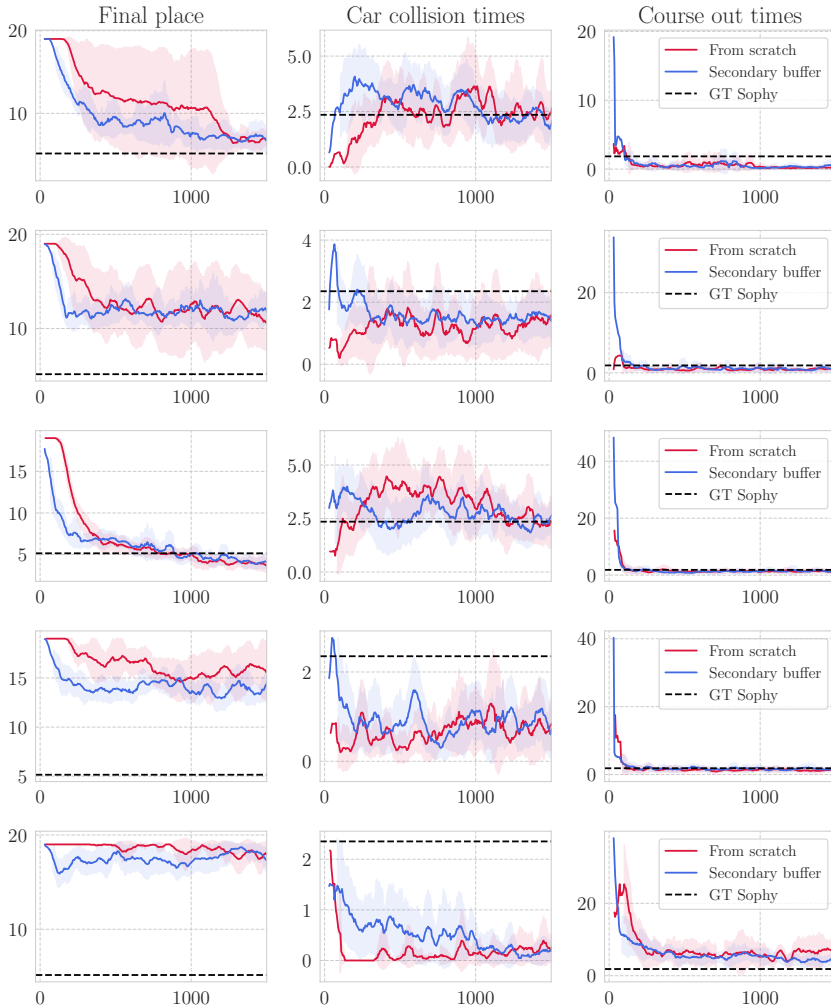


Figure 6: Training curves for five different generated final reward functions from **Ours** (rows). Each column shows a different training metric. Each reward function is trained with identical hyperparameters, except one is trained with a secondary replay buffer from the previous iteration, and one is trained from scratch. The x -axis for all plots are the number of epochs.

B Reward Generation Prompts

We separate the task of generating rewards into two prompts. First, we ask the LLM to generate an english-readable overview of the rewards. Then, the answer is provided to the LLM again to generate the corresponding Python code. There is also a separate set of prompts for generating the initial reward functions, and generating subsequent reward functions which improve upon the previous iteration.

First the prompt used to generate the english-readable overview of the initial reward functions for the first iteration.

Initial English Rewards Generation

You are an expert in motorsports, Gran Turismo 7, and reinforcement learning. I would like to design a reward function to train an agent to race in Gran Turismo 7. The goal of the agent is to {goal}.

Help me design a reward function for the above behavior by separating the total reward into several component rewards to facilitate future evaluations and improvements. Each individual reward function component should generate a separate reward for each of the N cars on the track. Each reward function component should have the following function signature, and output a numpy array of shape (N, 1):

```
““def reward_component(current_obs: GTObs, prev_obs: GTObs, course: Course) -> np.ndarray:““
```

Describe in words the high level function of each different reward component. Do not provide any code.

The code describing the attributes of the GTObs and Course objects in the reward input is provided below. You are not obligated to use all attributes.

```
““\n{obs_code}\n““
```

Directly answer in the following format for each reward function component. The reward name should follow python function definition naming conventions.

```
<|reward name|> Name here <|reward name|>  
<|reward description|> Description here <|reward description|>
```

Additional information:

- Do not provide any code, only provide the reward names and descriptions.
- The game designers have defined being out of bounds as having any of the four tires off the racing area.
- Keep the complexity and number of components to a minimum. It will be possible to further iterate and improve on the rewards at a later time.

Then, the output from the previous prompt is given to the following prompt to generate the actual Python code used for RL training.

Initial Code Rewards Generation

You are an expert in motorsports, Gran Turismo 7, reinforcement learning, and python code generation.

I would like to implement the following reward components. Each component has a function name wrapped by <|reward name|>, and a description of the reward function wrapped by <|reward description|>.

```
{reward_components_desc}
```

Each individual reward function component should generate a separate reward for each of the N cars on the track. Each reward function component should have the following function signature, and output a numpy array of shape (N, 1) with a float data type:

```
““def reward_name(current_obs: GTObs, prev_obs: GTObs, course: Course) -> np.ndarray:““
```

Generate code for each of the described reward function components. The final reward will be a weighted sum of each of these component reward functions. Within the block of code, after all function definitions, suggest a weight coefficient for each component in the following format:

```
““weights = {"first_reward_name": weight, "second_reward_name": weight, ...}““
```

The code describing the attributes of the GTObs and Course objects in the reward input is provided below.

```
““\n{obs_code}\n““
```

Additional information:

- All reward function components should be contained within a single code block
- There should be no helper functions. The only function definitions within the generated block should be the reward components themselves and the weights dictionary. There should be no other function definitions.

For all subsequent iterations, the following pair of prompts are used in a similar fashion to generate an english overview of the improvements to implement, and to generate the Python code once more.

Subsequent Improvements Suggestions

You are an expert in motorsports, Gran Turismo 7, and reinforcement learning. We would like to train a reinforcement learning agent in Gran Turismo 7. The goal of the agent is to {goal}
Read the following code describing the attributes of the GTObjs and Course objects, which are used to design reward functions.

```
“python  
{obs_code}  
““
```

To achieve the above goal, we design reward functions iteratively by improving on the previous reward functions. After every iteration, the reward functions are used to train an agent, and improved following a set of instructions based on feedback from the trained agent. Thus far, we have designed {num_iters} round(s) of reward function components based on the following instructions.

```
{all_english_rewards}
```

Finally, the most recent set of instructions were used to write the following most recent reward function components.

```
“python  
{reward_code}  
““
```

These reward components were used to train a reinforcement learning agent to {goal}. The total reward used to train the agent is a weighted sum of these rewards using the weights dictionary at the bottom of the code block.

We evaluated our agent every 100 epochs during training. Evaluations are done by placing the agent in a race against 19 other cars. The other cars are controlled by a built-in AI provided by the video game. The agent always begins the race in last place. An evaluation episode terminates after completing {max_laps} laps, or after {max_episode_duration} seconds has elapsed. The evaluation course has a maximum fixed_course_v of {course_max_v} meters. We track the total reward for each individual reward component over an evaluation episode. The rewards reported are already multiplied by their corresponding reward weights. The total values for each reward component for one evaluation episode is shown below as a list. They are recorded at intervals of 100 training epochs, showing the agent's progress during training for each reward component. Additionally, the return is also shown, which is simply the sum of all weighted reward components.

```
““  
{moi_data}  
““
```

After examining videos of the trained agent, we also provide the following *ground truth* feedback :

```
{human_feedback}
```

Carefully analyze the current reward function code in conjunction with the training progress for each reward component. First, provide a short summary of the training run given the available information. This summary should be wrapped by a <|training summary|> tag. Then, take everything into consideration to suggest improvements to the current reward functions. You may do one of four things, which should follow the format:

1. Modify reward weights: <|modify reward weight|> existing reward function name <|modify reward weight|>\n <|description|> {{Current weight}} --> {{New weight}}, followed by a short explanation. <|description|>
2. Modify the implementation of existing reward functions: <|modify reward|> existing reward function name <|modify reward|>\n<|description|> Description of modifications <|description|>
3. Completely remove existing reward functions: <|remove reward|> existing reward function name <|remove reward|>
4. Add new reward function: <|new reward|> new reward function name <|new reward|> \n<|description|> Description here <|description|>

At the end, summarize all suggestions wrapped by a <|suggestions summary|> tag.

Additional information:

- Do not provide any code.
- The game designers have defined being out of bounds as having any of the four tires off the racing area.
- Limit the quantity and complexity of suggested changes, as performing too many significant changes at once can lead to unexpected results. For example, significantly modifying existing reward components or adding new reward components should be avoided unless absolutely necessary.

Subsequent Improvements Implementation

Generate code for each reward function component once more, this time incorporating all suggested changes. Remember to include the weights for the reward components in a dictionary at the end of the code block.

Additional information:

- All reward function components should be contained within a single code block
- There should be no helper functions. The only function definitions within the generated block should be the reward components themselves and the weights dictionary. There should be no other function definitions.

C Eliciting Preferences from the VLM

Evaluation videos of the trained agents range anywhere from 10 to 30 minutes to complete four laps of the designated track. We only provide the first lap or 3 minutes (whichever comes first) of video data for the VLM to evaluate at a frame rate of 10 frames per second. In order to reduce the context length input to the VLM, we further split the video into 15 seconds clips. The VLM then provides a description of the two agents' behaviors in each clip, then all descriptions are fed into an LLM as text for the LLM to decide which agent was preferred. We found this to be the best compromise in order to analyze long-form videos at a relatively high frame rate, evaluating proper behavior in Gran Turismo 7 requires high frame rate fidelity. The prompt used by the VLM to generate text summaries of a pair of 15 seconds clips is given below.

Summarizing Videos Into Text

You are an expert in motorsports and Gran Turismo. We trained two agents to play the Gran Turismo video game. The primary goal of both agents is to {goal}.

We trained two agents to play the Gran Turismo video game. The primary goal of both agents is to {goal}. You will be given information about both agents' trajectories, by receiving frame-by-frame data coming directly from the environment's API. The data will be coming at {fps} frames per second, and represents a {clip_length} second video clip of each agent's behavior. Documentation about each frame's data is given here:\n {variables_documentation}\n

The first video clip is shown below.
{frames_and_observations_1}

Now watch this second video clip
{frames_and_observations_2}

Describe and compare in detail any and all of the agents' behaviors in each clip as it pertains to the desired goal.

\nAdditional information: \n

- The game designers have defined being out of bounds as having any of the four tires off the racing area.\n
 - There should be enough detail to evaluate the agents without watching the video. \n
 - Remain as objective as possible for all descriptions,
- as to not bias anyone towards whether the goal was achieved or not for each agent.\n
- Directly answer in the following format: description of agent 1 in clip 1, description of agent 2 in clip 2, and a comparison between the two clips.

Then, the text descriptions of all 15 seconds clips are combined to create a long description of the entire 3 minute video. The prompt used to gather the preference from an LLM over the stitched text-based descriptions of the two agent's behaviors is shown:

Getting Preference over Agents

You are an expert in motorsports and Gran Turismo. We trained two agents to play the Gran Turismo video game. The primary goal of both agents is to {goal}.

The final behavior of the two agents were recorded into two {video_duration} second videos, recorded at {fps}. The two videos, one for each agent, were then analyzed and compared clip by clip for each agent. The first clip always describes the first agent, and the second clip always describes the second agent.

The result of comparing every pair of clips in the complete video is shown below.
\n\n''{full_comparison_text}''\n\n

After carefully analyzing the descriptions of each agent’s behavior for each clip, declare which agent is overall preferable with respect to achieving or making progress towards the desired goal for the entire video.

Answer in the following format: \n
'("preferred_agent": 1) if the first agent is preferred or\n'
'("preferred_agent": 2) if the second agent is preferred.\n'
Include also a brief explanation for your answer.\n

The LLM was not given the option to provide ties. We found it was important to use a contrastive approach to gather clip descriptions. To do so, the VLM was always fed two clips side-by-side for more accurate descriptions of an agent’s behavior, since inaccuracies are not as important as getting a relative sense of which agent was better.

D Replacing Human Feedback with an LLM’s

In order to conduct the closed-loop experiments from the **No Feedback** set of runs, we replaced all human feedback in the prompts with one generated automatically by an LLM following the agent descriptions obtained in Appendix C. The prompt used to gather this description is given here.

Getting Agent Behavior Descriptions

Carefully analyze the descriptions of each agent’s behavior for each clip. Then, provide a detailed summary of agent {agent_to_describe}’s behavior only, especially as it relates to the stated goal.

Additional notes: \n

- The game designers have defined being out of bounds as having any of the four tires off the racing area.\n

- Remain objective for all descriptions."

Do not give opinions on whether certain behaviors are 'good' or 'bad'.\n

- Make no mention of the other agent in the summary.\n

- Start your response with '## Agent Summary'.

Simply refer to the agent as 'the agent', instead of 'agent 1' or 'agent 2'.

E Reinforcement Learning Hyperparameters

We present the relevant hyperparameters used to train all agents in Table 3. Hyperparameters were selected from the baseline GT Sophy agent. Additional hyperparameters used only for our method (secondary replay buffer) were selected following a small grid search.

F GT Sophy Rewards

The following list describes the reward components used by the baseline Sophy. In general, we found that the generated reward components captured the same categories of rewards and penalties: progress, collisions, track violations, and passing. Penalties for steering smoothness were generated occasionally.

- **progress**: A reward for the amount of progress made along the track center-line between observations.

Table 3: Relevant Training Parameters

Hyperparameter	Value
Optimizer	Adam
Policy learning rate	1.25e-5
Critic learning rate	2.5e-5
Entropy coefficient α	0.01 (fixed)
Discount	0.9896
n-step	7
Number of Quantiles	32
Parameter copy τ	0.005
Secondary replay buffer sampling ratio	0.2
Secondary replay buffer size	100 000
Replay buffer size	1 000 000

- **off_course_penalty**: Penalty for the maximum amount of time that any one tire was off course since the last observation.
- **wall_penalty**: Penalty for hitting the wall.
- **rel_v2_collision_penalty**: A complex penalty that tries to determine collision fault. If our car is at fault it is penalized by the square of the difference in velocity vectors.
- **constant_car_collision**: A penalty for any collision with other cars.
- **passing_reward**: A reward for gaining position on opponents and a penalty for losing position on opponents. Only considers opponents within certain distance range.
- **steering_change_cost**: A penalty for the magnitude of change in steering angle. Used to discourage large sudden changes.
- **acted_steering_history_cost**: A penalty for making many changes in steering direction in a short window of time.

G LLM Generated Reward Functions

Instead of providing the raw reward functions generated by the LLM, we conduct an informal qualitative analysis of the generated functions to better understand the diversity of components, and how they relate to the human designed reward functions in the previous section. We identified the following ten most popular categories of reward components generated by the LLM. For each category, there can be several different implementations and approaches to penalizing or rewarding the agent.

- **collision_avoidance**: Reward to discourage collisions through instantaneous penalties or rewarding continuous safe driving.
- **track_boundary**: Penalty for going out of bounds.
- **forward_progress**: Reward for making forward progress through the track.
- **driving_smoothness**: Penalty for inhuman-like movements, or jerky movements.
- **clean_overtake**: Reward for overtaking an opponent while avoiding any collisions.
- **rank_improvement**: Reward for improving the agent’s current position on the track relative to other opponents.
- **optimal_path_following**: Reward for staying near the centerline of the track.
- **finish_line_bonus**: Reward for crossing the finish line.
- **gap_closure**: Reward for closing in on the next closest opponent ahead of the agent.
- **off_track_overtake**: Penalty for overtaking opponents while outside the track boundaries

Category	Percent of Runs
collision_avoidance	100.0%
track_boundary	100.0%
forward_progress	100.0%
driving_smoothness	91.0%
clean_overtake	59.0%
rank_improvement	58.5%
optimal_path_following	55.0%
finish_line_bonus	15.5%
gap_closure	8.5%
off_track_overtake	3.5%

Table 4: Different categories of reward components generated by the LLM from all seeds in **Ours**, and the percentage of generated reward functions containing some form of the specified component. In total, 200 reward functions were generated and analyzed, twenty for each of the ten seeds run for our method.

H Additional Notes on Reward Correlations and Performance

H.1 Sampling Trajectories for Correlation Experiments

Trajectories were generated using one of the 300 epochs runs generated by our system. For every 10th epoch we used the saved policy to run a one lap race against 19 built-in-ai agents, with our target policy starting in 10th position. This gives a range of experiences resulting from the varying level of skill produced at different stages of training. Further, by placing our agent in the middle of the track this encourages more interaction with other cars even when the racing skill might be low.

H.2 On Tuning the Reward Weight Coefficients

As mentioned in Section 5.4, one potential improvement to the framework is to focus on tuning the weight coefficients of the reward components themselves instead of the qualitative nature of the rewards. This is due to the observation that nearly all reward functions generated by the LLM, even those with poor performance, can have their weights modified to achieve near perfect correlation with a reward function with *good* performance. In this section, we test the hypothesis that tuning the reward weights alone can be sufficient to achieve more consistent results.

Experimental setup First, we select the best performing seed from our method and use its reward function as a reference reward. Then, we tune the reward weights from the final reward function of two other random seeds (out of ten) with poor performance with respect to the reference rewards with good performance. The reward weights are tuned with stochastic gradient descent from the trajectories collected in Appendix H.1 in two stages: the weights are initially optimized to maximize correlation, then the weights are multiplied by a constant that is optimized to minimize the mean squared error to account for differences in scale. Finally, the agent is re-trained with the tuned reward weights.

Results Figures 7 and 8 report the results for two sub-optimal seeds. Although the tuned rewards are almost perfectly correlated with the reference reward, there may still be a gap in performance in the tuned rewards performance. For example, the results in Figure 8 show improved placements, but at the cost of increased collisions that do not quite match the expected performance from the reference reward function.

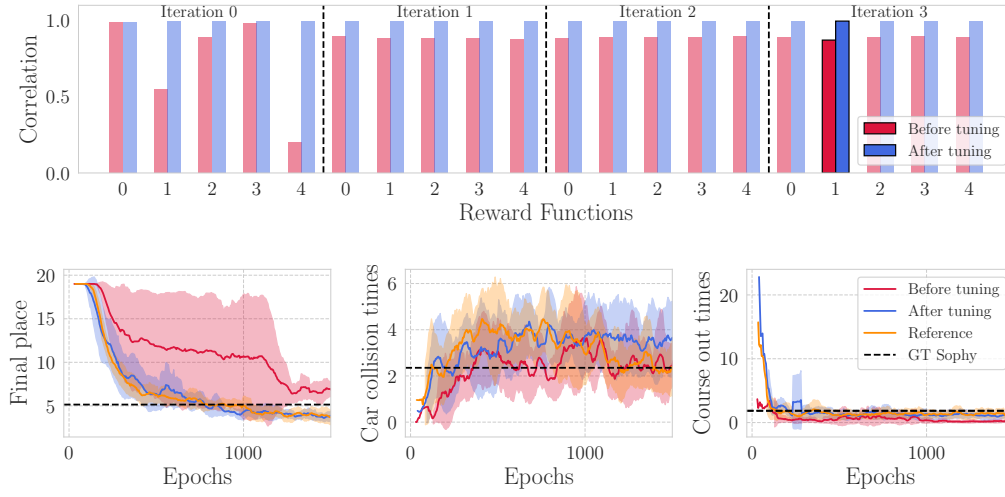


Figure 7: **Tuning weight coefficients can improve performance.** *Top:* Correlations of generated reward functions for all iterations of a single seed of **Ours** both before and after tuning the weights with respect to the best LLM-based agent’s reward function. The darker bars represent the chosen reward function for the final iteration. *Bottom:* Training curves of the final chosen reward function of the same seed. We report the mean final place, car collision times in seconds, and course out times in seconds across three seeds (\pm std)

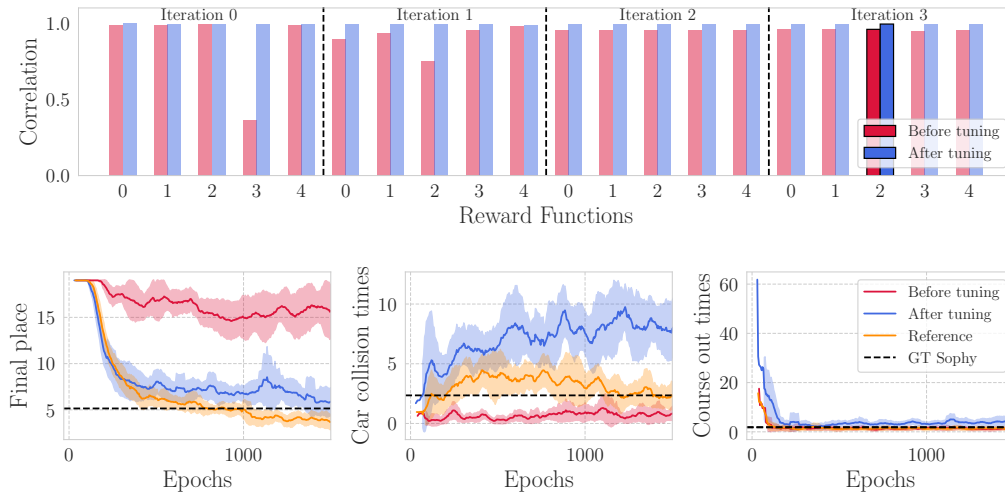


Figure 8: **Tuning weight coefficients is not always sufficient.** *Top:* Correlations of generated reward functions for all iterations of a single seed of **Ours** both before and after tuning the weights with respect to the best LLM-based agent’s reward function. The darker bars represent the chosen reward function for the final iteration. *Bottom:* Training curves of the final chosen reward function of the same seed. We report the mean final place, car collision times in seconds, and course out times in seconds across three seeds (\pm std)