Entertainment Computing 16 (2016) 95–104

Contents lists available at ScienceDirect

Entertainment Computing

journal homepage: ees.elsevier.com/entcom

Evolving models of player decision making: Personas versus clones ☆,☆☆

Christoffer Holmgård [a,*], Antonios Liapis [b], Julian Togelius [c], Georgios N. Yannakakis [b]

[a] Center for Computer Game Research, IT University of Copenhagen, Denmark
[b] Institute of Digital Games, University of Malta, Malta
[c] Game Innovation Lab, New York University, New York, United States

ARTICLE INFO

Article history:
Received 10 March 2015
Revised 6 August 2015
Accepted 25 September 2015
Available online 24 October 2015

Keywords:
Decision making
Procedural content generation
Evolutionary computation
Player modeling

ABSTRACT

The current paper investigates multiple approaches to modeling human decision making styles for procedural play-testing. Building on decision and persona theory we evolve game playing agents representing human decision making styles. Three kinds of agents are evolved from the same representation: procedural personas, evolved from game designer expert knowledge, clones, evolved from observations of human play and aimed at general behavioral replication, and specialized agents, also evolved from observation, but aimed at determining the maximal behavioral replication ability of the representation. These three methods are then compared on their ability to represent individual human decision makers. Comparisons are conducted using three different proposed metrics that address the problem of matching decisions at the action, tactical, and strategic levels. Results indicate that a small gallery of personas evolved from designer intuitions can capture human decision making styles equally well as clones evolved from human play-traces for the testbed game MiniDungeons.

1. Introduction

This paper investigates how to create models of human decision making styles in games using generative, game-playing agents for procedural play-testing. It proposes an evolution based framework for representing player decision making in games and a simulation based method for evaluating human likeness of game playing agents at three different levels. The framework is applied in two different ways: evolving in a top-down manner from designer-driven intuitions and evolving in a bottom-up, data-driven manner from play-traces. The evaluation method is then used on both applications of the framework to evaluate their performances. Finally, a possibility for combining the two applications of the framework, allowing for hybrid top-down/bottom-up decision modeling through generative agents is suggested.

Generative, game-playing agents that represent and replicate human decision making may be useful in games for many purposes e.g. as believable stand-ins for human players or as benchmark rivals for players to surpass. The work presented here focuses on using game playing agents representing human decision making styles as stand-in players, supporting the traditional process of human play-testing.

Play-testing is typically an integral part of game development [1]. The complexity and cost of the play-testing depends on the kind of game under development, the stage in the games development process, and the objectives of the play-testing. At one extreme play-testing may be conducted by the game designer herself by simply imagining how players might interact with the game, a feature or a piece content. At the other extreme play-testing may be conducted under highly instrumented laboratory conditions or at a massive scale in the wild by telemetrically collecting data from players after the launch of the game [2].

In this paper, we suggest there may be an opportunity for methods using generative agents to support designers in situations where new content is being developed, but access to human play-testers is limited or impossible. For example, when a level designer is implementing a new level for a game or making changes to an existing one, these changes might not be large enough to mandate a full play-test with human players. Still, it might be useful for the level designer to observe how different kinds of players would interact with the level.

In situations like these, generative game playing agents based on models of human decision making might provide designers with surrogate play-traces to inform their design process and explore what parts of the game space players are likely to interact with and how, effectively delivering *procedural play-testing*.

☆ This paper has been recommended for acceptance by Nikitas Marinos Sgouros.
☆☆ This paper is part of the virtual special issue on "Selected Papers from the 2014 International Conference on Entertainment Computing", edited by Dr. Nikitas Marinos Sgouros and Dr. Matthias Rauterberg.

* Corresponding author.
E-mail address: holmgard@itu.dk (C. Holmgård).
URL: http://christoffer.holmgard.org (C. Holmgård).

http://dx.doi.org/10.1016/j.entcom.2015.09.002
1875-9521/

When agents sufficiently simulate a particular archetypal human decision making style we call agents *procedural personas*. Integrated with content creation tools, we envision that procedural personas will allow for mixed-initiative game design tools that yield immediate feedback during the design process, even if this feedback is not a completely accurate representation of how human players might play the game. Additionally, play-traces from procedural personas can be used as input for procedural content generation systems shaping the output in response to the generative player models [3].

In other words, agents that play like humans can help understand content by playing it as it is being created.

1.1. Research questions

An important question then arises with respect to which sources of information about player decision making styles are useful for constructing believable procedural personas that simulate human decision making with sufficient accuracy. Do we need some amount of low level behavioral data from actual players or can we derive the same information from the expert knowledge of a game designer?

A second question is how general we can make the resulting models. Can we ensure that they perform consistently on unseen content that either no play-traces were sampled from or that the game designer was not explicitly considering?

The work presented here addresses these questions by comparing two particular methods for realizing procedural personas, one drawing on designer expert knowledge and one using empirically gathered play test data, in order to evaluate which method produces the best models for generating synthetic play-test data.

1.2. Prior work

In previous work we have designed a simple turn-based, tile-based dungeon crawling game, *MiniDungeons*, which features monsters, treasures and potions in mazes [4]. 38 players played 10 levels of this game and we recorded their every action. Next, we analyzed the design of the game to extract a number of possible affordances which we translated into partially conflicting objectives that a player might seek to fulfill (e.g. kill all monsters, avoid danger or get to the exit quickly). Using these *affordances* we trained agents to play the game rationally for each objective. Both Q-learning [4] and evolutionary algorithms [4] were used to train high-performing agents; the evolved agents have the benefit that they generalize to levels they were not trained on in contrast to Q-learning agents which where unable to perform on levels they previously had not seen.

1.3. Metrics and methods for comparing agents to human players

The agents' behaviors were compared to play-traces of the human players through a metric we call the *action agreement ratio* (AAR) which compares agents and humans at the action level – comparing every action of the player and the agent and asking if the agent would pursue the same *next action* as the player. But is this really the right level of analysis for comparing players to agents? It could be argued that the microscopic level of comparing actions gives a biased view of how well an agent's behavior reproduces player behavior, and that it is more interesting to look at behavior not on the level of atomic decisions, but rather at the level of tactical or strategic decisions. Further, are we right to assume that players exhibit boundedly rational behavior given some set of objectives? It might be that with the same agent representation, we could train agents that reproduce player behavior better by using the actual play-traces as training data instead of focusing on player objectives. The current paper tries to answer these two questions.

Expanding on previous work [4], we propose two new play-trace comparison methods, *tactical agreement ratio* (TAR) and *strategic agreement ratio* (SAR) that, instead of asking whether an agent would perform the same singular action as the player in a given state, ask whether it would choose to pursue the same *next affordance* or the same *overall outcome*, respectively.

We also train a second class of agents to behave as similarly as possible to human players on *unseen* levels using play-traces as objectives, again evaluated on the three levels of comparison: the action level, the tactical level, and the strategic level. We call such agents *clones*.

Finally, we train a third class of agents to behave as similarly as possible to human players on *previously seen* levels in order to explore the maximal performance of our chosen representation. We call such agents *specialized agents* as they are likely to be the closest fit of the representation to an individual play-trace, but are trained for just one particular level.

1.4. Modeling bounded rationality

Grounded in contemporary decision science, this paper has two central assumptions about human players' decision making: The first is that players' decisions are guided by their expected *utility* for a given decision; i.e. the amount of experienced value they expect to derive from the consequences of a decision.

The second is that human players exhibit *bounded rationality* i.e. players allocate limited amounts of cognitive resources to decisions in games either due to innate limitations or because they only apply part of their cognitive capacity to the decision due to conscious or subliminal reasons. Decision making style in games thus depends not only on preferences in outcomes, but also the resources the player is willing or able to allocate to the decision making task. Our approach to simulating a decision maker in the form of a generative agent is to represent these two characteristics, the player's rational utility function and the player's cognitive bounds, in the implementation of the agent.

In the following we outline the relations between persona theory, decision theory, player modeling, and the resulting concept of procedural personas. We briefly describe our testbed game, MiniDungeons, and the methods we used to create game playing personas and clones, before we present the results from comparing the resulting agents to the human players.

2. Related work

In this section, we review decision theory, the concept of personas as applied to (digital) games, player modeling, and the relations between the three areas in this study.

2.1. Decision theory and utility

The personas used for expressing designer notions of archetypal player behavior in MiniDungeons are structured around the central concepts of decision theory. Decision theory states that whenever a human makes a *rational decision* in a given situation, the decision is a result of an attempt to optimize the expected *utility* [5]. Utility describes any positive outcome for the decision maker and is fundamentally assumed to be idiosyncratic. This means that in principle no definite assumptions can be made about what can provide utility to the decision maker. The problem is further complicated by the fact that the effort a decision maker directs toward attaining maximum utility from a decision can be contingent on the expected utility itself. For problems that are expected to provide low utility