Intrinsically Motivated General Companion NPCs
Extended Abstract

Christian Guckelsberger[1], **Christoph Salge**[2] and **Simon Colton**[1,3]
[1]Computational Creativity Group, Goldsmiths, University of London, UK
[2]Adaptive Systems Research Group, University of Hertfordshire, UK
[3]The Metamakers Institute, Falmouth University, UK
c.guckelsberger@gold.ac.uk

Abstract

Non-player characters (NPCs) in games are traditionally hard-coded or dependent on pre-specified goals, and consequently struggle to behave sensibly in ever-changing and possibly unpredictable game worlds. To make them fit for new developments in procedural content generation, we introduce the principle of *coupled empowerment maximisation* as an intrinsic motivation for game NPCs. We focus on the development of a general game companion, designed to support the player in achieving their goals. We evaluate our approach against three intuitive and abstract companion duties. We develop dedicated scenarios for each duty in a dungeon-crawler game testbed, and provide qualitative evidence that the emergent NPC behaviour fulfils these duties.We argue that this generic approach can speed up NPC AI development, improve automatic game evolution and introduce NPCs to full game-generation systems.

Introduction

Ellie from *The Last of Us* or the unnamed pet from *Nethack* – memorable companions are an important part of our gaming experience. But companions can also be a great source of annoyance or shatter our game immersion, especially when their behaviour fails miserably (Cerny 2015). The vast majority of companions are hard-coded by means of e.g. finite state machines or behaviour trees, and consequently struggle to produce believable or even plausible behaviour in unforeseen contexts (Merrick 2008; Forgette and Katchabaw 2014). More advanced companions can adapt their behaviour by means of planning, or by learning a policy via neural networks or traditional reinforcement learning. Nevertheless, these methods require intense training or pre-specified rewards, which again renders them inflexible when facing dynamic game worlds. In the future, the demands on Non-Player Character (NPC) AI in general are likely to increase further (Smith 2014). This is particularly emphasised by progress in procedural content generation, which not only focusses on game elements such as levels and game mechanics (Liapis, Yannakakis, and Togelius 2014; Smith 2014), but ultimately aims at generating entire games (Cook, Colton, and Gow 2016). How can NPCs deal with these ever-changing and potentially unpredictable game worlds?

We suggest to motivate NPCs intrinsically. Instead of relying on pre-defined goals which might become meaningless when the game changes, intrinsically motivated agents perform "an activity for its inherent satisfactions rather than for some separable consequence" (Ryan and Deci 2000). We work with the intrinsic motivation of *empowerment* (Klyubin, Polani, and Nehaniv 2008), a measure of how much an agent is in control of the world it can perceive. But while empowerment might allow to produce an intrinsically motivated *general NPC*, we have to look specifically into how to turn it into a good *companion*.

Background

Merrick (Merrick 2008) investigated how intrinsically motivated reinforcement learning (Singh, Barto, and Chentanez 2004) can support NPCs in learning complex tasks in a dynamic game world. They propose two models of motivation as reward signals for Q-learning: an agent's interest in a new situation, given past experiences, and its competence based on the error in learning policy updates. Qualitative studies in *Second Life* and a quantitative analysis of behavioural variety and complexity in dedicated RPG testbeds confirm that intrinsic motivation allows agents to adapt their behaviour in a changing environment. In contrast to our study, their NPCs act in solitude, and not in favour of the player.

Empowerment as intrinsic motivation has so far only been employed for general game-playing, but not to drive the behaviour of companion or enemy NPCs. Anthony, Polani and Nehaniv analysed empowerment maximisation to drive player behaviour in Sokoban and Pac-Man, and in the same course proposed several optimisation methods (Anthony, Polani, and Nehaniv 2014). Mohamed and Rezende focus primarily on optimisation, with a likely application in general game playing (Mohamed and Rezende 2015).

Method

Empowerment (Klyubin, Polani, and Nehaniv 2008) is defined between an agent's actuators and sensors. In a deterministic environment, it quantifies the options available to an agent in terms of availability and visibility. In a stochastic setting, this generalises to the potential influence of an agent's actions on its environment and to the extent to which the agent can perceive this influence afterwards.

Empowerment does not measure an agent's actual, but rather their potential influence on the environment. The empowerment maximisation hypothesis (Salge, Glackin, and

Polani 2014) suggests that an agent should, in the absence of any explicit goals, choose actions which likely lead to states with a higher influence on the environment, i.e. potentially more options. *Coupled empowerment maximisation* (CEM) is an extension of this principle to the multi-agent case.

In a shared game world, each agent can affect the others either explicitly or implicitly, which can be quantified by empowerment. Given the previous intuition, we suggest that increasing the empowerment of a goal-directed agent can be considered as supporting them in performing and achieving their tasks. We consequently hypothesise that equipping an NPC with an action selection policy which not only maximises their own- but also the player's empowerment leads to the emergence of companion-like behaviour.

Evaluation

We claim that maximising coupled empowerment realises companion-like behaviour. In our evaluation, we let the agent choose actions in an ad-hoc fashion according to the CEM policy. We analyse several scenarios, each illustrating one of the following companion duties:

1. *Player Integrity*: Ensure that the player can continue playing the game. Act against any force that would constrain these abilities.
2. *Support*: Support and do not hinder the player in achieving their goals. Maintain *operational proximity*, i.e. act towards states where you can support the player.
3. *Companion Integrity*: Secure your own existence and ability to act and support the player in the long term.

We did not define any goal-specific companion duties which could constrain the NPC's adaptivity. Instead, their goal directedness arises from its interaction with the player.

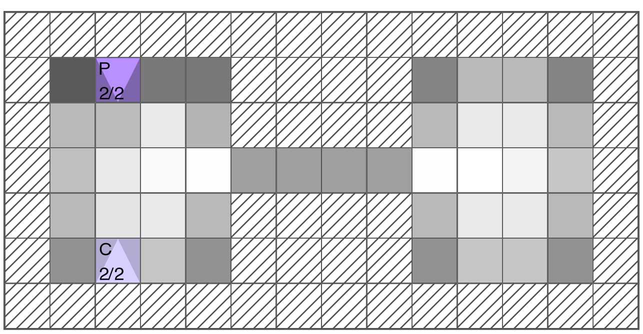

Figure 1: Snapshot of the evaluation environment

The experiments were set in a minimal *Dungeon-Crawler* game in which the player, supported by a companion NPC, has to navigate through rooms connected by corridors and defeat enemies in order to reach a goal state (Fig. 1). This game type traditionally relies on procedural content generation and elements of chance, and therefore poses interesting challenges to a general NPC policy. Classic examples such as *Nethack* illustrate how our testbed can be extended to introduce new challenges to the formalism. Dungeon crawlers are traditionally discrete in time and space and thus simplify the computation and analysis of agent behaviour.

Results

Our findings confirm that CEM establishes a sufficiently general frame for companion-like behaviour by inducing the player's goal into the companion's policy. Given our experimental evidence and the universality of the formalism, we hypothesise that the principle generalises to a wide range of game scenarios and genres.

If this proves correct, the flexibility and adaptivity of CEM may make NPCs fit for the most recent challenges in the games industry and academic research. It could allow industry to save efforts and reduce costs of manually authoring NPC behaviour, especially in games with a strong focus on procedurally generated content. Even if much scripting is required, our formalism can help in establishing a default mode of interaction with the player and other agents.

Acknowledgment

CG is funded by EPSRC grant [EP/L015846/1] (IGGI). CS received funding from the H2020-641321 socSMCs FET Proactive project. SC is funded by ERA Chair Project GRO (621403).

References

Anthony, T.; Polani, D.; and Nehaniv, C. L. 2014. General Self-Motivation and Strategy Identification: Case Studies based on Sokoban and Pac-Man. *Trans. CIAIG* 6(1):1–17.

Cerny, M. 2015. Sarah and Sally : Creating a Likeable and Competent AI Sidekick for a Videogame. In *Proc. AIIDE Workshop EXAG*.

Cook, M.; Colton, S.; and Gow, J. 2016. The ANGELINA Videogame Design System, Part I. *Trans. CIAIG* 1–12.

Forgette, J., and Katchabaw, M. 2014. Enabling Motivated Believable Agents With Reinforcement Learning. *Games Media Entertainment* 1–8.

Klyubin, A. S.; Polani, D.; and Nehaniv, C. L. 2008. Keep Your Options Open: An Information-Based Driving Principle for Sensorimotor Systems. *PloS one* 3(12):1–14.

Liapis, A.; Yannakakis, G. N.; and Togelius, J. 2014. Computational Game Creativity. In *Proc. 5th Conf. Computational Creativity*.

Merrick, K. 2008. Modeling Motivation for Adaptive Nonplayer Characters in Dynamic Computer Game Worlds. *CiE* 5(4):1.

Mohamed, S., and Rezende, D. 2015. Stochastic Variational Information Maximisation. *Proc. 29th Conf. NIPS* 1–9.

Ryan, R., and Deci, E. 2000. Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions. *Contemporary Educational Psychology* 25(1):54–67.

Salge, C.; Glackin, C.; and Polani, D. 2014. Empowerment – an Introduction. *Guided Self-Organization: Inception* 67–114.

Singh, S.; Barto, A.; and Chentanez, N. 2004. Intrinsically Motivated Reinforcement Learning. *Proc. 18th Conf. NIPS* 17:1281–1288.

Smith, G. 2014. The Future of Procedural Content Generation in Games. In *Proc. AIIDE Workshop EXAG*.