



HAL
open science

Understanding reinforcement learned crowds

Ariel Kwiatkowski, Vicky Kalogeiton, Julien Pettré, Marie-Paule Cani

► **To cite this version:**

Ariel Kwiatkowski, Vicky Kalogeiton, Julien Pettré, Marie-Paule Cani. Understanding reinforcement learned crowds. 2022. hal-03780979

HAL Id: hal-03780979

<https://hal.science/hal-03780979v1>

Preprint submitted on 19 Sep 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Understanding reinforcement learned crowds

Ariel Kwiatkowski^{a,*}, Vicky Kalogeiton^a, Julien Pettré^b, Marie-Paule Cani^a

^aLIX, École Polytechnique/CNRS, Institut Polytechnique de Paris, Palaiseau, France

^bINRIA Rennes, France

ARTICLE INFO

Keywords: Reinforcement learning, Crowd simulation, Character animation, Multiagent reinforcement learning, Artificial intelligence

ABSTRACT

Simulating trajectories of virtual crowds is a commonly encountered task in Computer Graphics. Several recent works have applied Reinforcement Learning methods to animate virtual agents, however they often make different design choices when it comes to the fundamental simulation setup. Each of these choices comes with a reasonable justification for its use, so it is not obvious what is their real impact, and how they affect the results. In this work, we analyze some of these arbitrary choices in terms of their impact on the learning performance, as well as the quality of the resulting simulation measured in terms of the energy efficiency. We perform a theoretical analysis of the properties of the reward function design, and empirically evaluate the impact of using certain observation and action spaces on a variety of scenarios, with the reward function and energy usage as metrics. We show that directly using the neighboring agents' information as observation generally outperforms the more widely used raycasting. Similarly, using nonholonomic controls with egocentric observations tends to produce more efficient behaviors than holonomic controls with absolute observations. Each of these choices has a significant, and potentially nontrivial impact on the results, and so researchers should be mindful about choosing and reporting them in their work.

1. Introduction

Simulating virtual human crowds is a common task when creating lively, populated scenes for graphics applications. This includes a large variety of scenarios – ranging from small, artificially structured scenes used in research, to large scale simulations with thousands of virtual agents. They can be used either for real-life applications (e.g. testing evacuation scenarios in airports or sport stadiums) or targeting automatic content creation in films and games, featuring a diversity of background characters, each with their own goals and motivations.

While the approaches to simulating such crowds typically use hand-crafted or data-driven algorithms, in recent years there has been an increasing research interest in using Deep Reinforcement Learning (DRL) methods. These algorithms have a promise of a flexible, problem-agnostic training process that can autonomously produce diverse behaviors. However, they also bring many new challenges, which must be

explored independently of the previous knowledge of crowd simulation algorithms.

It is worth noting that just “crowd simulation” is somewhat of an underdefined problem. Although we have an intuition as to what a crowd is and how it behaves, it is not clear how to formalize it, which is exemplified by the variety of descriptions in prior work. At the same time, as Reinforcement Learning (RL) is designed to optimize a given scalar reward function, it excels when the objective is clearly stated. For this reason, we investigate the various ways to specify the crowd simulation problem.

There are three main components of an RL system – observations, actions, and rewards. For each of them, it is not necessarily obvious what is the appropriate level of abstraction, or even the implementation, to properly simulate human crowds. Take actions, for example – it is infeasible to perform a fully accurate biomechanical simulation of each muscle movement, so we must use a simplified model. Similarly for observations, performing a full rendering of the scene to emulate human vision would not only be expensive to perform, it would also be difficult to train.

Designing the reward function is arguably even more complex. While navigating in our daily life, we balance several, often mutually contradictory objectives, such as getting to

*Corresponding author.
e-mail: ariel.kwiatkowski@polytechnique.edu (A. Kwiatkowski)

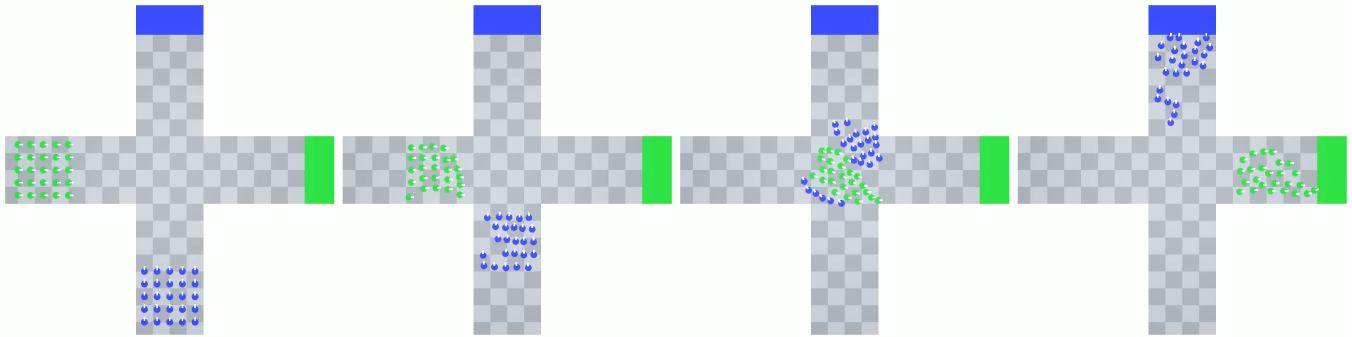


Fig. 1: Two groups of virtual agents crossing a corridor

the destination efficiently, moving at a comfortable speed or avoiding bumping into others. Even though some of them are simple to formalize as a single scalar reward, it is not obvious how they should be balanced. Is it acceptable to reach the destination two seconds earlier, but increase the risk of bumping into someone by 5%? Can we take a shorter path which leads through a group of people, or should we go around, reducing risk of collision, but increasing the distance? The answer to each of this question will vary from person to person, and many external factors. If going a bit slower would cause us to barely miss a train, we are likely to accept a higher risk of colliding than if we are just taking a walk around the park.

In this work, we intend to bring clarity to the intersection of crowd simulation and DRL, by exploring in detail the impact of these design choices on the generated virtual crowds. We perform a theoretical analysis of controlling the agents’ velocities, and an empirical investigation of various observation and action spaces. We evaluate them in terms of optimizing the reward function, but also consider the energy expenditure, and various quantitative properties of the movement. Our contributions are:

1. Empirical evaluation of raycasting versus direct agent perception in RL crowd simulation
2. Empirical comparison of various implementations of observations and dynamics in RL crowd simulation
3. Theoretical and empirical analysis of the properties of reward functions for efficient navigation with RL

2. Background & Related Work

Microscopic simulation of virtual crowds has garnered significant research interest in recent years. Various non-learning techniques have been introduced in the previous decade [1], which typically involve designing rule-based systems to update each agent’s velocity based on their context. There also exists a body of work using RL for controlling virtual characters, including crowd scenarios [2]. In this section we describe the elements of related work which are the most relevant to applying RL in crowd simulation.

Reinforcement Learning. RL is a study of sequential decision making, where one or more agents act in an environment, affecting its state, and receiving rewards. Modern RL

widely uses neural network-based algorithms [3], using them e.g. as a policy function, mapping observations to actions taken by the agent. This network is then optimized with a procedure based on the Policy Gradient Theorem, as introduced by the REINFORCE algorithm [4]. A more modern version that follows the same principle is the Proximal Policy Optimization (PPO) algorithm, introduced by Schulman et al. [5]. PPO is now the de facto standard on-policy algorithm used in many DRL applications due to its simplicity and efficiency.

Design Choices in RL. The importance of various, seemingly minor design choices in RL is a common phenomenon across many domains of the field. In physics-based animation, Reda et al. [6] explore how the performance of RL agents is affected by parameters like the initial state distribution, components of the state representation, or the control frequency. Engstrom et al. [7] perform a large-scale study of implementation details and code-level optimizations that affect the performance of TRPO and PPO in robotic control tasks. Similarly, Andrychowicz et al. [8] investigate how different design choices in on-policy algorithms affect the agent’s performance.

Crowd Simulation. Microscopic simulation of crowds is typically done by either using **force-based** methods [9] where the positions of nearby agents, obstacles, as well as the agent’s destination, all contribute to Social Forces that drive the acceleration of the agent at the next time step, or alternatively **velocity-based** methods such as Optimal Reciprocal Collision Avoidance (ORCA) [10]. The latter construct obstacles in the agent’s velocity space, which correspond to velocities that would result in a collision if the agent were to take them. This leads to effective solutions to collision avoidance, able to capture anticipation behaviours in crossing scenarios. Recently, **vision-based** and **data-driven** algorithms were explored as well, promising even more human-like results [1].

Crowd Simulation via DRL. There is a number of prior papers which train DRL agents on the task of crowd simulation. Long et al. [11] use a multiagent robotic navigation setup, which shares certain properties with crowd simulation. Lee et al. [12] train an RL agent on a variety of crowd scenarios, showing that a single trained model can be used to control multiple agents acting in a shared environment, on a diverse range of scenarios. Sun et al. [13] train groups of agents by making them follow specially-trained leader agents.

Xu et al. [14] combine DRL with an ORCA layer that ensures collision-free movement.

DRL is also used to generate more interesting, higher-quality trajectories. Xu and Karamouzas [15] use real-world human trajectory data to train a supervised model judging the human-likeness of a generated trajectory. Then, the output of that model is used as an additional component of the reward function, encouraging agents to act in a human-like manner. Hu et al. [16] use a parametric RL approach to generate heterogeneous behaviors with a single shared policy network. Each agent has its own preferred velocity value, and is trained to move according to it. Similarly, Panayiotou et al. [17] vary the weights in the reward functions of different agents in order to give them unique and configurable personalities.

While each of these prior works tackles the same problem of crowd simulation, many elements of their fundamental setup differ in potentially significant ways, which makes direct comparison infeasible. Therefore, in this work, we compare the basic design choices used in all these papers, in order to analyse their impact.

3. Environment Design Choices

We identify three design choices which can impact the properties of virtual crowds trained with a standard DRL algorithm – observations, actions, and the reward function. In this section, we set the problem in standard multiagent RL formalism, and describe the variants of observations and action spaces explored in this work.

3.1. Problem Formulation

We model the problem of crowd simulation as a Partially Observable Stochastic Game (POSG) [18]. A POSG is defined as a tuple $(\mathcal{I}, \mathcal{S}, \{\mathcal{A}_i\}, \{\Omega_i\}, \{O_i\}, T, \{R_i\}, \mu)$, where \mathcal{I} is the set of agents, \mathcal{S} is a set of states of the environment, \mathcal{A}^i is a set of actions for agent i ($\mathcal{A} = \times_{i \in \mathcal{I}} \mathcal{A}^i$ is the joint action set), Ω^i is the set of observations, $O^i: \mathcal{S} \rightarrow \Omega_i$ is the observation function, $O: \mathcal{S} \rightarrow \Omega$ is the observation function, $T: \mathcal{S} \times \mathcal{A} \rightarrow \Delta \mathcal{S}$ is the environment transition function, $R^i: \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the reward function, and $\mu \in \Delta \mathcal{S}$ is the initial state distribution.

In a POSG, all agents simultaneously make decisions based on their own private observations. Then, the environment is updated according to the joint action of all agents, and each agent receives its own reward that it tries to maximize. The reward is computed the same way for each agent, but based on its individual situation (i.e. no reward sharing). We additionally specify a time limit $T_{max} \in \mathbb{N}$ which is the maximum number of steps the environment is allowed to take before resetting.

3.2. Observation Space

In order to navigate through the environment, each agent must perceive its environment in some way. However, it is not obvious in what form agents should receive this information, or in fact, what the information should be.

The simplest human-inspired design is to give the agent information in its own frame of reference, and to have it perceive the environment through raycasting – a simple approximation of human vision. The intention is that if humans can effectively navigate using this type of information, then it should also suffice for virtual agents, which should then act more human-like. However, it is not necessarily the case that an anthropomorphic structure is indeed optimal for virtual agents, especially with it being only a rough approximation. More realistic rendering of the agent’s vision is an option, but it would result in very large observation sizes, subject to the curse of dimensionality. Thus, it is worthwhile to explore other possibilities.

An agent must receive information about its surroundings (other nearby agents and obstacles, i.e. **Environment Perception**), but also about its own internal state and knowledge (e.g. its current velocity, or its current destination, i.e. **Proprioception**). In both of these cases, it is also relevant what is the **reference frame** in which they are observed. For the Environment Perception, we consider two types of perception: Raycasting and Direct Agent Perception (AP) (Section 3.2.1). For the reference frame, we have three representations: Absolute, Relative, and Egocentric (Section 3.2.2).

3.2.1. Environment Perception

Raycasting refers to a method where several rays are cast from the center of the agent, in a plane parallel to the ground. Each of those rays then provides information on whether or not it collides with any object within a predefined distance, and what is the distance to the collision location.

Direct Agent Perception, similar to the method used by Xu and Karamouzas [15] is an alternative approach, where the agent directly receives the positions of other agents within a certain range, along with other relevant parameters. The reference frame in which the information is passed follows the chosen proprioception model (i.e. Absolute, Relative or Egocentric). This has the possible benefit of directly giving access to relevant information, but it introduces two important complications. Firstly, this method cannot canonically represent obstacles. While small, human-sized obstacles can be treated as stationary agents, large obstacles like walls need a different approach. Secondly, the number of neighboring agents is variable, and can grow very large in high density scenarios. The standard Multi-Layer Perceptron architecture cannot handle variable-sized inputs, so this variability has to be accounted for in some way. Furthermore, the order in which neighboring agents are observed is irrelevant, so a permutation-invariant architecture is necessary to decrease the effective size of the observation space.

In this work, we also test a multimodal approach which combines both Raycasting and Agent Perception, described in more detail in Section 5.2. With this hybrid method, the raycasting is only used to perceive static obstacles such as walls, ignoring other agents, whose positions are instead observed directly. We hypothesize that this might enable gaining the benefits of both methods – the agent has accurate knowledge of others, as well as a general idea of the surrounding layout,

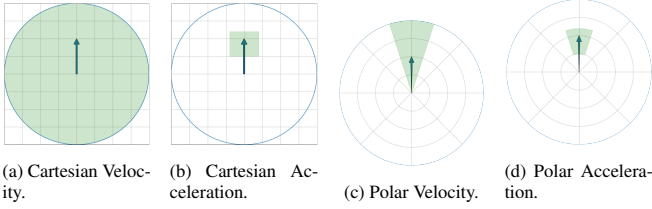


Fig. 2: A schematic representation of the available action spaces. In each case, we take a bird’s-eye view of an agent moving in the positive Y direction at an intermediate speed, represented by the blue arrow. The blue circle represents the space of all physically possible velocities (i.e. below the maximum speed). The green area represents the velocities that the agent is able to have in the following timestep under the specific action space.

sufficient for navigation.

3.2.2. Reference Frames

Absolute observations [13] use a bird’s-eye view on the global scene. The agent observes a vector consisting of its position \mathbf{p} and the position of its goal \mathbf{p}_g in the global coordinate frame, its current orientation ϕ (which is relevant for some choices in Action Spaces), and its current velocity \mathbf{v} . Using the Agent Perception approach, the agent observes the positions \mathbf{p}_i and velocities \mathbf{v}_i of nearby agents in the global frame.

Relative observations [15, 11, 14], like Absolute, use the global frame, however it is translated so that it is centered on the agent. It again receives its absolute position \mathbf{p} and velocity \mathbf{v} in order to retain the information about the large context, but the goal position is given relatively to the agent’s own position, as $\mathbf{p}_g - \mathbf{p}$. Similarly, using Agent Perception, other agents’ positions are given as $\mathbf{p}_i - \mathbf{p}$, and velocities are given in the absolute form \mathbf{v}_i .

Egocentric observations [16, 12] use the agent’s local frame according to its orientation. The agent observes its position \mathbf{p} and orientation angle ϕ in the global frame. We write R_ϕ to be the rotation matrix associated with the agent’s current orientation. The agent observes its goal position as $R_\phi(\mathbf{p}_g - \mathbf{p})$. Using Agent Perception, the positions of other agents are also represented as $R_\phi(\mathbf{p}_i - \mathbf{p})$, and their observed velocities are $R_\phi\mathbf{v}_i$.

While each of these has a theoretical justification (moving from absolute information about the scene, towards a more human-like first-person view), it is not immediately obvious which one is the best. On the one hand, absolute observations can give a high-level overall view of the scene, potentially aiding coordination. On the other hand, a relative or egocentric view allows agents to better reuse experiences between different positions and situations. If the agent is heading towards its goal, it is not that important whether it is to the left or to the right, looking from a bird’s eye view. Since the Egocentric observations lose this information, navigation might be expected to be learned more efficiently. Note that the choice of the reference frame also affects the structure of Agent Perception.

3.3. Action Space and Dynamics

Human motion is highly complex, and a biomechanically accurate simulation of human motion is a challenging research problem in of itself, so for the purposes of creating virtual crowds, we use a simplified model. The simplest choice is holonomic locomotion [19], where at each step, the agent can choose its velocity constrained only by its magnitude. However, this approach does not correspond well to the motion constraints of real humans. Arechavaleta et al. [20] propose a nonholonomic model, in which an agent can move in the direction of its current orientation, and incrementally change its orientation for the next timestep.

Allowing the agent to freely choose its velocity at every timestep gives it much more flexibility in choosing its behavior. However, from the perspective of Newtonian mechanics, it is more physically justified for the agent to directly choose its acceleration. This would mean that the velocity change at each timestep is incremental.

Similarly, there is a choice in how the actions should be represented. We can take the bird’s eye view, where the agents choose their actions according to an absolute reference frame, moving up or down, left or right. Alternatively, we can take a more individual perspective, where the agents operate in a polar frame, choosing their linear movement and the direction of that movement.

For this reason, we consider four different dynamics models of the environment: Cartesian Velocity, Cartesian Acceleration, Polar Velocity, and Polar Acceleration.

Cartesian control (used by [15, 16]) implies that the agent separately chooses the x and y components of its motion - either of its velocity or acceleration.

Polar control (used by [12, 11, 13, 14]), implies that the agent separately updates its orientation angle, and its linear speed. The linear speed is again updated either by choosing an arbitrary value below a certain magnitude (velocity controls), or by incrementally updating it according to the acceleration chosen by the agent.

Acceleration controls (used by [16, 15]), are modeled in our implementation using a linear damping model. This means that given an acceleration a_t chosen by the agent as its action, the updated velocity will be

$$v_t = v_{t-1} + (a_t - \lambda v_{t-1})\Delta t$$

where λ is chosen so that we obtain the same maximum speed as in velocity controls, which is equal to $2\frac{m}{s}$.

With velocity controls (used by [12, 11, 13, 14]), the agent can choose an arbitrary speed of a magnitude lower than $2\frac{m}{s}$.

In Cartesian controls, the agent’s orientation is defined to be parallel to its current velocity. In Polar controls, the orientation is directly controlled, and the velocity is parallel to the orientation.

4. Reward Function Design

Designing the reward function is arguably the most impactful, and the most difficult part of creating an RL-driven crowd

simulation. The two components which are present in all existing work are a positive term correlated with reaching the goal, and a negative term correlated with collisions. Beyond that, various elements may be included to promote certain behavior characteristics, or to improve training performance through reward shaping. In this work, we consider the following reward components based on prior work:

1. Reward for reaching the goal $R_g = c_g$ (once) [11, 13, 14]
2. Reward for approaching the goal $R_p = c_p(d_t - d_{t-1})$ (every timestep) [15, 16, 12, 11, 13, 14]
3. Reward for maintaining a comfortable speed $R_v = -c_v|v - v_0|^{c_e}$ (every timestep) [15]
4. Penalty for collisions $R_c = -c_c$ (every collision, every timestep) [15, 16, 12, 11, 13, 14]
5. Reward for urgency $R_t = -c_t$ (every timestep) [13]

where $c_g, c_p, c_v, c_e, c_c, c_t$ are arbitrary (typically positive) coefficients. Their roles are as follows: R_g is the main (sparse) reward representing the agent’s destination. R_p provides a dense reward for the navigation objective to enable faster training. R_v incentivizes moving at a comfortable speed v_0 . We propose raising the absolute difference of speeds to an arbitrary power c_e in order to further shape the behavior. R_c makes agents avoid colliding with obstacles and with one another. R_t is a commonly used reward component in goal-based RL environments, as it incentivizes the agents to reach their goal sooner rather than later.

4.1. Energy and Metrics

When considering different reward functions, it is important to have metrics that are independent of the specific reward formulation in order to have a meaningful comparison. For this reason, we consider two types of metrics.

Firstly, for each component of the reward function, we compute an unnormalized value measuring its performance, which can then be compared between different training runs. For example, we consider the total number of collision, regardless of the size of the collision penalty in the reward function.

Secondly, we compute the mean energy expenditure of all agents in the scene. We use the following formula [21, 22]:

$$E_t = e_s + e_w v^2$$

where E_t is energy spent per second, per kilogram body mass (J/kg·s), e_s and e_w are constants, and v is the current velocity (m/s). We use the values of $e_s = 2.23$ and $e_w = 1.26$ of a typical human. This induces an optimal walking speed of $v^* = \sqrt{\frac{e_s}{e_w}} = 1.33$ m/s which we use as the preferred walking speed. Work by Bruneau et al. [23] suggests that when navigating to avoid collisions, people tend to choose a path that minimizes the energy usage, so we consider it to be a valuable metric describing the efficiency of the generated trajectory.

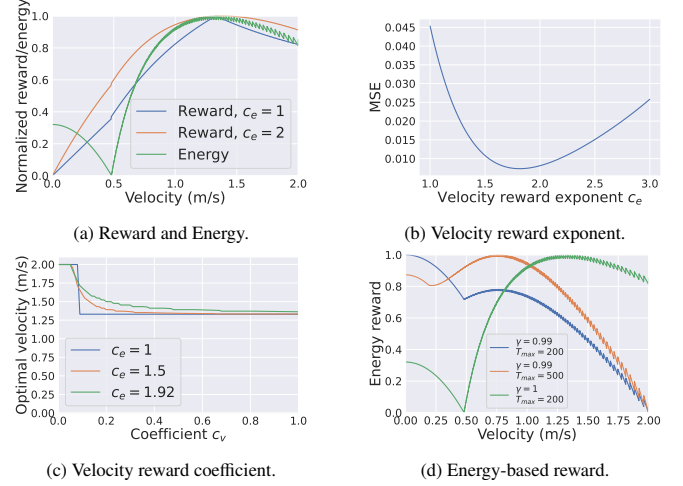


Fig. 3: (a) Rewards and energy for an agent moving at a constant speed, in the simplified model described in Section 4.2. All curves are normalized to be in the $[0, 1]$ range in order to enable direct comparison. We consider energy values with the opposite sign, because energy is supposed to be minimized, while the reward is maximized. (b) MSE between the reward and the energy, as a function of the velocity reward exponent. (c) Optimal velocity as a function of the velocity reward coefficient c_v , varied by the exponent c_e . (d) Discounted negative energy expenditure as a function of velocity.

4.2. Reward and Preferred Velocity

As we see from the energy optimization mechanism, humans tend to move at a certain speed which is below their maximum possible speed. This must be reflected in the reward function that the RL agent optimizes. However, it goes against the typical RL incentives to obtain rewards sooner due to the discounted utility model [3].

The interaction between the values of $c_g, c_p, c_v, c_e, c_t, \gamma$, and T_{max} , make the effective preferred velocity nontrivial to predict, and as a consequence – to design. Consider a simplified environment model in which the agent must travel a distance d towards its goal, with no other agents or obstacles. In this case, the only decision it makes is the velocity throughout its motion, under the assumption that it will travel to its goal in a straight line at a constant velocity. In this model, we can express the total obtained reward as:

$$R(v) = \gamma^T c_g + \sum_{i=0}^T \gamma^i (c_p v \Delta t - c_v |v - v_0|^{c_e} - c_t) \quad (1)$$

where $T = \min(T_{max}, \lceil \frac{d}{v \Delta t} \rceil)$. Due to the discounted sum whose bounds are dependent on the velocity itself, it is difficult to investigate this expression analytically. Nevertheless, we can gain some insight through numerically analyzing this model. Note, however, that this model does not capture the full complexity of RL optimization, and only serves to build general intuition.

Consider the following set of parameters values – based on prior work and then manually adjusted to produce reasonable behaviors – as a starting point of our analysis: $c_g = 10$, $c_p = 1$, $c_v = 0.75$, $c_e = 1$, $c_t = 0.005 v_0 = 1.33$, $\gamma = 0.99$, $T_{max} = 200$, $d = 8$, $\Delta t = \frac{1}{12}$.

Let us investigate the full reward as a function of velocity,

alongside the negative energy expenditure as defined in Section 4.1. This relation is shown in Figure 3a. We can see that while there is a correlation between reward and energy, there are two main discrepancies. Firstly, the energy has a local optimum at $v = 0$, which is caused by the velocities that cause agents to not reach their destination within the allotted time. Secondly, due to the absolute difference term in Equation 1, there is a sharp decrease in the reward (blue curve) which does not occur in the energy.

To improve this, we evaluate the impact of the exponent c_e from the reward function. In Figure 3a, we also show $c_e = 2$ (orange). There, the curve is smoother and closer to the corresponding energy values when the energy is near the optimum. To quantify this, we vary the parameter c_e and compute the mean square error between the two normalized curves in the range $1 \frac{m}{s} < v < 2 \frac{m}{s}$, as we consider lower velocities to be less relevant due to their low efficiency. As we show in Figure 3b, the optimal value under this simple model is $c_e = 1.92$. We further validate this in Section 6, where we use different values of c_e for training actual RL agents.

It is worth mentioning that with $\gamma = 0.99$, using $c_e = 2$ increases the effective optimal velocity to $v^* = 1.39 \frac{m}{s}$. This can be further adjusted using other parameters. While modifying γ does not affect the optimal velocity when $c_e = 1$, any higher values of c_e make it so that decreasing γ , increases v^* .

The other reward coefficients also have an impact on the optimal velocity. When adjusting c_v with $c_e = 1$, there is a threshold around $c_v = 0.09$ below which the optimal velocity is the maximum value of $2 \frac{m}{s}$. Above that threshold, the optimal velocity is the preferred value of $v_0 = 1.33 \frac{m}{s}$. However, with $c_e > 1$, the transition between these two realms becomes more gradual, as we show in Figure 3c.

4.3. Energy as reward

Finally, it is worth considering using energy directly as a reward function for training RL agents. At a glance, it seems like it would incentivize efficient motion at the optimal velocity. However, there are two apparent problems which arise in this paradigm. Firstly, as we already see in Figure 3a, even in our simplified model there is an attractive local optimum at $v = 0$. This is likely to be even more impactful in practical scenarios, because moving at the right speed would lead to a very low reward if the direction of the movement is wrong. The reason for this is the time limit present in the environment – moving at a non-zero speed only pays off in terms of energy if the agent eventually reaches the goal. Otherwise, it will expend energy until the end of the episode, and by reducing its velocity, it can reduce the energy expenditure. Secondly, the commonly used method of reward discounting has a significant impact on the optimal policy. Using a discount factor of $\gamma = 0.99$ leads to a situation where the global optimum of the discounted energy-based reward function is standing still with $v = 0$. Potential solutions to these problems include using a discount factor $\gamma = 1$ or a nonexponential discounting mechanism, increasing the time limit, and using a curriculum-based approach.

Conclusion. Even in the absence of collision avoidance and other, more complex tasks, one must pay attention to the

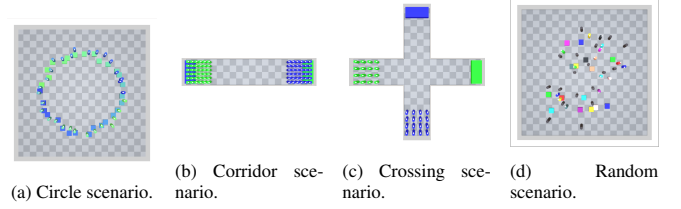


Fig. 4: Agent’s initial positions and goals in four scenarios: (a) Circle with 30 agents. (b) Corridor with 72 agents. (c) Crossing with 32 agents. (d) Random with 15 agents.

parameters defining the reward function. Notably, using an exponent in the velocity reward term causes other parameters to nontrivially affect the effective optimal velocity. For this reason, when designing the reward function, it is worthwhile to validate its parameters using a simpler model to ensure it has desired properties, whether that is closeness to the energy expenditure, or a specific value of the preferred velocity.

5. Experimental setup

In order to evaluate the impact and quality of the various design choices, we apply them on four commonly used crowd scenarios, in order to provide a wide range of interactions between agents: Circle, Corridor, Crossing, Random (see Figure 4). In the Circle scenario, agents start on the perimeter of a circle, with a random noise applied independently in both Cartesian directions. Their goals are placed on the antipodal points of the circle, with an independent noise of the same magnitude applied. In Corridor, agents start at two ends of a straight corridor whose width is 4 meters and length is 20 meters. They start either in a regular grid or in a random formation, and their goal is to reach the opposite side of the corridor. In Crossing, the agents start at the ends of two corridors intersecting at a right angle, with the same size as in Corridor. Similarly, they spawn either in a regular grid or a random formation, and must reach the other end of their respective corridors. In Random, the agents’ starting positions and goals are generated according to a uniform distribution with a given maximum size. In each of these scenarios, the area available to the agents is a square of 20x20 meters. In both Circle and Random, there are optional small obstacles placed randomly in the scene, represented as immovable agents.

Implementation. The code¹ used in this work is available online. We use identical agents represented as circles of radius 0.2 m. Their collisions are treated as rigid body collisions, processed by the PhysX engine default in Unity 2021.3, used via the Unity ML-Agents framework [24]. We use a decision timestep $\Delta t = \frac{1}{12}$ s, similar to the values used in prior work. To obtain a more accurate simulation, the physics of the scene are updated 10 times after each decision, for an effective simulation timestep of $\frac{1}{120}$ s. The agent’s action is repeated during

¹The code for the environment is available at <https://github.com/RedTachyon/CrowdAI>, and the training code is available at <https://github.com/RedTachyon/coltra-rl>.

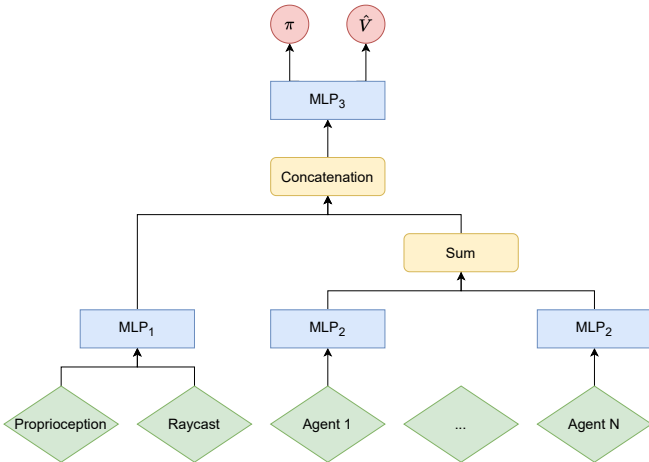


Fig. 5: The neural architecture used as the policy. Green blocks represent inputs, blue blocks represent feed-forward neural networks, yellow blocks represent vector operations, red blocks represent outputs. Depending on the observation model used, certain elements of the architecture are disabled.

each of these updates. With agent perception observations, each agent can see the 10 nearest agents. While these implementation details (i.e. agent size, collision handling, physics engine, timestep) can also affect the resulting simulations and the training performance, we do not explore their impact in this work, because we expect it to be lower compared to the other choices listed in this paper. It is nevertheless important to be aware of these choices for reproducibility purposes. A single training takes between 1 and 5 hours on GPU, depending on the number of agents and the difficulty of the scene, while running 8 independent training runs in parallel.

5.1. Policy Optimization

In this work, we train RL agents using PPO with Generalized Advantage Estimation (GAE) [25] to estimate the advantages. The agents are trained in an independent paradigm with parameter sharing, that is they share the same policy network, but each agent takes its own action based on its private observations. The neural network outputs the mean of a Gaussian distribution, and the standard deviation is kept as a trainable parameter of the network.

5.2. Network Architecture

In order to appropriately process the Agent Perception observations, we use a neural architecture depicted in Figure 5. It is inspired by prior work such as Deep Sets [26] and Mean Embedding [27], and extends the architecture used by Xu and Karamouzas [15].

The main desirable property of our architecture is permutation invariance – given multiple identical nearby agents, it should not matter in what order their representations are input into the network, as this order is completely arbitrary. Without this property, the agent would need to learn this invariance itself, which quickly becomes expensive as the number of observed agents grows. Furthermore, the architecture should be able to accept a variable number of observed agents, as this

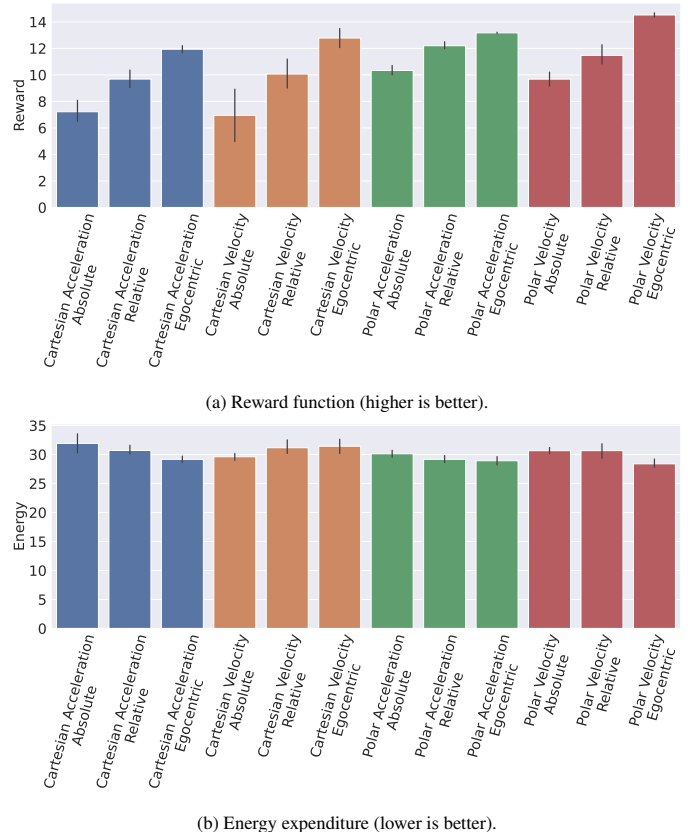


Fig. 6: Comparison of training results after a hyperparameter search in the Circle 12 scenario. (a) Mean episodic reward (b) Mean energy expenditure. Black bars represent the standard error of the mean.

quantity will vary throughout the episode. For this reason, we use the following model as an embedding of nearby agents:

$$\phi \left(\sum_i \psi(x_i) \right)$$

where ϕ and ψ are regular MLP neural networks, and x_i is the observed information about an agent i . The summation is performed over all agents visible to the agent observing the scene. Because of the summation operator, this architecture fulfills both previously stated desiderata, as the ordering information is lost, and any number of agents can be processed into a fixed-size embedding. This embedding is then concatenated with the main stream of the neural network, which processes the proprioceptive observations, as well as optionally the raycasting.

6. Experiments

In this section we describe the specific experiments we performed, along with their results and interpretation. Experiments are primarily evaluated in terms of their obtained rewards and energy usage, but also other behavior characteristics when necessary.

6.1. Dynamics and Observations performance

In order to robustly evaluate the performance of various dynamics and observation models, we run a hyperparameter

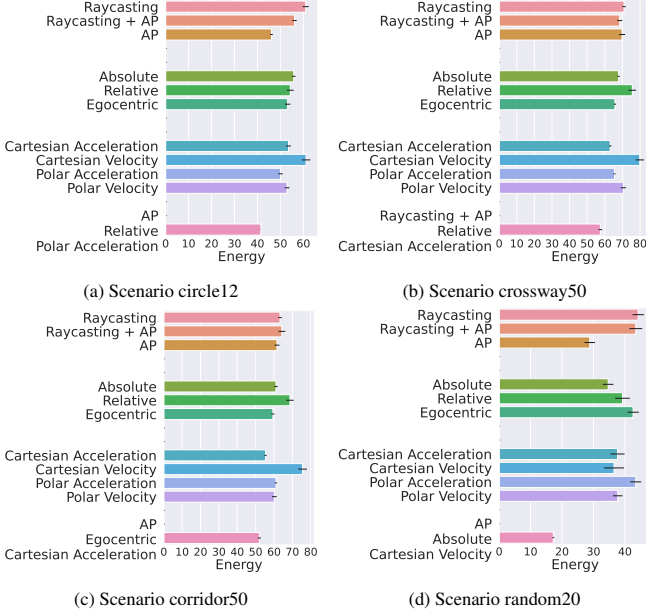


Fig. 7: Comparison of various design choices in a given environment. The last bar corresponds to the best-performing agent across all design choices. All values are averaged across 8 independent training runs with different random seeds and otherwise identical parameters. Lower is better. AP stands for Agent Perception as defined in Section 3.2.1.

search with each action space, with each observation reference frame, on a Circle scenario with 12 agents and a radius of 4 meters. Then, we use the best-performing hyperparameters of each model for further experiments used in this section. We sample 150 sets of hyperparameters for each (dynamics, observations) pair for a total of 1800 training runs. We use the default sampler available in Optuna [28]. These runs use a fixed reward function, with parameters $c_g = 10$, $c_p = 1$, $c_v = 0.75$, $c_e = 1$, $c_c = 0.05$, $c_t = 0.005$.

We show the results of the hyperparameter search in Figure 6. We consider the 5 best-performing hyperparameter sets for each model, and report their aggregate performance. There is a clear trend where agents with Egocentric observations perform better than Absolute and Relative versions. Similarly, Polar controls perform better than Cartesian controls, with Polar Velocity controls with Egocentric observations performing the best out of all investigated variants. Interestingly, in the case of Cartesian Velocity controls, while the reward follows the same trend, the energy usage is in fact the lowest with Absolute observations. This highlights the discrepancy between the reward function and the energy metric, showcasing the need for careful evaluation of emergent behaviors.

6.2. All Scenarios

Using the hyperparameters obtained in the experiment described in Section 6.1, we evaluate the various design choices on several environments. Specifically, we train agents in a Circle with 12 agents and a radius of 6 meters (circle12), a Corridor with 50 agents (corridor50), a Crossing with 50 agents (crossway50), and a Random scenario with 20 agents and random obstacles (random20).

Due to the large number of possible combinations (3 architectures \times 4 dynamics \times 3 observations \times 4 scenarios), we consider the performance of each individual choice, averaging the remaining ones, in each environment separately. The results are in Figure 7.

We run each training for 1000 PPO iterations to ensure sufficient time for convergence. Because the neural networks are relatively small (2-6 layers of 32-128 units, depending on the architecture and the model, found via hyperparameter optimization), the main bottleneck is the CPU power required to run many copies of the environment as opposed to the GPU memory.

While the exact results vary between scenarios, there are some regularities. Most notably, Agent Perception consistently outperforms raycasting-based approaches with otherwise well-performing settings. In the case of the crossway50 scenario, it is beneficial to include raycasting information around the surrounding walls, however overall it does not seem to provide a large benefit as compared to the pure AP approach. This is likely due to the static nature of the evaluated environments. Given the current position in the global frame, the agent can determine its proximity from obstacles. Unsurprisingly, including raycasting to perceive walls, deteriorates the performance in scenarios without a significant presence of walls, as this information effectively becomes an additional source of noise.

In most cases, it is best to use Egocentric or Relative observations, with some Relative runs sometimes performing the best. A notable exception is the random20 scenario, where a combination of Absolute observation with Cartesian Velocity dynamics outperforms any other option. The difference is the fact that in all other scenarios, the agents need to predominantly move in a general “forward” direction, whereas in the random scenario, the goal can be in an arbitrary position relative to the agent’s orientation.

An interesting observation is the common “failure mode” of Raycasting models, particularly in Circle scenarios. They generally perform worse than Agent Perception models, but their qualitative behavior may be more desirable due to its asymmetry. Because they do not manage to reach the perfectly symmetrical trajectories achieved by Agent Perception, there is a higher variability in the individual trajectories, making them look more realistic. This indicates that simply training a strong RL algorithm on any objective which does not explicitly reward human-likeness is likely to lead to overly perfect, unrealistic trajectories.

Conclusion. We recommend using AP as opposed to the more commonly used raycasting for providing agents with the information about their surrounding. In scenarios where walls are a prominent feature, it may be beneficial to add raycasting which only perceives the distance to walls, and ignores other agents. In scenarios where agents may need to make sharp turns to reach their destination, Cartesian Velocity controls with Absolute observations are favorable. Otherwise, nonholonomic controls combined with egocentric or relative observations typically perform better.

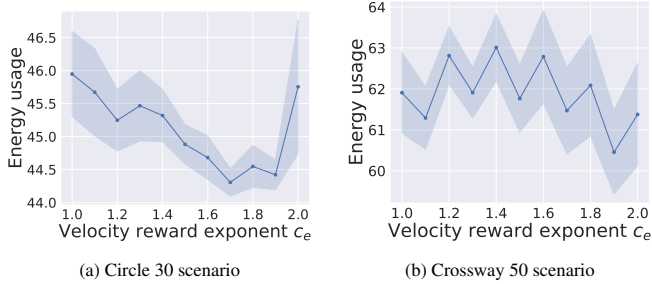


Fig. 8: Comparison of energy usage in agents trained with a different exponent in the velocity term of the reward function. Lower is better.

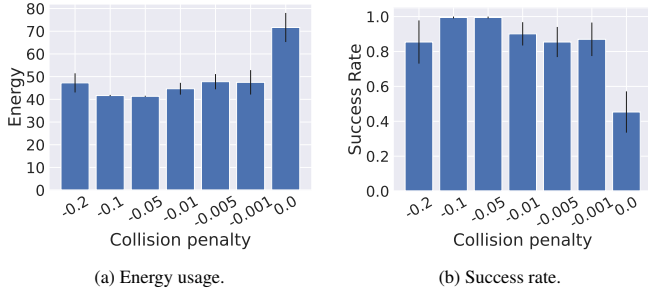


Fig. 9: Comparison of energy usage and success rate in agents trained in Circle 12 scenario, with a varied collision penalty in the reward function.

6.3. Velocity Reward Exponent

In Section 4.2, we show that using an exponent in the velocity reward term makes it match more closely to the energy consumption. To validate this, we train Egocentric Polar Velocity agents on a Circle 30 scenario, and Egocentric Polar Acceleration agents on a Crossway 50 scenario. In both cases, an exponent $c_e > 1$ can lead to a higher efficiency in the trained agents, as compared to the simple $c_e = 1$ (see Figure 8). The exact optimal value of c_e depends on the scenario and must be determined on a case-by-case basis.

6.4. Importance of collision penalty

In a pursuit of simplicity in the design of the reward function, one might be tempted to eliminate the collision penalty altogether. After all, if collisions result in unfavorable physical results (i.e. unexpected change of velocity), sufficiently intelligent agents should learn to avoid them by themselves, at least to the extent that is necessary for effective navigation.

In Figure 9 we show how the collision penalty affects the energy usage, and how often the agents reach their destination. There is an optimum around -0.05 , where the agents reliably reach their goals. Using a penalty that is too high or too low leads to a deterioration of the agents’ performance in terms of the navigation task.

6.5. Common Failure Modes

Due to the stochasticity inherent to RL training, the trained agents often exhibit various types of suboptimal behaviors. This can be identified via tracking the performance (in terms of the reward and energy), but also by observing emergent behaviors that the agents learn to execute. Here, we describe

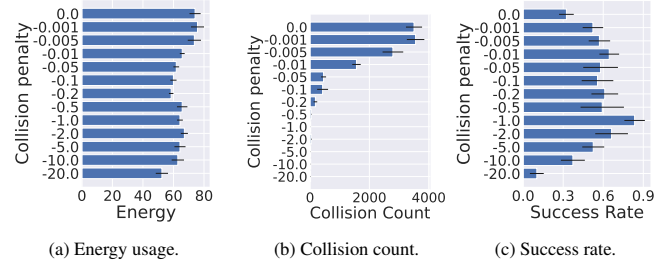


Fig. 10: Comparison of energy usage, collision count and success rate in agents trained in Crossway 50 scenario, with a varied collision penalty in the reward function.

some of the common ways in which the RL-trained crowds are suboptimal.

Instability. When training an RL agent using PPO, the trained policy is stochastic. This is required both for training, to ensure that the agent takes sufficiently diverse actions; but it is also at the core of the resulting policy. In order to deploy or evaluate the agent, we must choose a method of sampling actions from the output of the policy. The two natural options are taking the mean for the “optimal” action, or simply sampling from the distribution. The former does not fully correspond to the optimization objective, and in a crowd scenario, it can get stuck on an obstacle or another agent. The latter necessarily causes the policy to act randomly, which leads to potential unpredictable mistakes due to small erratic movements.

Value alignment failure. As we describe in Section 4, the reward function used in this and prior work is a weighted sum of several components with largely arbitrary coefficients. This implies that there is no guarantee it will correspond well to the actual objective we intend the agents to achieve. Furthermore, while the energy usage can be a useful, non-arbitrary metric, it is also sensitive to some details of the practical RL setup, as we show next.

Consider the experiment described in Section 6.4, investigating the importance of the collision penalty. We perform the same analysis on the Crossway 50 scenario with Polar Velocity dynamics and Egocentric observations, extending the range of evaluated collision penalties to -20 . We present these results in Figure 10.

When considering the energy usage and collisions, the result seems to follow the intuition with a collision penalty up to -1 . When the collision penalty reaches -20 , we expect the performance to deteriorate, since the penalty is so steep that the agents will never put themselves at any risk of collision. We confirm this by investigating the success rate, which decreases with large collision penalties. However, at the same time, the energy usage also decreases – in fact, using a very large collision penalty leads to the lowest energy usage among the evaluated options. So while energy has its value as a metric, it clearly breaks down in extreme cases, where most agents remain stationary.

This phenomenon is caused by a relatively short time limit. In general, energy minimization induces a simple optimal ve-

locity (see Section 4.1), in RL we must set a time limit after which the episode is terminated. The energy penalty for not finishing the episode is relatively low when compared to the energy cost of actually navigating to the destination, and an undesired emergent behavior may turn out to be optimal according to the metrics.

Depending on the design choices made, this same problem can manifest itself differently. When training on the same setup, but using different design choices, a commonly occurring emergent behavior is that one group of agents efficiently navigates to their destination, while the other one remains stationary. Similarly to the previous situation, the energy usage metric is very low in that situation, but the success rate plateaus at 50%.

7. Discussion

In this work, we present an analysis of various design choices made by designers of RL-trained crowd simulation systems. We show that many of these choices, typically ignored by researchers, can in fact significantly impact the resulting simulation, in particular when evaluated in terms of the energy efficiency. Here we summarize the main findings of this paper.

We show that the commonly used raycasting underperforms when compared to a method we call Agent Perception where the information about neighboring agents is directly available. This is likely a consequence of a much simpler and accessible representation of that information. Even when raycasting uses frame stacking which enables movement perception, the reasoning needed to infer the positions and velocities of nearby agents is rather complex.

Designing the right reward function is also important for obtaining desired properties of the motion. Navigation, speed control, and collision avoidance rewards, all have to be balanced in just the right way that each of them contributes to the agent’s decisions. Crucially, quantitative comparisons of crowds are non-trivial, because even ignoring the question of believability, the arbitrary reward function and energy usage are flawed in certain scenarios.

Qualitative properties of the obtained motion in the Circle scenario indicate that a naive approach of reward maximization for any reward that does not explicitly incentivize human-like behavior is likely to create trajectories that look artificial, or even too carefully choreographed to pass for natural human behavior. While decreasing the model’s capacity and performance might lead to more believable behavior in the short term, we believe a more deliberate approach is necessary to truly approach human-like behavior.

In summary, our main findings are as follows:

1. Direct agent perception outperforms simple raycasting
2. Egocentric controls tend to outperform absolute ones
3. The reward design is important and nontrivial
4. Many failure modes may still occur in RL trained crowds
5. Simple reward is not sufficient for human-like behavior

7.1. Limitations and Future Work

All experiments in this work are performed on relatively small, static scenarios with a single destination. The described design choices mostly affect local navigation, and more complex scenarios can be expressed as a sequence of partial objectives or checkpoints. That being said, naively implementing this would likely cause issues near the transition points where agents switch their destination. Therefore, a more complex training scenario would be beneficial so that the agent is exposed to these situations.

Furthermore, each agent is only trained on a single scenario. Prior work suggests that using various scenarios in the training process enables generalization, which was considered out of scope of this work. We also limit our analysis to the efficiency of the resulting trajectories, ignoring realism or believability.

Similarly to prior work, we train agents using an arbitrarily-designed reward function. While using the energy usage as a reward has certain problems (see Section 4.2), it might be a viable option by using a curriculum-based approach where the reward function changes as the training progresses; and a different discounting mechanism that improves the global optimization properties. Furthermore, by using recent work on evaluating the realism of generated trajectories [29], a promising direction is using a realism metric as a reward. This would allow going beyond efficiency, and creating crowds which behave in a believable way.

It is also possible to improve the dynamics available to the agent. In this work, we use relatively simple, 2-DoF models, but the RL paradigm makes it viable to implement arbitrary nonholonomic constraints like sidesteps or walking backwards, without having to change the learning logic. Thus, a promising option is introducing a more complex, human-like range of motion actions available to the agent, with the goal of improving the believability of motion.

7.2. Conclusions

Crowd simulation with RL is a complex problem, and despite recent advances, many challenges remain. Observations, actions, underlying physics, and especially the reward function, all have a significant impact on the results, and a lack of attention to these design choices makes it impossible to compare various approaches. We bring these issues to attention, and introduce a basic methodology for comparison between various approaches by comparing the energy expenditure under a specified time limit. In the absence of a standard benchmark, we call researchers to be more explicit and precise about their choices, and encourage them to explore different options in their work, to ensure robustness of their approach.

Acknowledgments

This work has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 860768

(CLIFE project). This work was performed using HPC resources from GENCI–IDRIS (Grant 2022-AD011013684). This work was also supported by the Hi!Paris institute of IP Paris, through the CREATIVE AI fellowship.

References

- [1] Toll, W, Pettré, J. Algorithms for Microscopic Crowd Simulation: Advancements in the 2010s. *Computer Graphics Forum* 2021;40(2):731–754. doi:10.1111/cgf.142664.
- [2] Kwiatkowski, A, Alvarado, E, Kalogeiton, V, Liu, CK, Pettré, J, van de Panne, M, et al. A Survey on Reinforcement Learning Methods in Character Animation. *Computer Graphics Forum* 2022;41(2):613–639. doi:10.1111/cgf.14504.
- [3] Sutton, RS, Barto, AG. *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: A Bradford Book; 2018. ISBN 978-0-262-03924-6.
- [4] Sutton, RS, McAllester, D, Singh, S, Mansour, Y. Policy gradient methods for reinforcement learning with function approximation. In: *Proceedings of the 12th International Conference on Neural Information Processing Systems. NIPS'99*; Cambridge, MA, USA: MIT Press; 1999, p. 1057–1063.
- [5] Schulman, J, Wolski, F, Dhariwal, P, Radford, A, Klimov, O. Proximal Policy Optimization Algorithms. arXiv:170706347 [cs] 2017;.
- [6] Reda, D, Tao, T, van de Panne, M. Learning to Locomote: Understanding How Environment Design Matters for Deep Reinforcement Learning. In: *Motion, Interaction and Games. Virtual Event SC USA*: ACM. ISBN 978-1-4503-8171-0; 2020, p. 1–10. doi:10.1145/3424636.3426907.
- [7] Engstrom, L, Ilyas, A, Santurkar, S, Tsipras, D, Janoos, F, Rudolph, L, et al. Implementation Matters in Deep Policy Gradients: A Case Study on PPO and TRPO. arXiv:200512729 [cs, stat] 2020;ArXiv: 2005.12729.
- [8] Andrychowicz, M, Raichuk, A, Stańczyk, P, Orsini, M, Girgin, S, Marinier, R, et al. What Matters In On-Policy Reinforcement Learning? A Large-Scale Empirical Study. arXiv:200605990 [cs, stat] 2020;ArXiv: 2006.05990.
- [9] Helbing, D, Molnár, P. Social force model for pedestrian dynamics. *Physical Review E* 1995;51(5):4282–4286. doi:10.1103/PhysRevE.51.4282.
- [10] van den Berg, J, Guy, SJ, Lin, M, Manocha, D. Reciprocal n-Body Collision Avoidance. In: Siciliano, B, Khatib, O, Groen, F, Pradalier, C, Siegwart, R, Hirzinger, G, editors. *Robotics Research*; vol. 70. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-642-19456-6 978-3-642-19457-3; 2011, p. 3–19. doi:10.1007/978-3-642-19457-3_1; series Title: Springer Tracts in Advanced Robotics.
- [11] Long, P, Fan, T, Liao, X, Liu, W, Zhang, H, Pan, J. Towards Optimally Decentralized Multi-Robot Collision Avoidance via Deep Reinforcement Learning. arXiv:170910082 [cs] 2018;ArXiv: 1709.10082.
- [12] Lee, J, Won, J, Lee, J. Crowd simulation by deep reinforcement learning. In: *Proceedings of the 11th Annual International Conference on Motion, Interaction, and Games. Limassol Cyprus*: ACM. ISBN 978-1-4503-6015-9; 2018, p. 1–7. doi:10.1145/3274247.3274510.
- [13] Sun, L, Zhai, J, Qin, W. Crowd Navigation in an Unknown and Dynamic Environment Based on Deep Reinforcement Learning. *IEEE Access* 2019;7:109544–109554. doi:10.1109/ACCESS.2019.2933492; conference Name: IEEE Access.
- [14] Xu, D, Huang, X, Li, Z, Li, X. Local motion simulation using deep reinforcement learning. *Transactions in GIS* 2020;24(3):756–779. doi:10.1111/tgis.12620.
- [15] Xu, P, Karamouzas, I. Human-inspired multi-agent navigation using knowledge distillation. In: *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2021, p. 8105–8112. doi:10.1109/IROS51168.2021.9636463.
- [16] Hu, K, Haworth, MB, Berseth, G, Pavlovic, V, Faloutsos, P, Kapadia, M. Heterogeneous Crowd Simulation using Parametric Reinforcement Learning. *IEEE Transactions on Visualization and Computer Graphics* 2022;:1–1doi:10.1109/TVCG.2021.3139031.
- [17] Panayiotou, A, Kyriakou, T, Lemonari, M, Chrysanthou, Y, Charalambous, P. CCP: Configurable Crowd Profiles. In: *Special Interest Group on Computer Graphics and Interactive Techniques Conference Proceedings. Vancouver BC Canada*: ACM. ISBN 978-1-4503-9337-9; 2022, p. 1–10. doi:10.1145/3528233.3530712.
- [18] Hansen, EA, Bernstein, DS, Zilberstein, S. Dynamic programming for partially observable stochastic games. In: *Proceedings of the 19th national conference on Artificial intelligence. AAAI'04*; San Jose, California: AAAI Press. ISBN 978-0-262-51183-4; 2004, p. 709–715.
- [19] Hughes, R, Ondřej, J, Dingliana, J. Holonomic collision avoidance for virtual crowds. In: *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation. SCA '14*; Goslar, DEU: Eurographics Association; 2015, p. 103–111.
- [20] Arechavaleta, G, Laumond, JP, Hicheur, H, Berthoz, A. On the nonholonomic nature of human locomotion. *Autonomous Robots* 2008;25(1-2):25–35. doi:10.1007/s10514-007-9075-2.
- [21] Whittle, MW. *Gait analysis: an introduction*. 4th ed., reprinted ed.; Edinburgh: Butterworth-Heinemann, Elsevier; 2008. ISBN 978-0-7506-8883-3.
- [22] Guy, SJ, Chhugani, J, Curtis, S, Dubey, P, Lin, M, Manocha, D. PLEdrians: A Least-Effort Approach to Crowd Simulation. *Eurographics/ ACM SIGGRAPH Symposium on Computer Animation 2010*;:10 pagesdoi:10.2312/SCA/SCA10/119-128; artwork Size: 10 pages ISBN: 9783905674279 Publisher: The Eurographics Association.
- [23] Bruneau, J, Olivier, AH, Pettré, J. Going Through, Going Around: A Study on Individual Avoidance of Groups. *IEEE Transactions on Visualization and Computer Graphics* 2015;21(4):9. doi:10.1109/TVCG.2015.2391862.
- [24] Juliani, A, Berges, VP, Teng, E, Cohen, A, Harper, J, Elion, C, et al. Unity: A General Platform for Intelligent Agents. arXiv:180902627 [cs, stat] 2020;ArXiv: 1809.02627.
- [25] Schulman, J, Moritz, P, Levine, S, Jordan, M, Abbeel, P. High-Dimensional Continuous Control Using Generalized Advantage Estimation. arXiv:150602438 [cs] 2018;.
- [26] Zaheer, M, Kottur, S, Ravanbakhsh, S, Póczos, B, Salakhutdinov, RR, Smola, AJ. *Deep Sets*. In: *Advances in Neural Information Processing Systems*; vol. 30. Curran Associates, Inc.; 2017, p. 1–11.
- [27] Hüttenrauch, M, Sosc, A, Neumann, G. Deep Reinforcement Learning for Swarm Systems. arXiv:180706613 [cs, stat] 2019;ArXiv: 1807.06613.
- [28] Akiba, T, Sano, S, Yanase, T, Ohta, T, Koyama, M. Optuna: A Next-generation Hyperparameter Optimization Framework. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. Anchorage AK USA*: ACM. ISBN 978-1-4503-6201-6; 2019, p. 2623–2631. doi:10.1145/3292500.3330701.
- [29] Daniel, BC, Marques, R, Hoyet, L, Pettré, J, Blat, J. A Perceptually-Validated Metric for Crowd Trajectory Quality Evaluation. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 2021;4(3):1–18. doi:10.1145/3480136.

Appendix

Here we provide the training plots of the experiments described in Section 6.2. For each scenario we plot together the three model types (Direct Agent Perception, Raycasting and both) together, for each combination of observations and dynamics. For each experiment, we show both a plot with the full range of rewards, as well as a version zoomed in to the high range of rewards obtained in at least some of the experiments. We also show the energy expenditure, where lower values correspond to better policies. Each line is averaged across 8 independent runs which differ only by the random seed. The shaded area corresponds to the standard error of the mean. Note that the rewards cannot be directly compared between different scenarios.

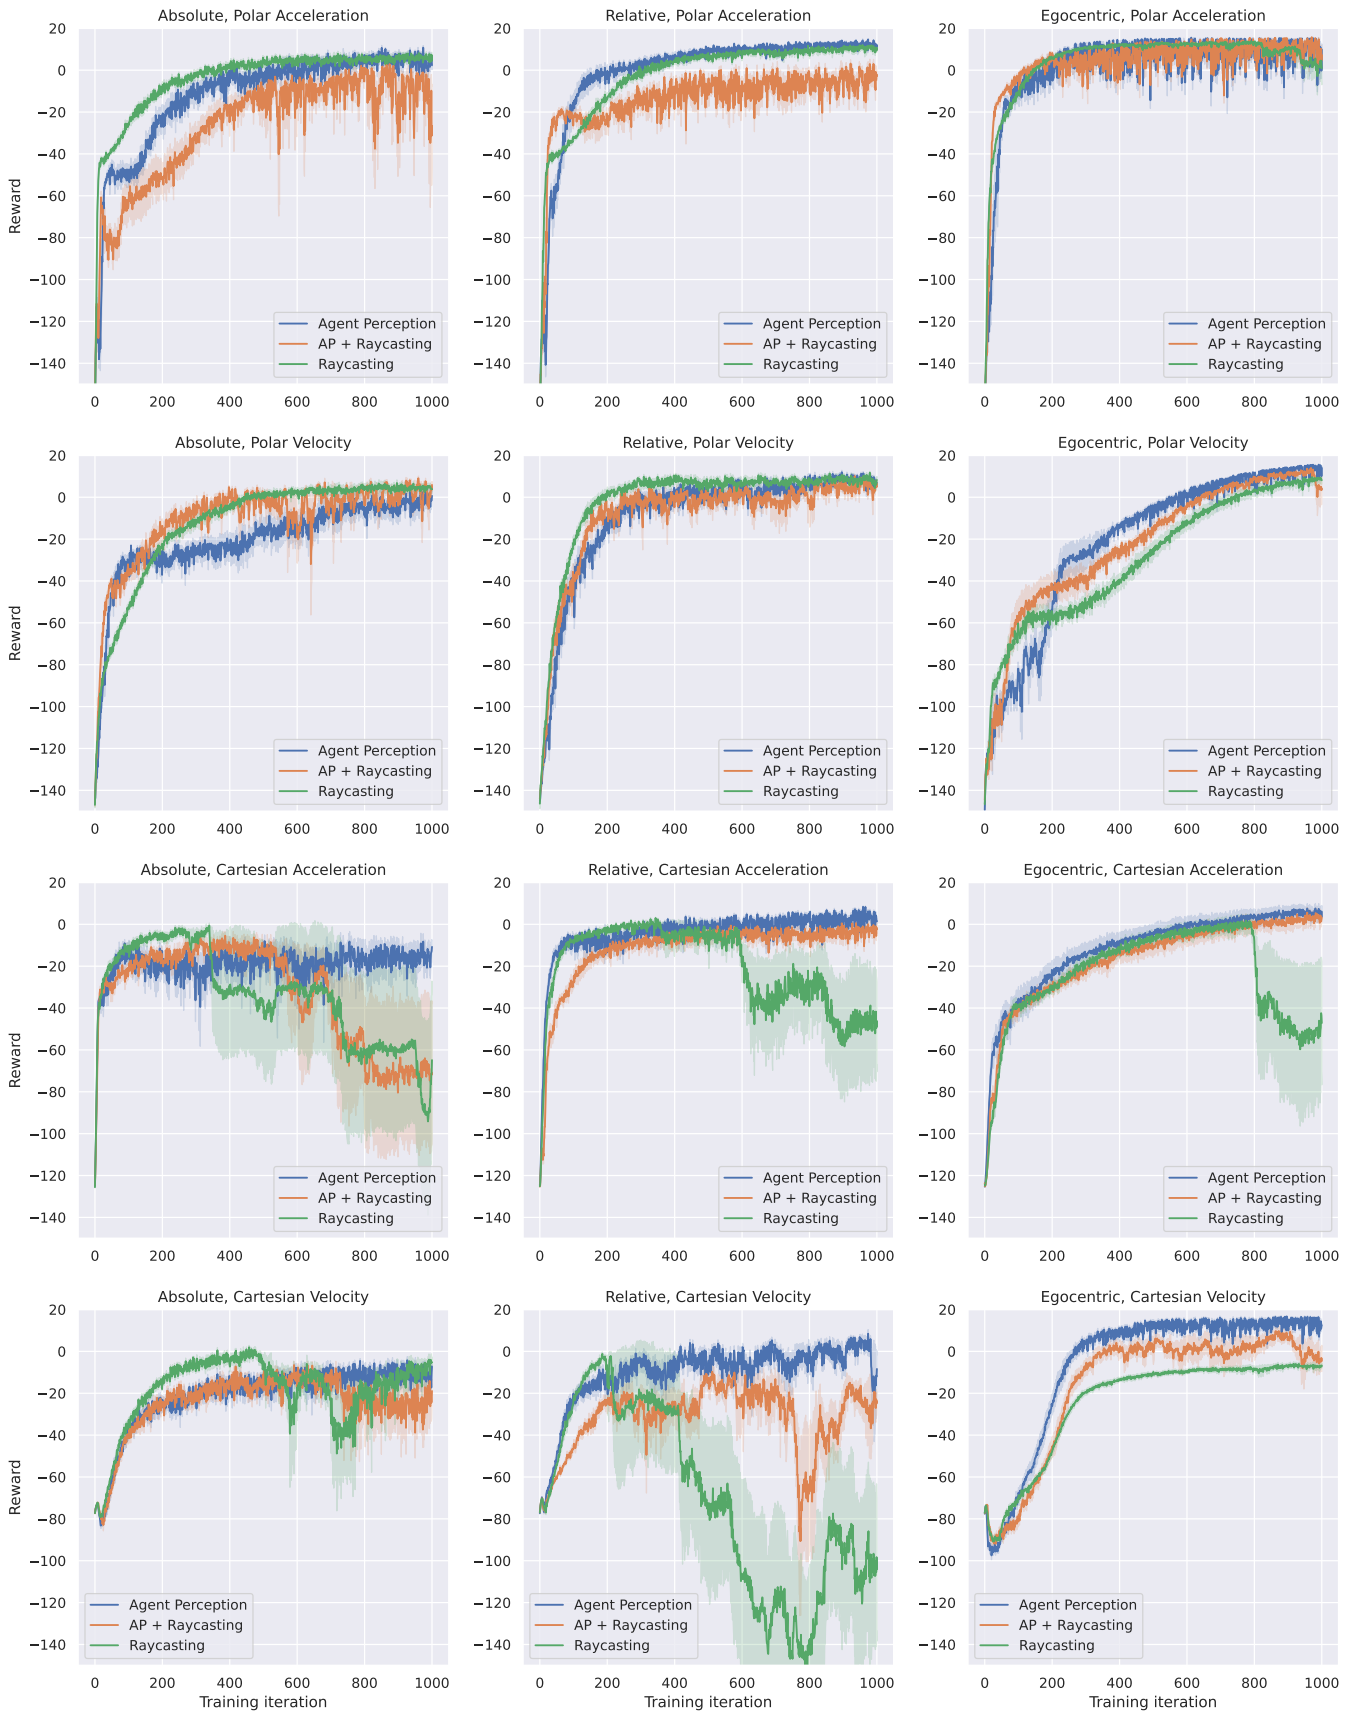


Fig. 11: Circle, full range of rewards

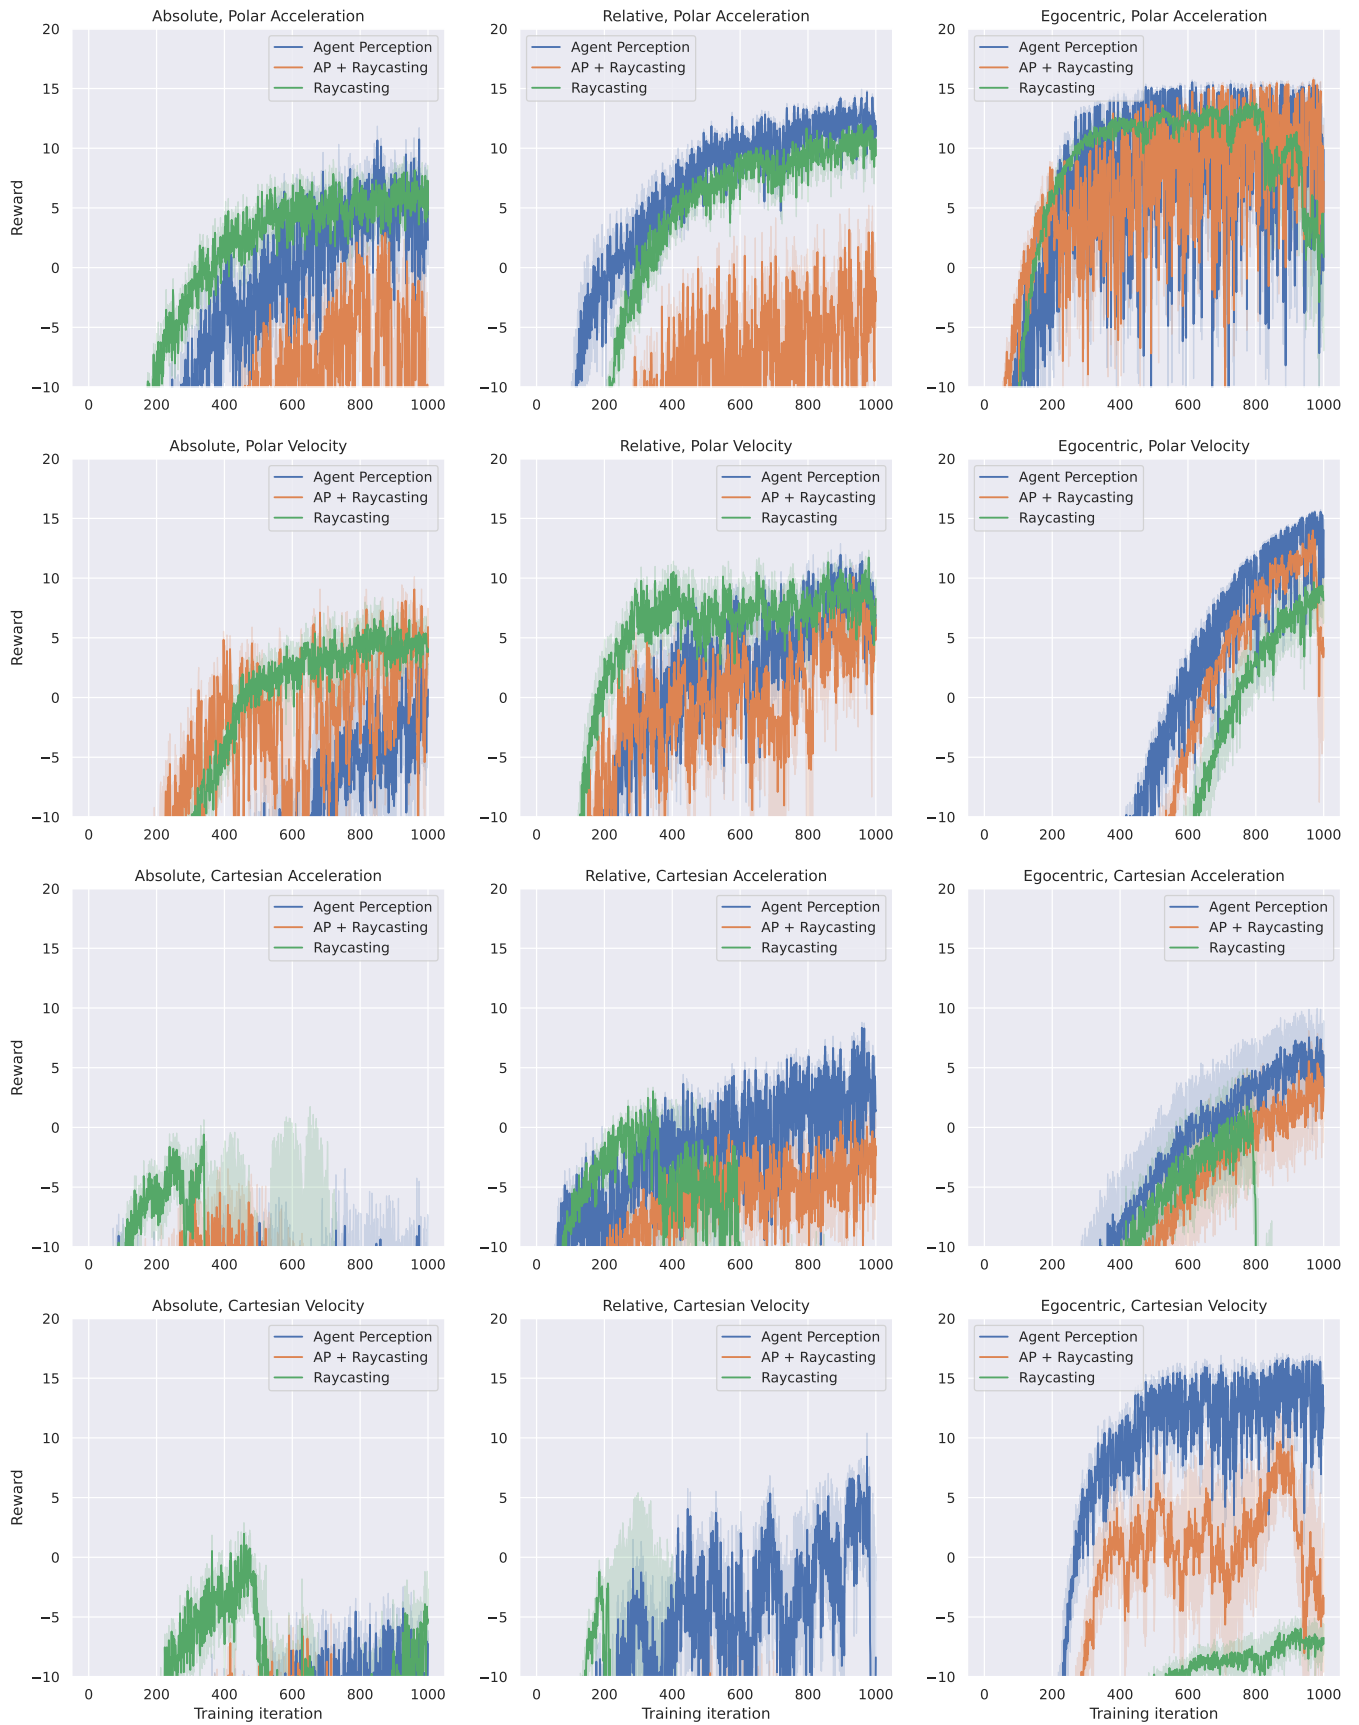


Fig. 12: Circle, high rewards

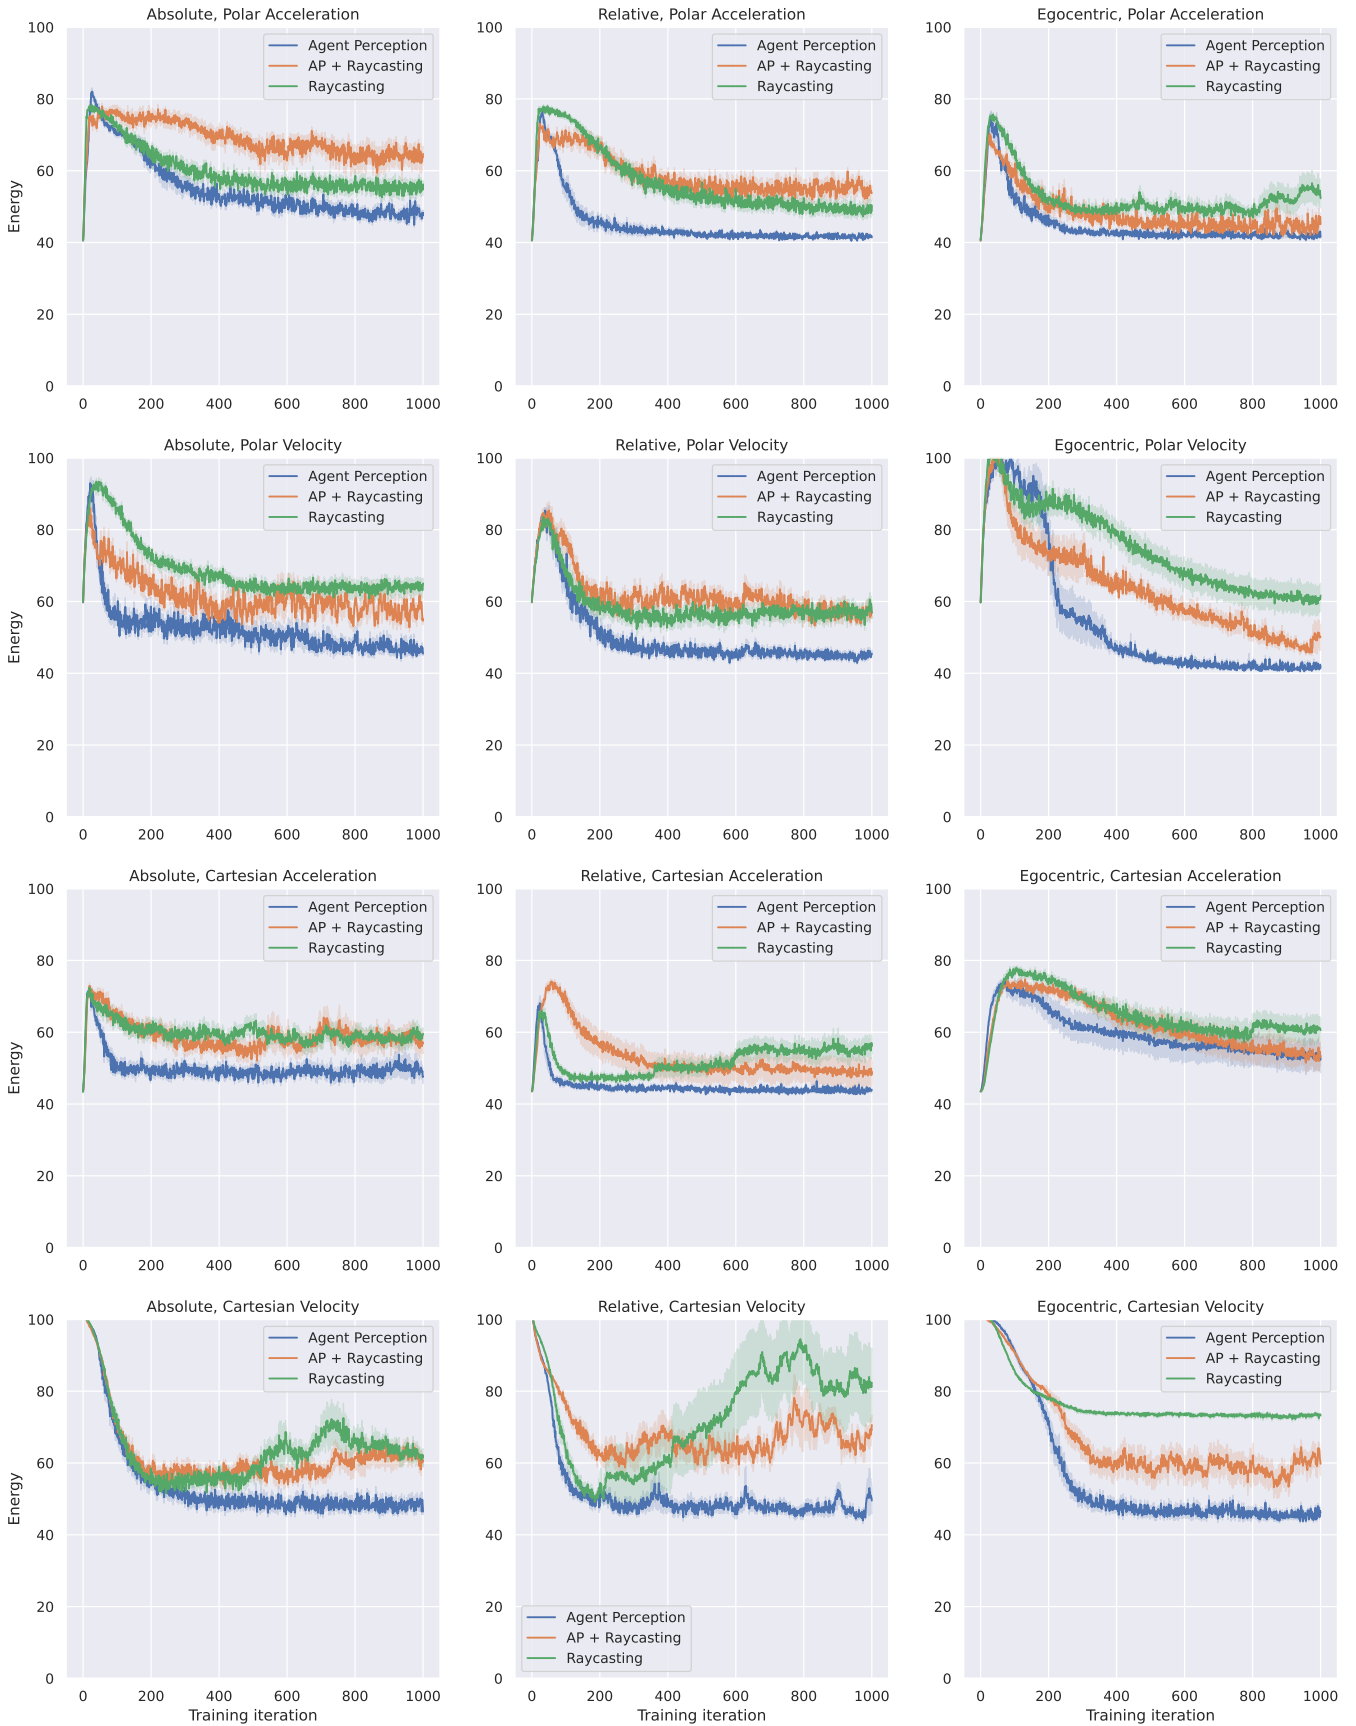


Fig. 13: Circle, energy

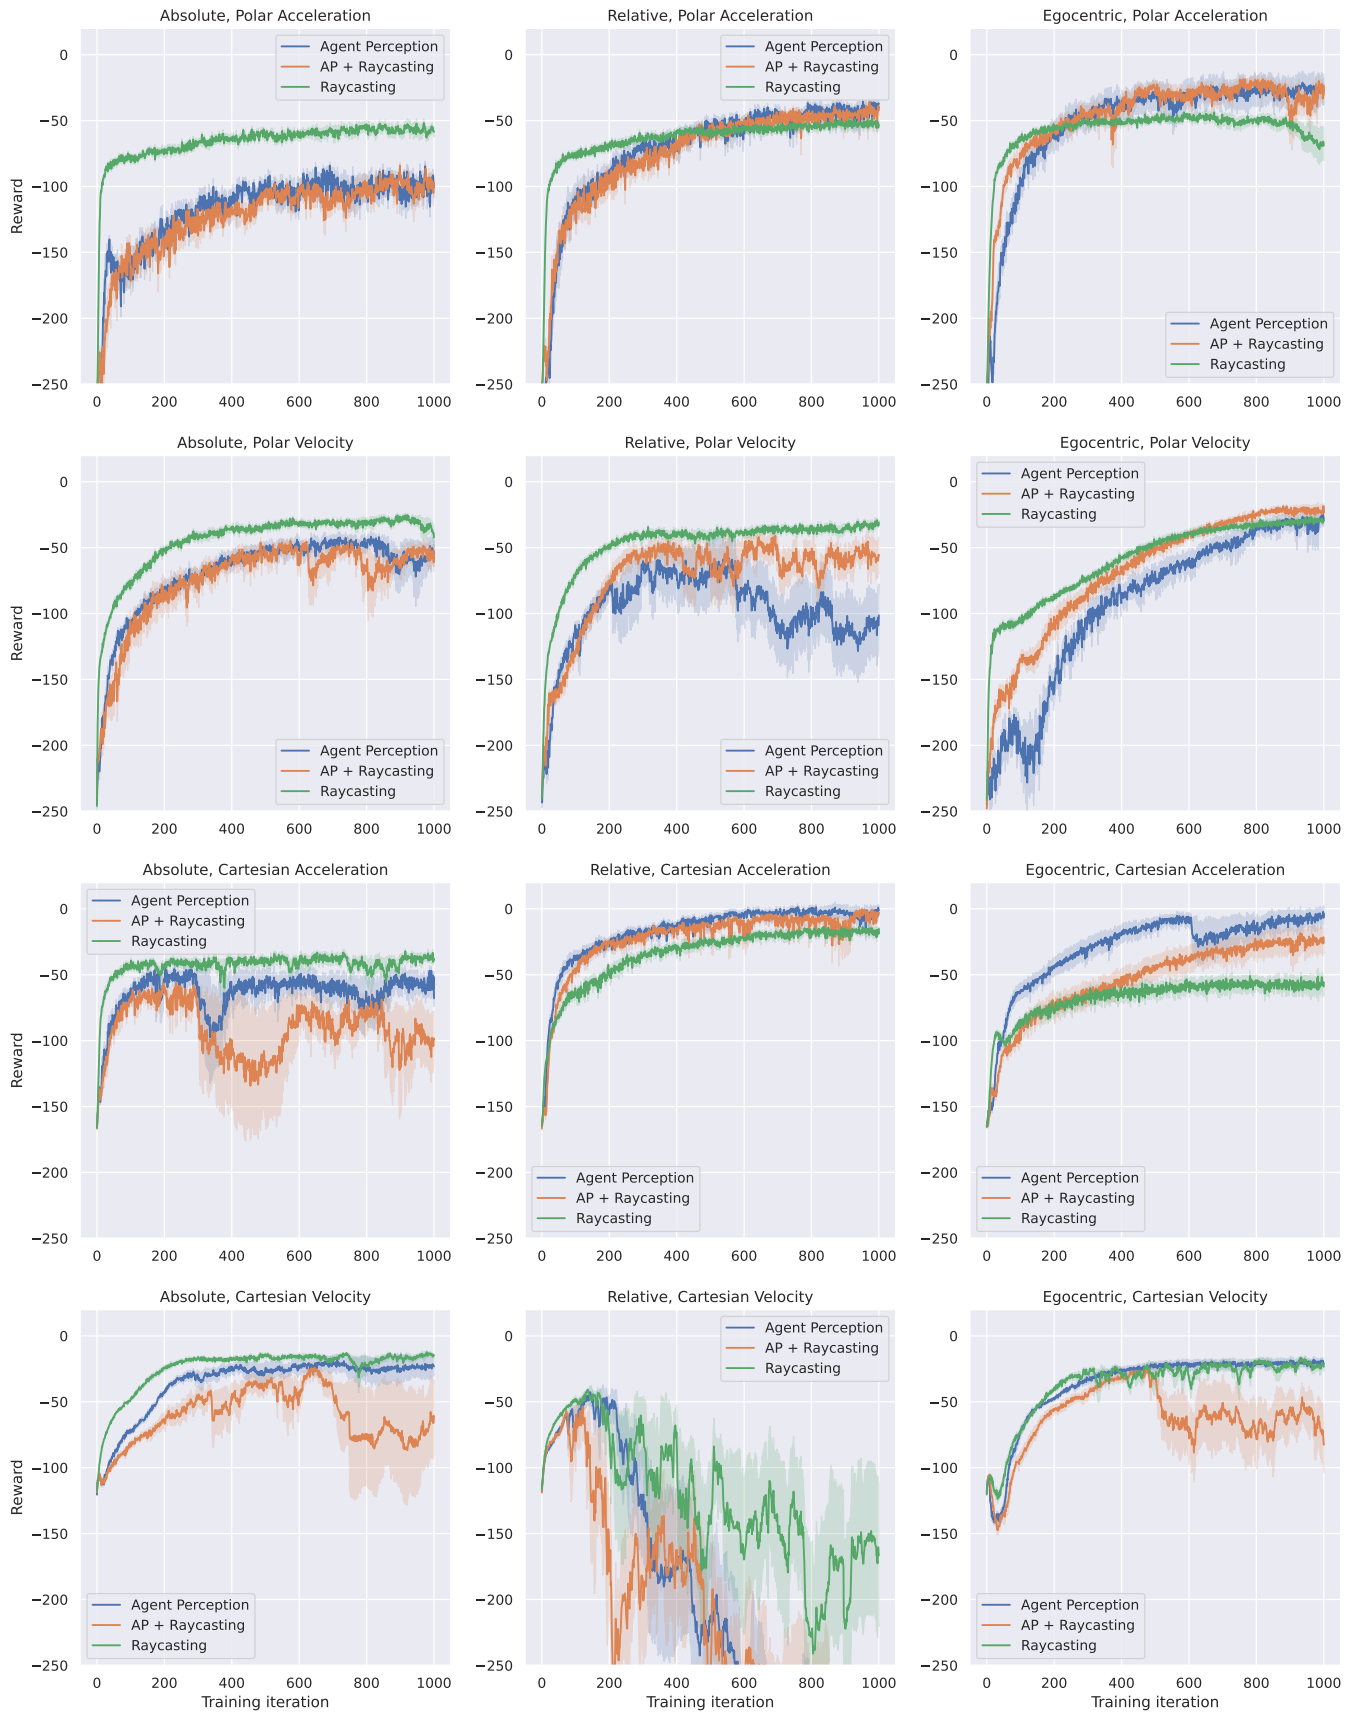


Fig. 14: Crossway, full range of rewards

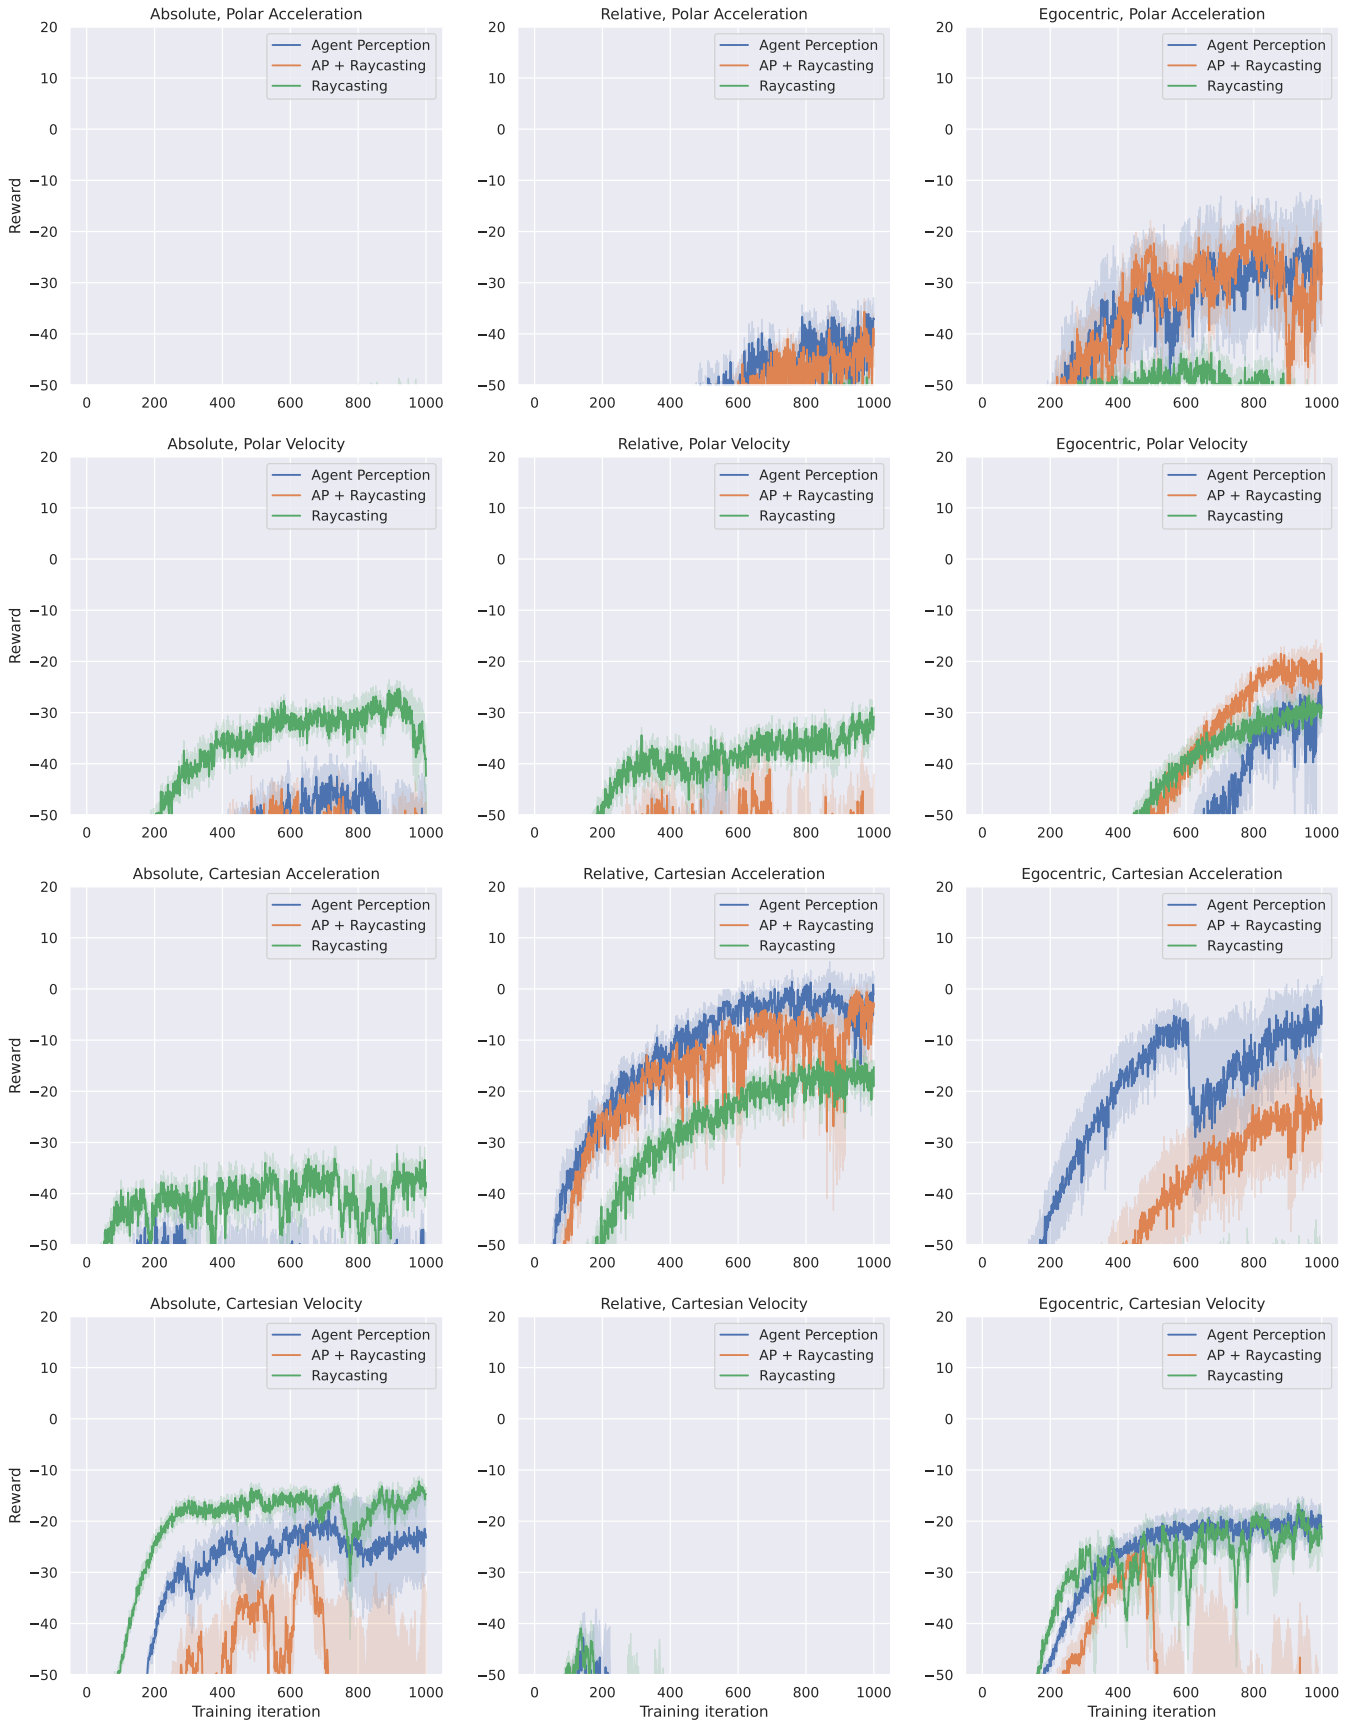


Fig. 15: Crossway, high rewards

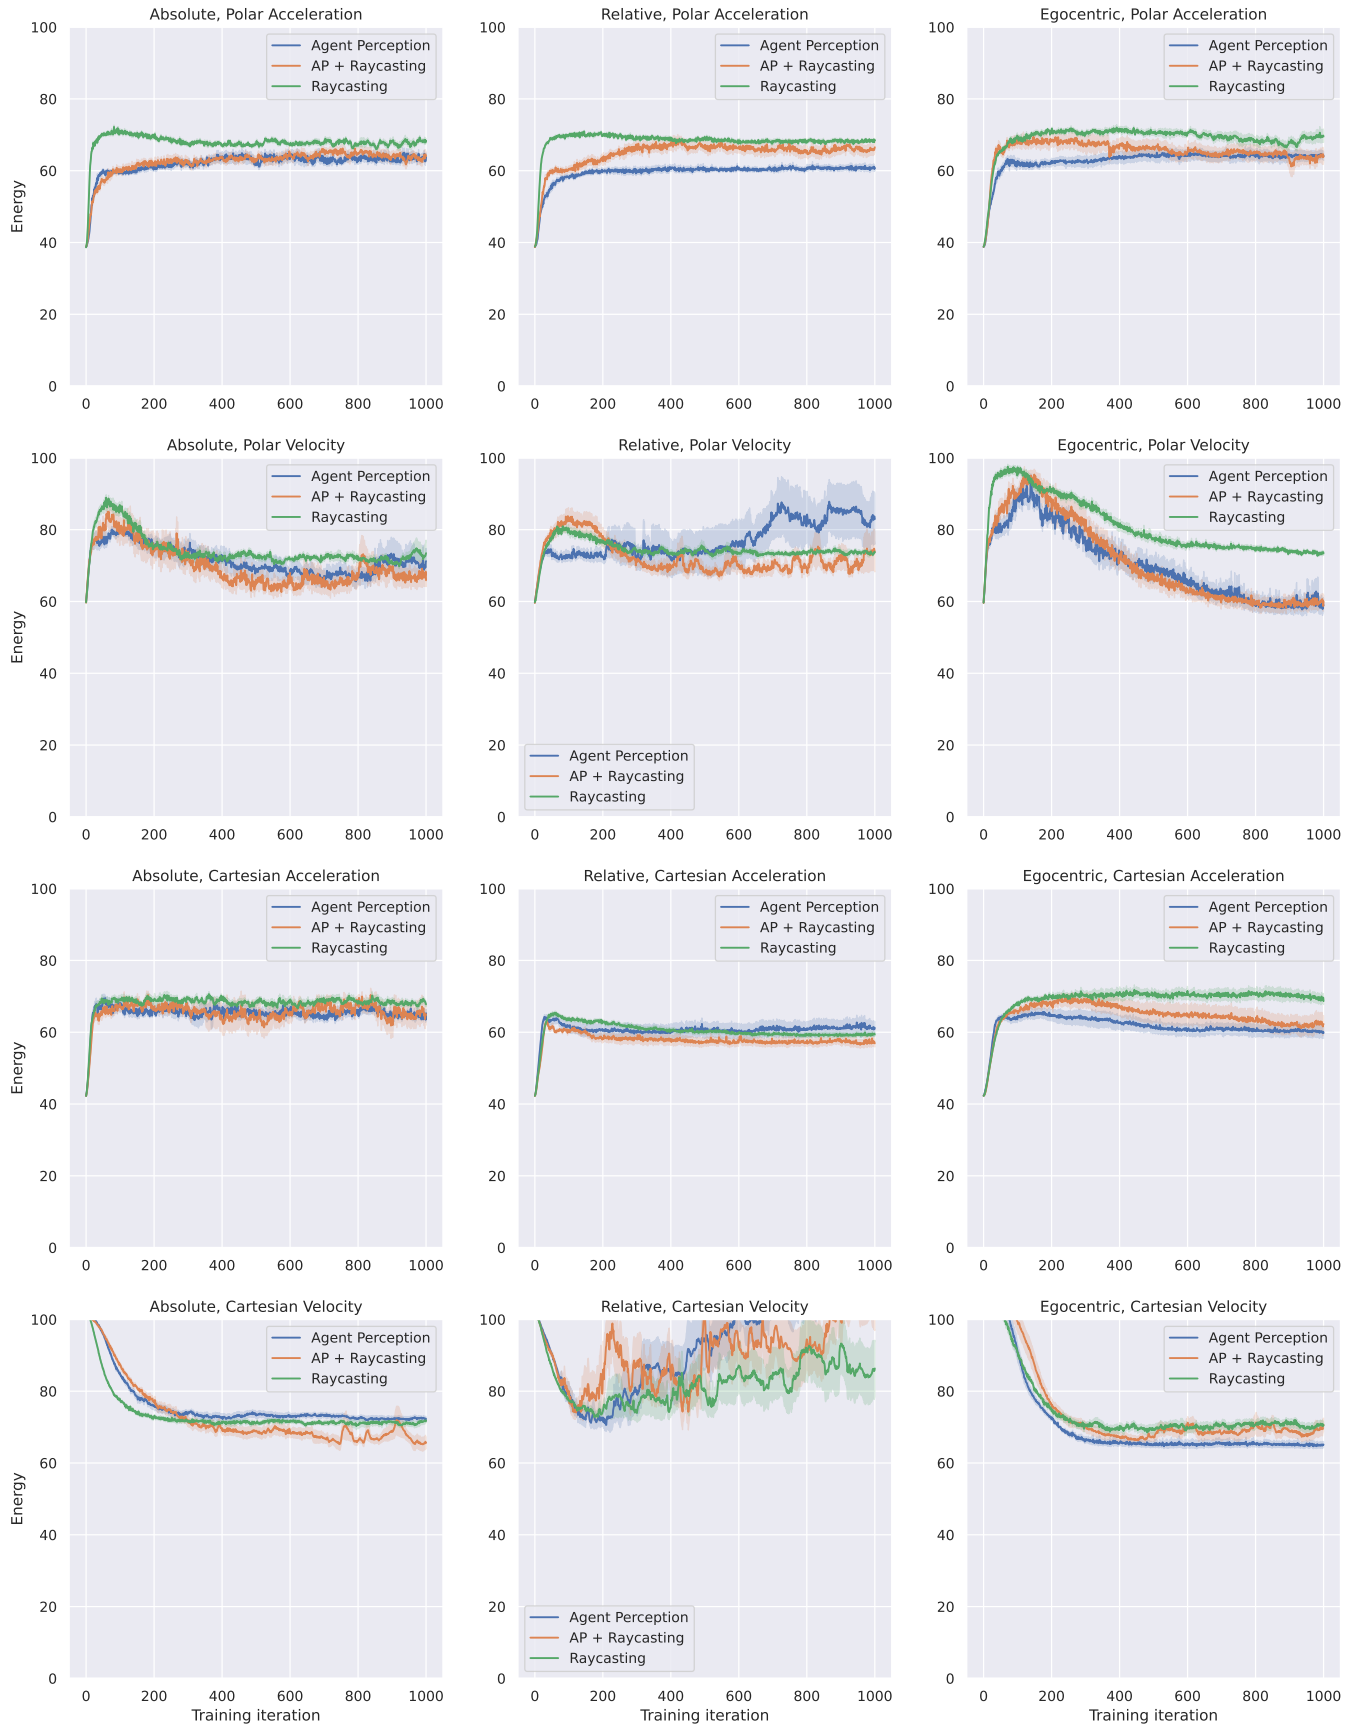


Fig. 16: Crossway, energy

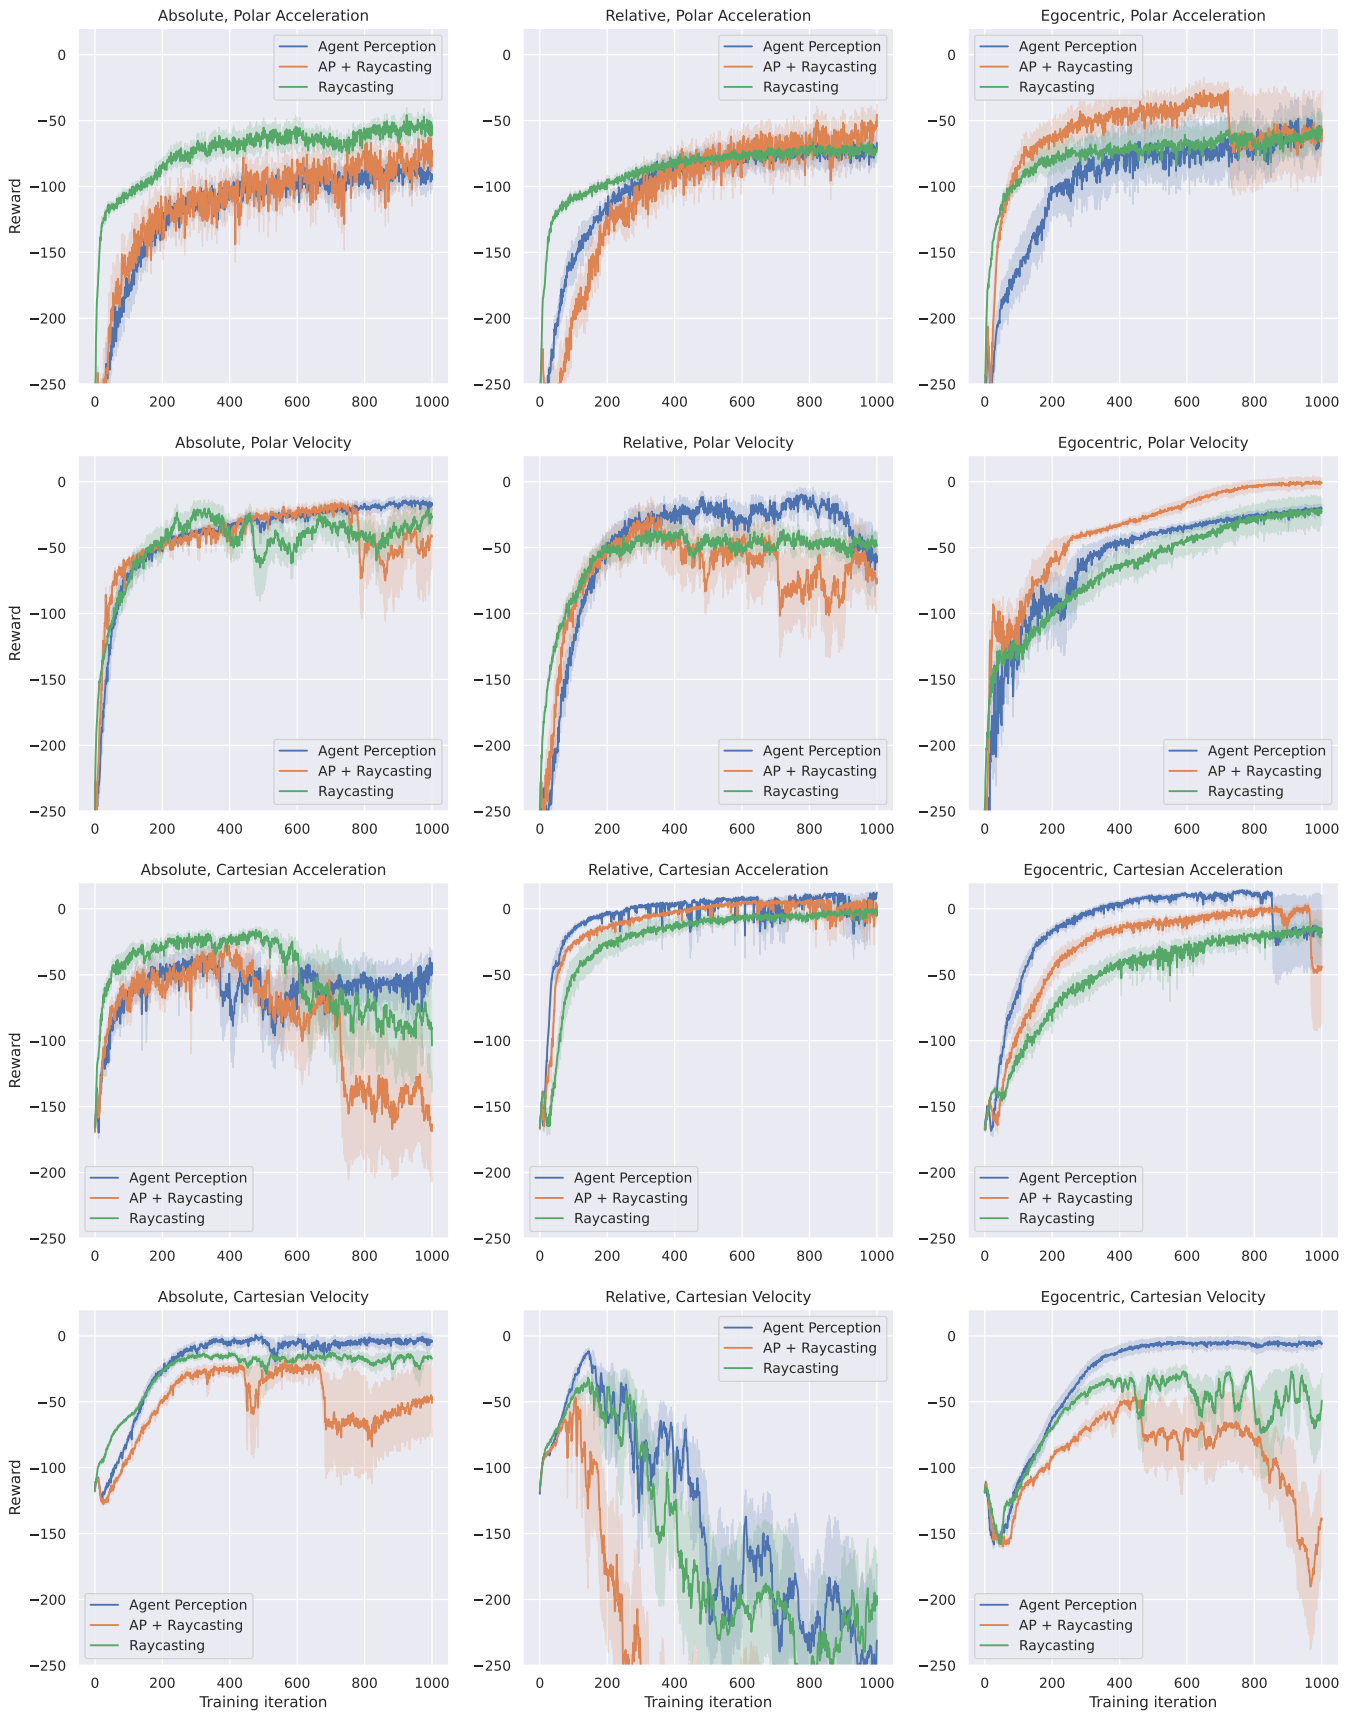


Fig. 17: Corridor, full range of rewards

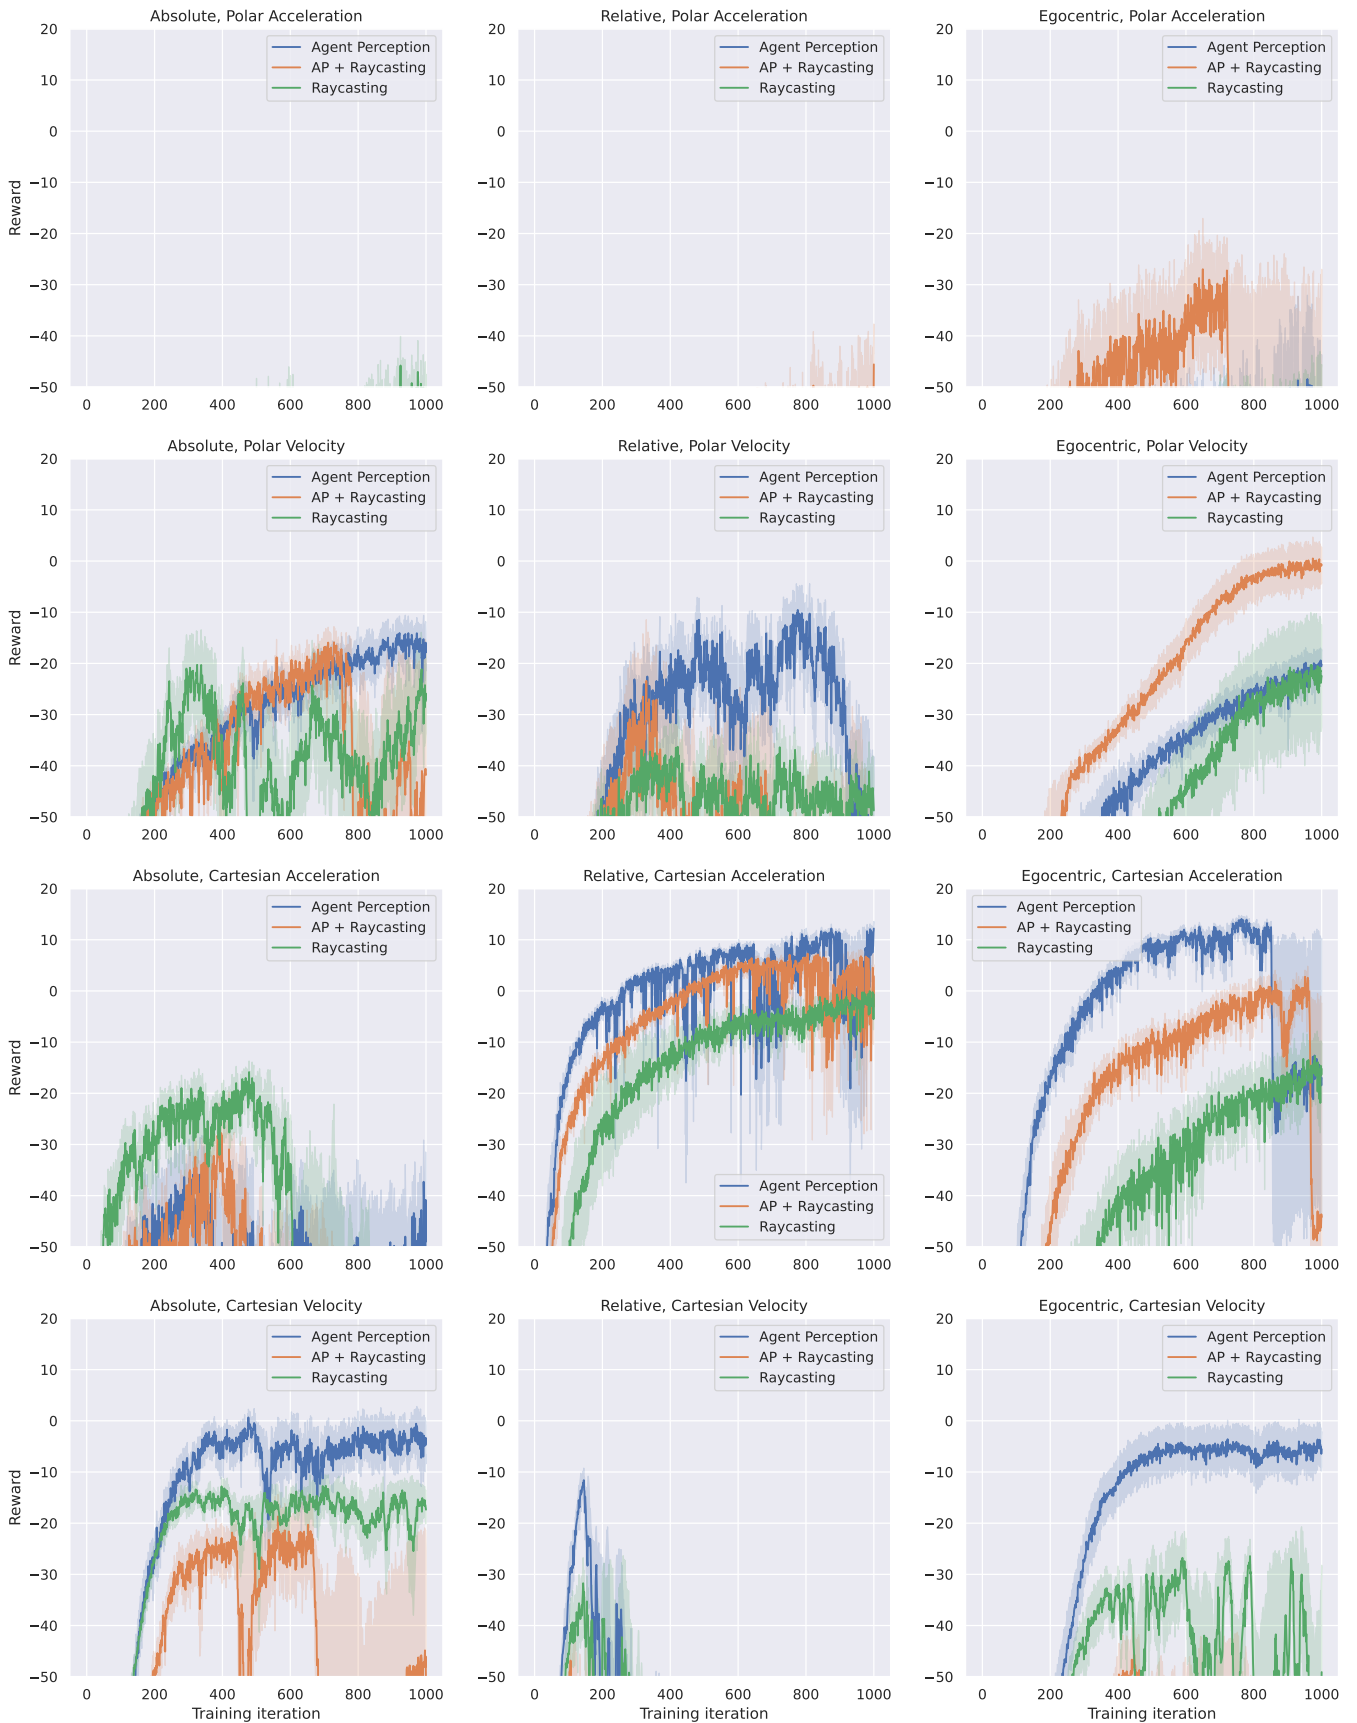


Fig. 18: Corridor, high rewards

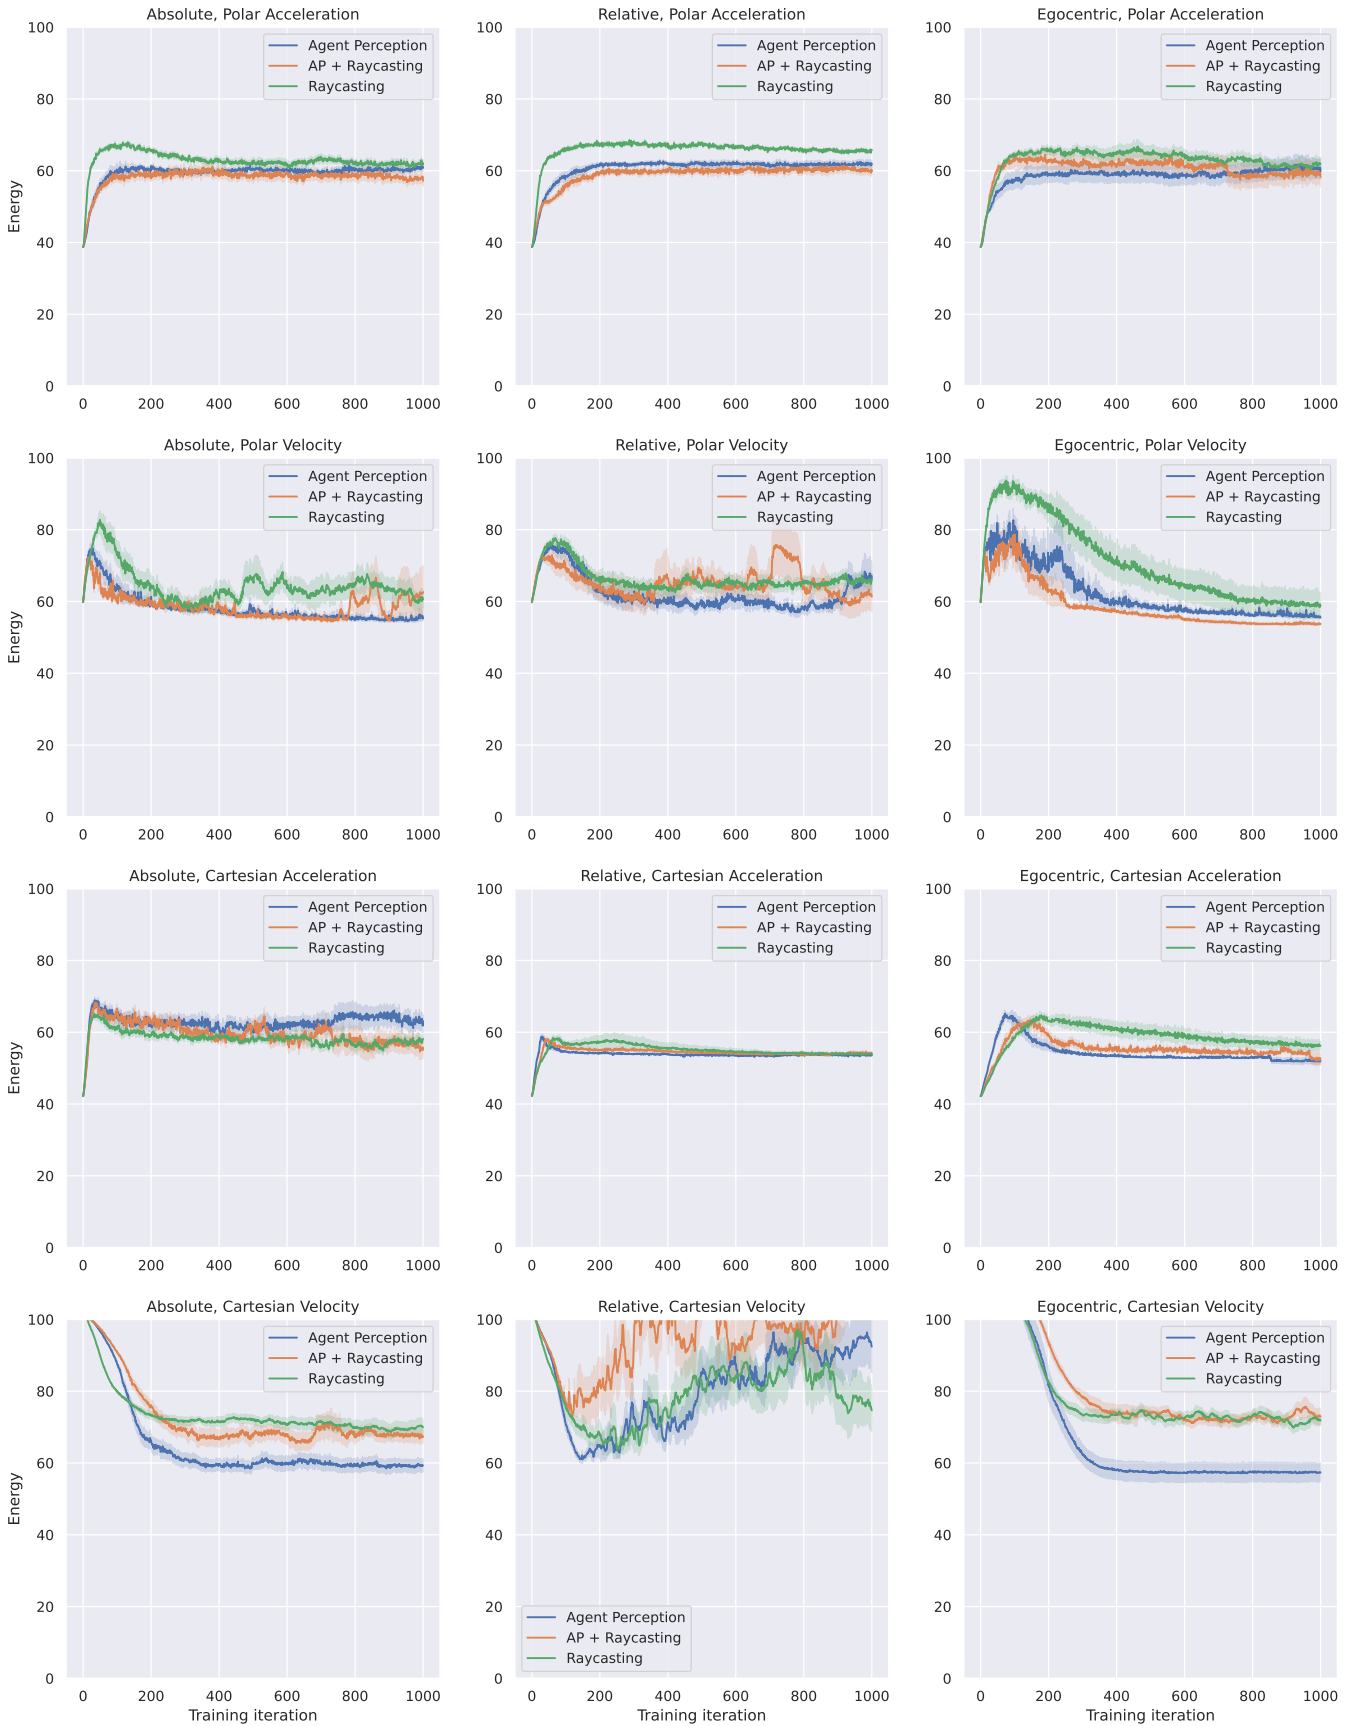


Fig. 19: Corridor, energy

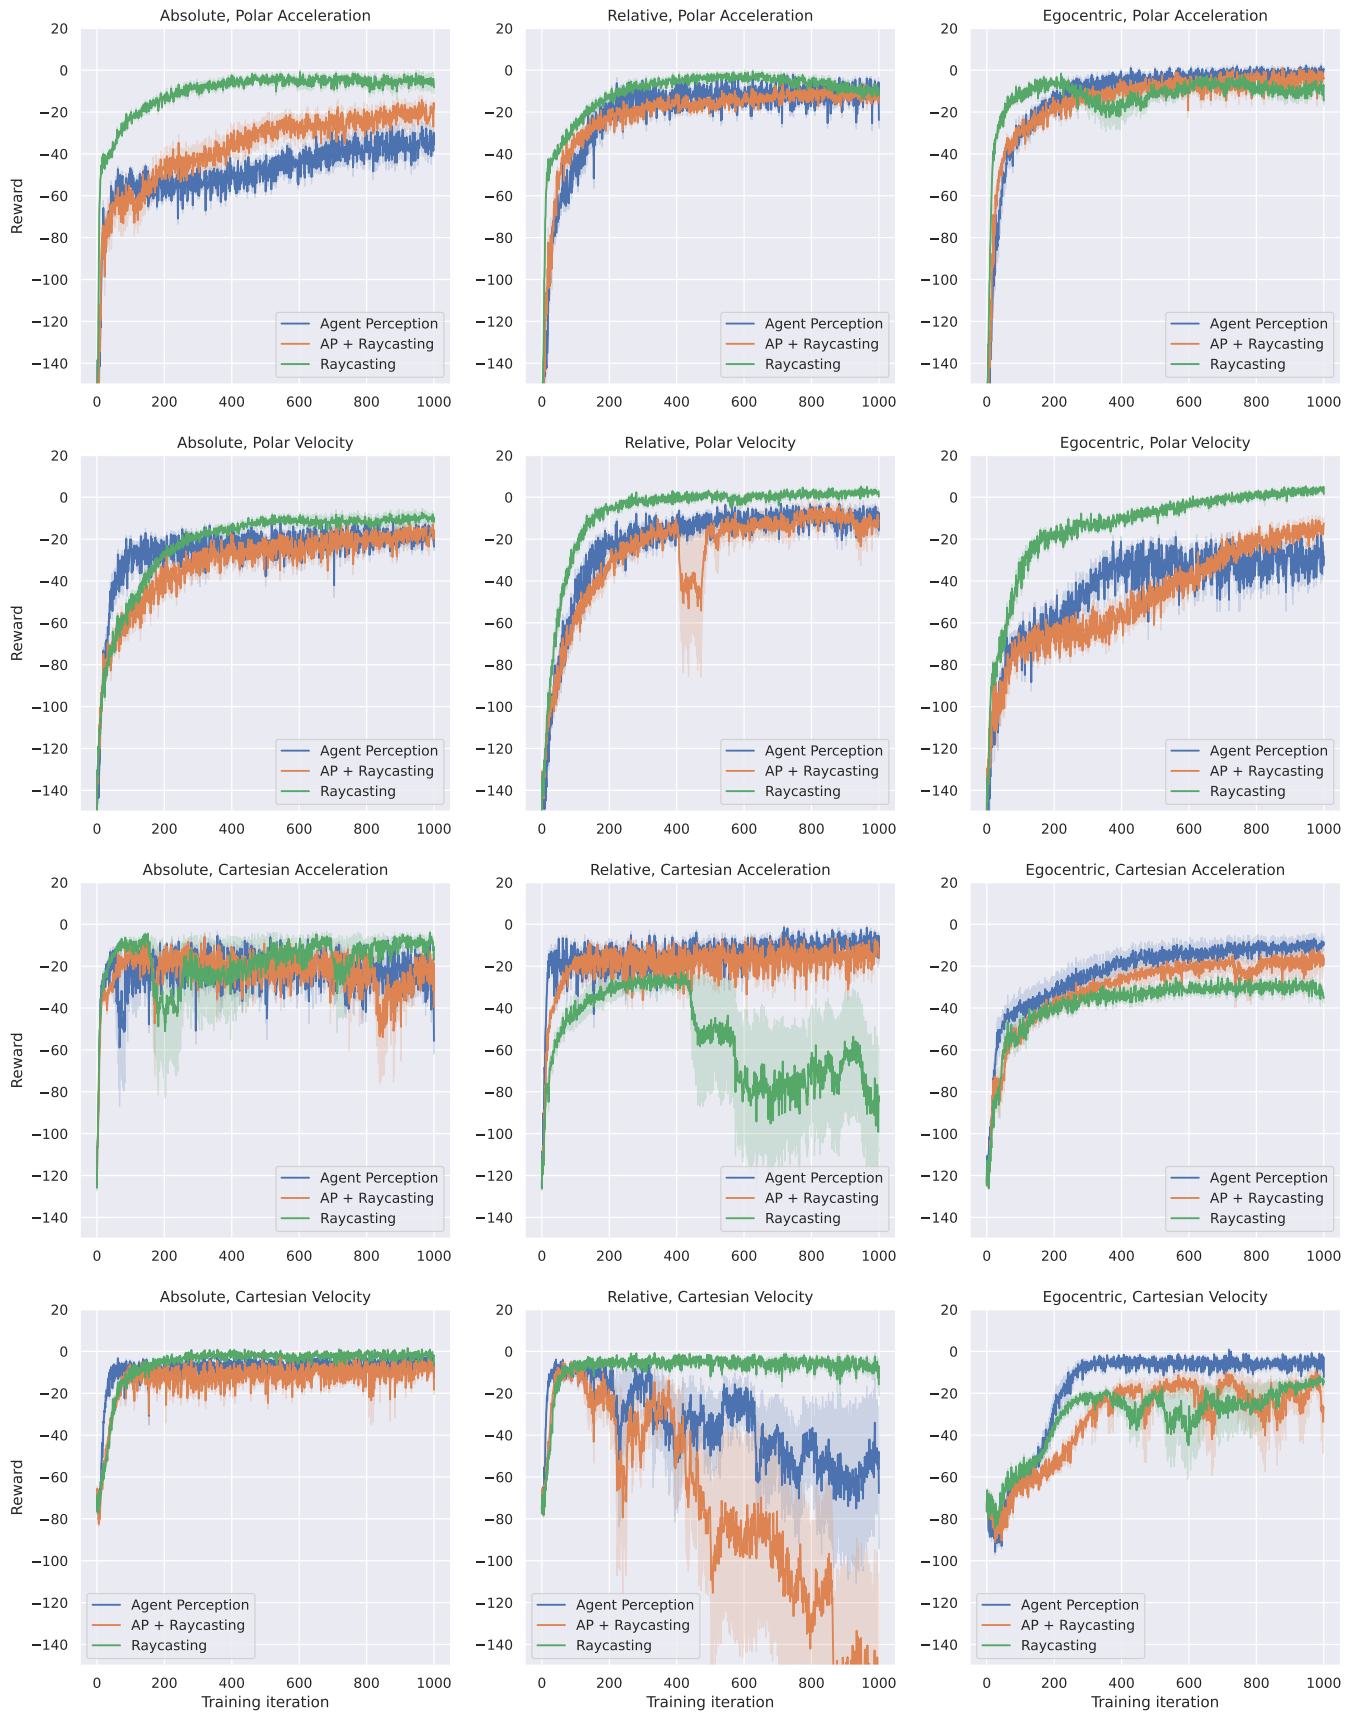


Fig. 20: Random, full range of rewards

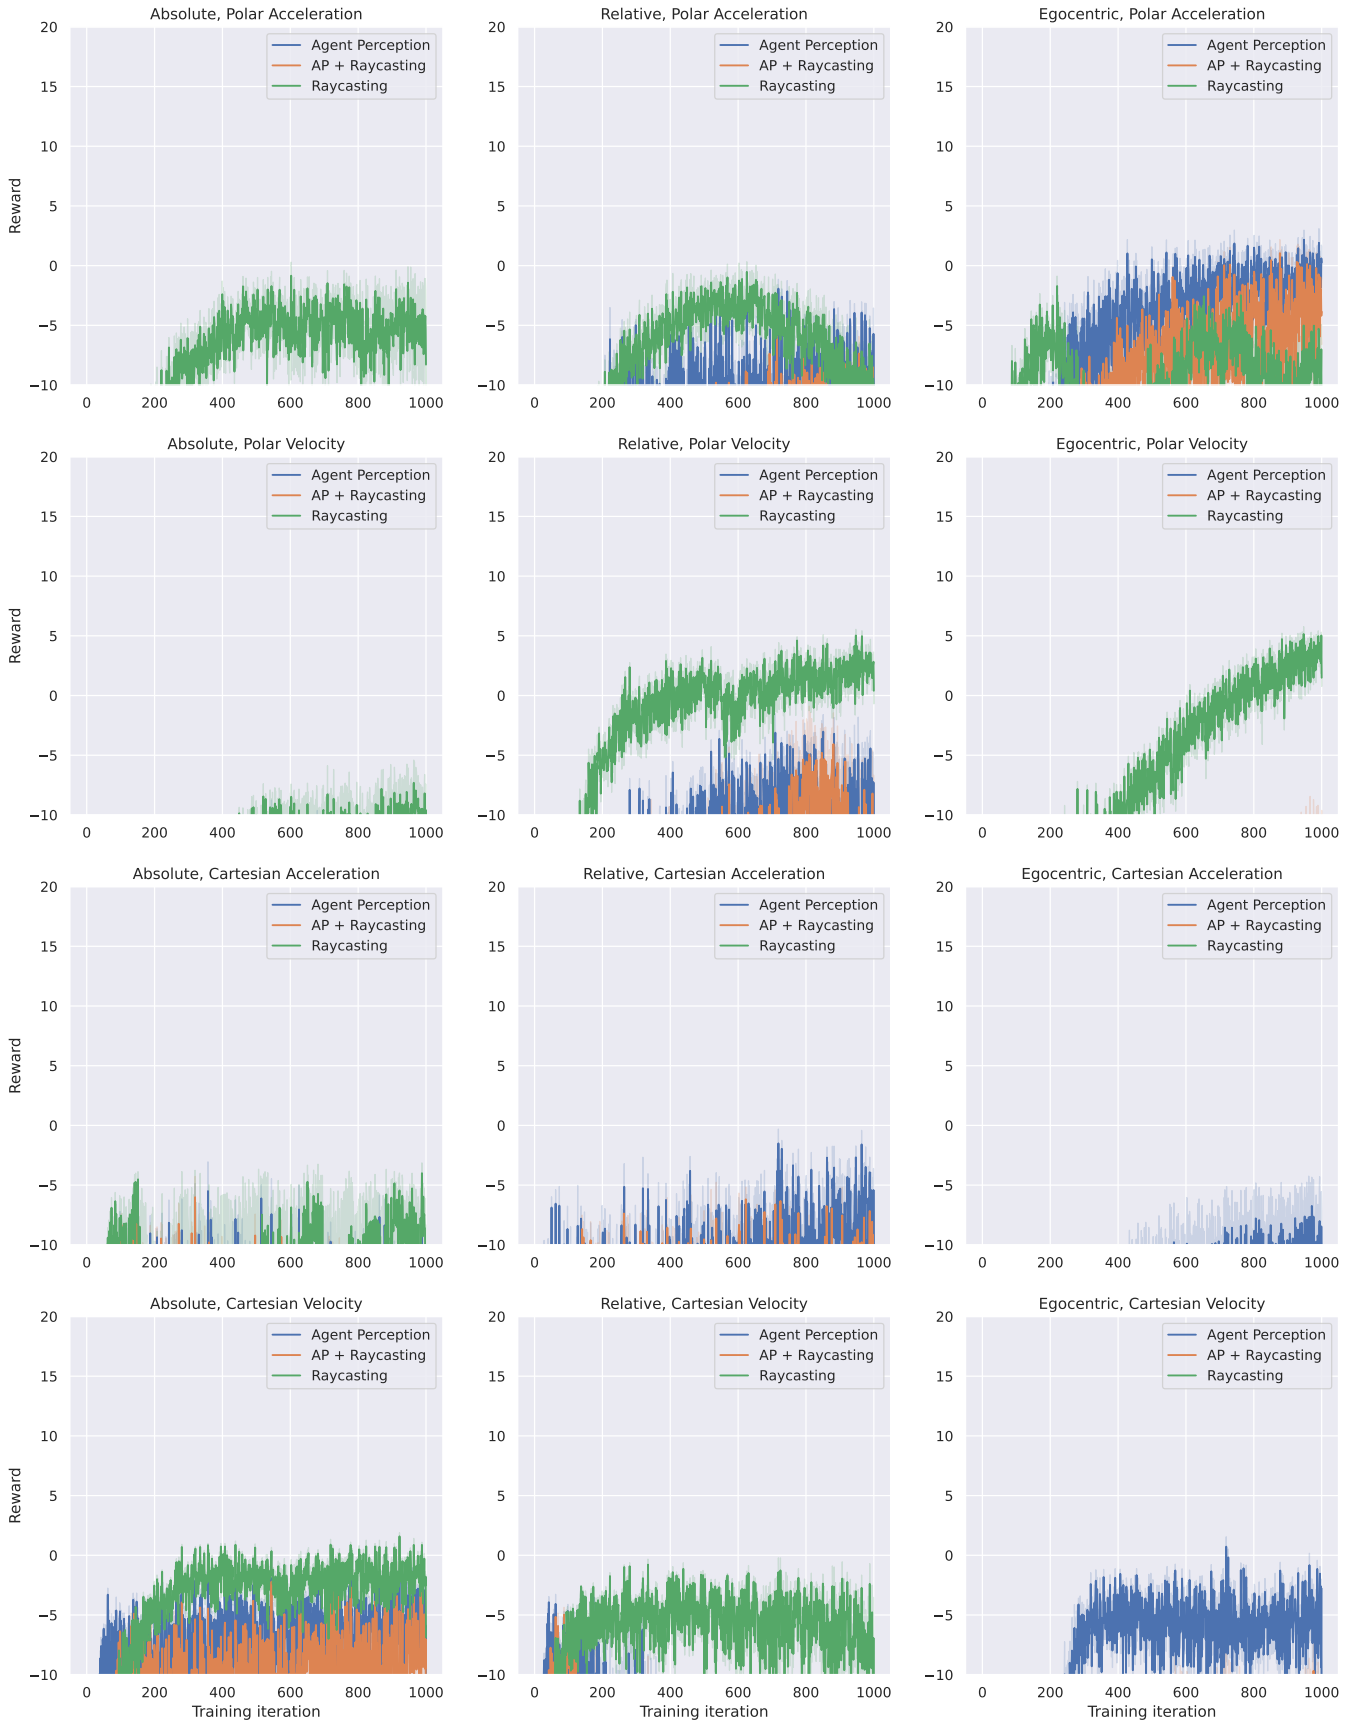


Fig. 21: Random, high rewards

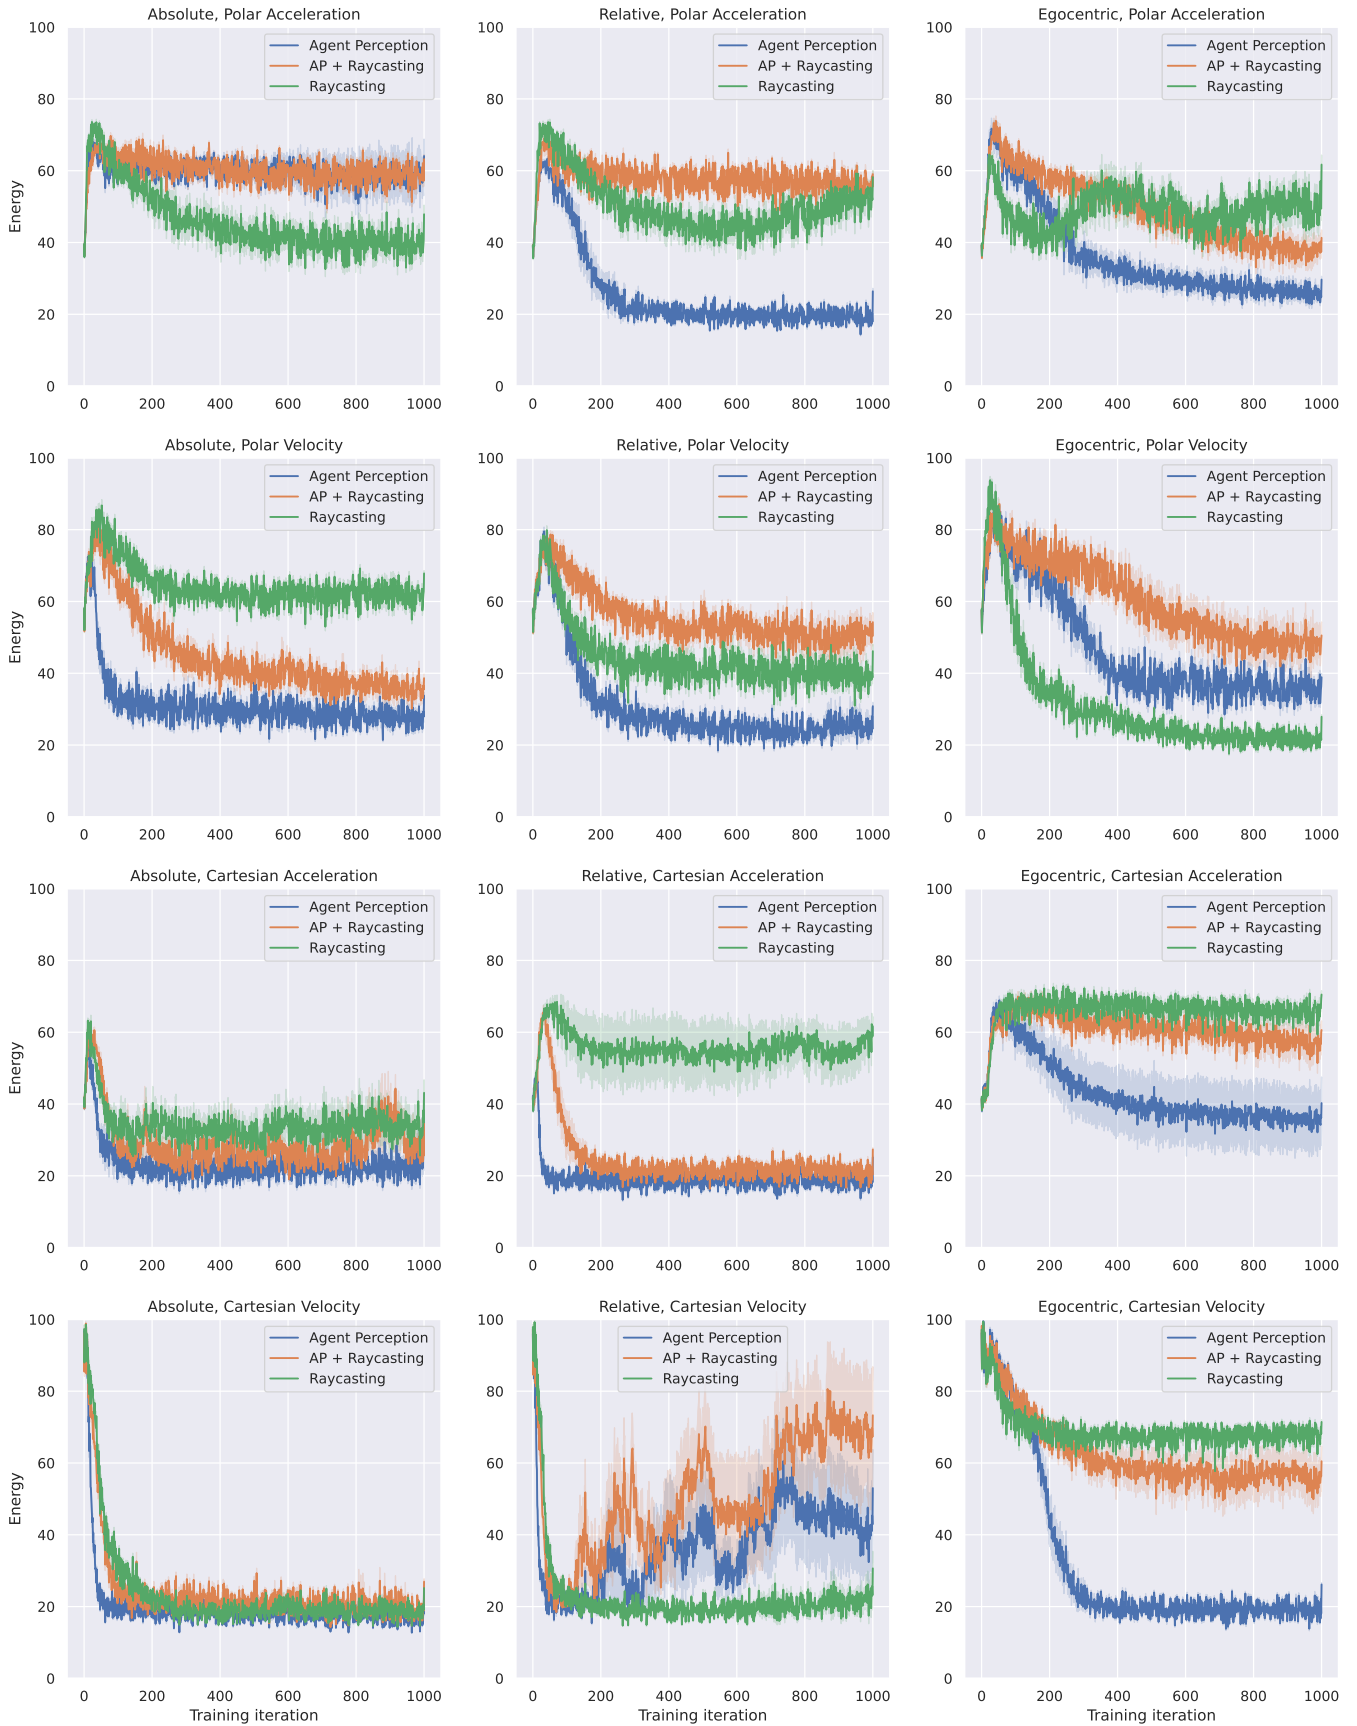


Fig. 22: Random, energy