Title : MAIIL - AI-driven character simulation based on Multi-Agents Interaction Imitation Learning

Host institution and proposed thesis advisor: Inria, Franck Multon

Context

Simulation of autonomous entities capable of collaborating or competing with real humans, as real other humans would in the same situation, is still a scientific and technological challenge. This is especially a complex problem when dealing with fullbody interaction between simulated and real entities, as it involves complex spatiotemporal constraints at different timing scales (continuous motion and sequence of actions), with a high-dimension research space (associated with the characters' degrees of freedom). Let us consider several characters, real and simulated humans, that have to collaborate to manipulate a digital mockup. Another interesting example is the coordination of (simulated and real) humans to compete against another team. In both cases, the simulated entities should behave to help or compete with real humans, as other real humans would in this situation, in a long-term close collaboration/competition. Hence, the problem is not to find an optimal solution to efficiently achieve the task, but act as expected by the participants, finding suboptimal but realistic solution.

Scientific objectives

The goal of MAIIL is to propose a framework to automatically design autonomous entities controllers, that can imitate the interaction behavior of humans in group tasks, either collaborative or competitive. To this end, the model should be able to learn the controller from a minimal set of interaction examples. The key point is to learn how humans solve the interaction problem, with a given style, which might be a suboptimal solution to achieve the common goal. One key problem is to imitate long term and continuous strategies, not limited to only reacting to instantaneous events.

Although the controller design step could lead to offline learning, the resulting controller should produce realtime response to the current state of the group. In this project, user evaluation of the solution is a key element of success, as we wish to obtain realistic behaviors that preserve the style given by the training examples.

Approach / Originality

MAIIL will explore how to adapt Multi-Agent Generative Adversarial Imitation Learning (MAGAIL) framework [Song et al., 2018] to continuous fullbody interaction, and to specific interaction style imitation. The goal is to make autonomous entities controlled by this approach be considered as real humans, or make participants be confident about its actions, as they would do for real humans. We will also explore how to extend natural language processing approaches to take multidimensional state representation of human-to-human interaction (including joint angles, or body-to-body coordination such as interactions meshes introduced in computer animation).

State of the art

When a few examples of interactions are available, reinforcement learning is a promising way to control physics-based characters. [Haworth et al. 2020] proposed a hierarchical policy that incorporates navigation, footstep planning, and bipedal walking skills, for controlling navigation of pedestrians. Unlike previous approaches, this method learns control policies that can handle interactions between multiple simulated humanoids. [Won et al. 2021] proposed a two-steps

approach that first learns an imitation policy from single-actor motion capture data, then transfers it into competitive policies. [Liu et al. 2022] trained football teams of physically simulated humanoids in a sequence of training stages using a combination of imitation learning, single/multi-agent reinforcement learning and population-based methods. However, these approaches have not been designed to imitate a specific style of interaction given by a few examples. To achieve this goal, we have proposed MAAIP Multi-Agents Adversarial Interaction Priors [Younes et al., 2022] for imitation of boxing interactions between two players. But it cannot handle middle term strategy, and is limited to reactive behaviors, not continuous collaborative or competitive behaviors, such as displacing a large object with a group. Moreover, up to our knowledge, no user study has been carried-out to evaluate the actual impact on collaborating with these AI-driven entities.

Priorité #3, PC3 MATCHING : Modélisation et compréhension des interactions collaboratives ou compétitives entre les humains et les entités pilotées par l'IA

Project organisation

The project is organized in three years, with 3 deliverables:

- Expand the MAAIP approach to more than two participants, continuous adaptation instead of reaction to events, and to middle/long term strategy prediction. (M1-24)
- Implement the method on autonomous virtual humans on a given collaborative or competitive scenario, to be used in VR with real users. (M24-M28)
- Carry-out user studies to collect feedback from real users, and analyze the impact on collaborative or competitive group behaviors mixing real and simulated participants. (M28-M36)

Co-supervisors: Richard Kulpa (M2S Lab: Virtual Reality, Autonomous Virtual Humans, Perception-Action Coupling), Ewa Kijak, Simon Malinowski (IRISA, LinkMedia: machine learning and computer vision, multimodal analysis)

References

Brandon Haworth, Glen Berseth, Seonghyeon Moon, Petros Faloutsos, and Mubbasir Kapadia. 2020. Deep integration of physical humanoid control and crowd navigation. In Motion, Interaction and Games. 1–10.

Siqi Liu, Guy Lever, Zhe Wang, Josh Merel, SM Ali Eslami, Daniel Hennes, Wojciech M Czarnecki, Yuval Tassa, Shayegan Omidshafi ei, Abbas Abdolmaleki, et al. 2022. From motor control to team play in simulated humanoid football. Science Robotics 7, 69 (2022), eabo0235.

Jiaming Song, Hongyu Ren, Dorsa Sadigh, and Stefano Ermon. 2018. Multi-Agent Generative Adversarial Imitation Learning. In Advances in Neural Information Processing Systems, Vol. 31.

Jungdam Won, Deepak Gopinath, and Jessica Hodgins. 2021. Control strategies for physically simulated characters performing two-player competitive sports. ACM Transactions on Graphics (TOG) 40, 4 (2021), 1–11.

Mohamed Younes, Ewa Kijak, Simon Malinowski, Richard Kulpa, Franck Multon. 2022. AIP: Adversarial Interaction Priors for Multi-Agent Physics-based Character Control. SIGGRAPH Asia 2022 Posters, 1-2