Game Playing

(Komal, Karuna Hazrati , Priyanka Mittal)

Abstract: Computer games are an increasingly popular application for Artificial Intelligence (AI) research, and conversely AI is an increasingly popular selling point for commercial games. Although games are typically associated with entertainment, there are many “serious” applications of gaming, including military, corporate, and advertising applications. There are also so-called “humane” gaming applications for medical training, educational games, and games that reflect social consciousness or advocate for a cause. Game AI is the effort of going beyond scripted interactions, however complex, into the arena of truly interactive systems that are responsive, adaptive, and intelligent. Such systems learn about the player(s) during game play, adapt their own behaviours beyond the pre-programmed set provided by the game author, and interactively develop and provide a richer experience to the player(s).

INTRODUCTION:

Computer games have been classified as the “Human-level AI’s Killer Application” (Laird & van Lent 2000). State- of-the-art computer games recreate real-life environments with a surprising level of detail. These environments are usually populated with many characters (allies or enemies) that require human-level intelligence and exhibit believable behaviours. However, even though there have been enormous advances in computer graphics, animation and audio for games, most of the games contain very basic artificial intelligence (AI) techniques. AI programming wisdom series (Rabin 2002; 2004) provides a good overview of current state of the art AI techniques used in the game industry. As a result of simplistic approaches for AI, the whole atmosphere created by the game can be broken when the game and characters situated within it behave in a non-believable manner. On the other hand, creating richer experiences requires a great deal of engineering effort on the part of game developers.

The first part of this paper will survey the research landscape of general game playing. We will see that a surprisingly broad range of classic AI fields have a role to play in general game playing. In each case we will show how existing approaches, methods and systems are contributing to the foundations for general game playing, to the improvement of the quality of existing general game-playing systems, and to the development of new methodologies. But even if your interest as an AI researcher does not lie in general game-playing systems themselves, they can be used as a non-trivial application for a broad range of more theoretically motivated AI research. We will also report on current research trends, identify some of the most pressing open questions, and look at the possibilities to gradually broadening today’s concept of general game playing to involve even more aspects of AI.

In the second part of this paper, we will give an overview of how general game playing has entered AI education, either in form of an advanced AI graduate course, with lectures and tutorials but where the special focus lies on practical work; or as part of a general introductory course to AI. Examples for the former can be found in the curricula at Stanford University, Reykjavik´ University,ık Bremen University and TU Dresden, where it was held by the author for four consecutive years starting in winter 2006/7. Examples for the latter include the general introduction to AI for undergraduate students at the University of New South Wales to which the author contributed in spring 2011. We will show why general game playing provides an excellent angle for teaching a variety of basic AI methods that also is a great motivator for students to design and implement their own AI systems. Our overview will include a survey of freely available teaching aides including slides, tutorial questions and programming tools, for the benefit of potential instructors.

The Research Landscape of General Game Playing

An outstanding characteristic of general game playing as an AI research topic is to involve a broad range of sub-disciplines with a focus on symbolic AI (as opposed to, say, RoboCup or the DARPA Grand Challenge for driverless vehicles). For this reason, general game playing has all the potential to become a rich source for interesting research problems in many different areas. As we will survey the research landscape, we will encounter several cases in which general game playing has been successfully used as an attractive—and challenging—application to demonstrate the viability of existing theories, methods, and systems. We will also see examples where the concept of general game playing has generated new research problems and solutions. Most importantly, it will become clear that it is not at all necessary to actually build a full-fledged, competitive player in order to make an original and significant contribution to this research challenge. Yet another characteristic of general game playing research is too often concern the combination and integration of two or more theories and methods, which naturally leads to collaborations involving different AI sub-disciplines.

In the relatively short time span since the first AAAI competition in 2005, at least four traditional AI disciplines have proved to be core aspects of research in general game playing:

1. Knowledge Representation
2. Search
3. Planning
4. Learning

In the remainder of this section, we will discuss in turn the role that each of these areas plays for general game playing: what interesting research problems they give rise to, which methods have been successfully applied, and what challenges lie ahead.

Requirements for Game AI

In previous work, Laird and van Lent (2000) analysed different game genres, and the AI challenges that each presents. In their report, they considered the following types of games: action, role playing, adventure, strategy games, god games, and individual and team sports games. In addition to those genres, we would like to consider two additional categories, namely, interactive drama (Mateas & Stern 2003) and educational games (Rieber 1996). Interactive dramas have a strong plot behind them that the author wants to communicate to the player, but where the player may have a strong influence on the plot. A key difference with the classical “adventure” genre is that adventures have a scripted plot, while interactive dramas are more open-ended and adapt to the player interaction as the story unfolds. Educational games have an additional rhetorical goal of teaching some particular content to the player.

By analysing the range of possible applications of computer game AI to different applications and game genres, we identify two different levels at which AI can be applied: 1) individual characters AI, with the goal of producing more intelligent or believable behaviours, and 2) a global AI that watches over the game or game-player interaction, influencing the directions that the game is taking. Thus, we can talk about character-level AI and game-level AI (the second being referred in some papers as the Drama Manager (Nelsonet al. 2006) or as the Director (Magerko et al. 2004)).

Different applications and game genres require a different mix of these two kind of AIs. For instance, real-time strategy games rely mainly on a game-level AI that controls all the units, while the individual unit behaviors can be scripted.

Role playing games, on the other hand, require believable character-level AI to provide an interesting player experience. Interactive dramas requires a mix of both kinds of AI: individual characters that are believable and a drama manager that leads the plot by guiding the individual character to take actions that can make the drama advance. Educational applications of gaming also require a game-level AI,similar to the drama manager, which monitors the interaction of the game as it unfolds, easing or complicating the tasks according to the learner’s expertise level, thereby making sure that educational purpose of the game is being met.

Each game genre presents particular requirements for character level and game level AI. For instance, god games usually require the game-level AI to solve resource allocation problems and solve long-term strategy problems, while interactive drama requires the game-level AI to adapt the story according to the player interactions in a way that it is more appealing to the player (thus, the latter requires user modelling and story planning). Moreover, adventures, interactive dramas and other genres with embodied characters usually require believability and natural language generation.

In the following section, we summarize a list of interesting challenges that computer games pose in general to the
AI community.

Challenges in Computer Game AI

Let us briefly describe some of the main issues that arise when developing artificial intelligence for computer games. This list is not exhaustive, but is intended to give a favour of the kind of problems that real computer games pose to the
AI community.

• Complex decision spaces: Most state-of-the-art com putter games involve complex strategic (real-time strategy games) or believable behaviours (interactive dramas).Both kind of behaviours share the characteristic of having huge decision spaces, and thus traditional search-based AI techniques cannot be applied. Learning techniques or higher level representations are required to deal with such complex games. Traditionally, computer games use handcrafted strategies coded by the game developers, but these tend to be repetitive, and players easily find holes and exploit them.

• Knowledge engineering: Even assuming that strategies or behaviours are handcrafted, authoring these behaviour sets in a game requires a huge human engineering effort. Game developers have to encode all the knowledge they have about a domain (either to achieve a strategic behaviour or a believable human behaviour) in some sort of behaviour language.

• Authoring support: Hand crafted behaviours are, ultimately, software code in a complex programming language, prone to human errors. The behaviour errors could be in the form of program “bugs” or not achieving the desired result. Tools are needed to support story authors, who are typically not artificial intelligence experts, to author behaviours in a computer programming language

• Unanticipated situations: It is not feasible to anticipate all possible situations and player strategies that can encountered during game play. This makes it difficult to craft believable behaviours that react in an appropriate manner to these unforeseen circumstances and player actions

• User-specific adaptation: Different players may enjoy different strategies to fight against (in the case of real time strategy games), or different styles of storytelling (in the case of interactive dramas), different types of story development, different kinds of character behaviours and interactions, or different educational problems. As game designers begin to include user modelling capabilities, the AI strategy and behaviour must, in turn, be adaptable based on the user model.

• Replay ability and variability: A player might get bored of seeing the same strategies and behaviours again and again.

Although simple variability can be achieved through stochastic selection of behaviours or strategies from a large repository, this increases the authoring burden. Furthermore, random selection begs the question of true interestingness.

• Rhetorical objectives: It is possible, even likely, that human-engineered behaviours or strategies do not achieve the game’s objectives adequately, especially in realistic, scaled-up domains or applications. These objectives could range from entertainment to education, training, etc. Thus, the game has to realize that the objectives are not being met on a per-use basis, and adapt accordingly. For example, a particular user may be getting bored, or no learning the intended lesson.

Behavior Modification for Believable Characters AI

In interactive games, embodied characters typically have their own personalities, affecting the way they act in the game. Authors usually create such characters by writing behaviours or scripts that describe the characters' reaction to all imaginable circumstances within the game world. This approach of authoring characters presents several difficulties.

First, when authoring a character's behaviour set, it is hard to imagine and plan for all possible scenarios it might encounter. Given the rich, dynamic nature of game worlds, this can require extensive programming effort. Second, over long game sessions, a character's static behavioural repertoire may result in repetitive behaviour. Such repetition harms the believability of the characters. Third, when behaviours fail to achieve their desired purpose, characters are unable to identify such failures and will continue to exhibit them. Ideally, we want a self-adapting behaviour set for characters, allowing characters to autonomously exhibit their author-specified personalities in new and unforeseen circumstances, and relieving authors of the burden of writing behaviours for every possible situation.

To address these issues, we have developed an approach in which agents keep track of the status of their executing behaviours, infer from their execution trace what might be wrong, and perform appropriate revisions to their behaviours. This approach to runtime behaviour transformation enables characters to autonomously adapt during execution to changing game situations, taking a first step towards automatic generation of behaviour that maintains desired personality characteristics. Our approach is related to plan revision (Cushing & Kambhampati 2005), with the added complexity that failure detection and behaviour modification must be performed during execution, enabling the game to continue seamlessly from the player's perspective.

Case-Based Planning for Strategy Games

AI techniques have been successfully applied to several computer games such as checkers, chess or Othello (Schaeffer 2001). However, in many computer games traditional AI techniques fail to play at a human level because of the characteristics of the vast search spaces this games require. For that reason, game developers need to invest significant effort in hand-coding specific strategies that play at a reasonable level for each new game.

For instance, previous research has shown that real-time strategy games (RTS) such as Wargus (a clone of the popular commercial game Warcraft II) have huge decision spaces (Aha, Molineaux, & Ponsen 2005; Buro 2003). In this section we present an architecture that uses case-based planning (Hammond 1990) to deal with such complex games.

In previous work, we have applied case-based reasoning (CBR) to RTS games (Sharma et al. 2007a). The idea there was to define a set of high level actions, and let a CBR system learn when each should be applied. In this section, we discuss a different approach that addresses the complexity of this domain by extract behavioural knowledge from expert demonstrations (i.e., an expert plays the game and our system observes). Then, at performance time, a case-based planning engine retrieves suitable behaviours observed from the expert and adapts them to the current game state. Adaptation is required since the game state may be different from the one in which the behaviour was originally demonstrated, with, for example, a new map, different units, or a different objective.

One of the main contributions of this approach is that it enables the game developers to specify the AI behaviour just by demonstration, i.e., instead of having to code the behaviour using a programming language, the behaviour can be specified simply by demonstrating it to the system. If the system shows an incorrect behaviour in any particular situation, instead of having to find the bug in the program and fix it, the game developers can simply demonstrate the correct action in the particular situation. The system will then incorporate that information in its case base, thereby improve its behaviours in the future.

Conclusion

General game playing is an exciting, still young but on the verge of maturing topic, which touches upon a broad range of aspects of artificial intelligence. In this paper we surveyed the research landscape of general game playing in an attempt to show its many facets and that it provides a rich source of interesting and challenging problems for many an AI researcher. We also showed that general game playing provides a unique approach to teaching a number of different topics in AI. Students who got exposed to the idea of a general game-playing AI system have repeatedly described it as "cool," and the author is inclined to agree.

References

Aamodt, A., and Plaza, E. 1994. Case-based reasoning: Foundational issues, methodological variations, and sys- tem approaches. Artificial Intelligence Communications 7(1):39–59.

Aha, D. W., and Wilson, D. 2005. Computer gaming and simulation environments: Papers from the ICCBR workshop. Technical Report, DePaul University, Chicago, IL.

Aha, D.; Molineaux, M.; and Ponsen, M. 2005. Learning to win: Case-based plan selection in a real-time strategy game. In ICCBR'2005, number 3620 in LNCS, 5–20. Springer-Verlag.

Aha, D. W.; Munoz-Avila, H.; and van Lent, M. 2005. Reasoning, representation, and learning in computer games:

Papers from the IJCAI workshop. Technical Report AIC-05-127, Washington DC: Naval Research Laboratory,

Navy Center for Applied Research in Artificial Intelligence. Bartneck, C. 2002. Integrating the occ model of emotions in embodied characters. In Proceedings of the Workshop on Virtual Conversational Characters: Applications,

Methods, and Research Challenges Buro, M. 2003. Real-time strategy games: A new AI research challenge. In IJCAI 2003, 1534–1535. Morgan Kaufmann.

Cheng, D. C., and Thawonmas, R. 2004. Case-based plan recognition for real-time strategy games. In Proceedings of the 5th Game-On International Conference, 36 – 40.

Cushing, W., and Kambhampati, S. 2005. Replanning: A new perspective. In Proceedings of ICAPS.

Epic Games.2004.Unreal tournament 2004, http://www.unrealtournament.com.

Gentry, M. S. 1998. Anchorhead. Available online at http://www.wurb.com/if/game/17.html.

Hammond, K. F. 1990. Case based planning: A framework for planning from experience. Cognitive

Science
14(3):385–443.
Isbister, K., and Doyle, P. 2001. Toward the holodeck: Integrating graphics, sound, character, and story. In Proceedings of the Fifth International Conference on Autonomous Agents, 409 – 416.Isbister, K., and Doyle, P. 2003. Web guide agents: Narrative context with character. In Mateas, M., and Sengers, P.,
eds., Narrative Intelligence, 229–243. Kolodner, J. 1993. Case-Based Reasoning. San Mateo,
CA: Morgan Kaufmann.
Laird, J. E., and van Lent, M. 2000. Human-level AI's killer application: Interactive computer games. In AAAI 2000, 1171–1178.
Lester, J., and Stone, B. 1997. Increasing believability in animated pedagogical agents. In First International Conference on Autonomous Agents, 16–21.
Lester, J., and Stone, B. 1998. Integrating reactive and scripted behaviors in a life-like presentation agent. In Proceedings of the Second International Conference on Autonomous Agents, 261–268.
Loyall, B. 1997. Believable Agents: Building Interactive Personalities. Ph.D. Dissertation, Carnagie Mellon University.
Magerko, B.; Laird, J.; Assanie, M.; Kerfoot, A.; and Stokes, D. 2004. AI characters and directors for interactive computer games. In Proceedings of the 2004 Innovative
Applications of Artificial Intelligence Confercence.
Mateas, M., and Stern, A. 2002. A behavior language for story-based believable agents. IEEE intelligent systems and their applications 17(4):39–47.
Mateas, M., and Stern, A. 2003. Integrating plot, character, and natural language processing in the interactive drama facade. In Proceedings of the 1st International Conference on Technologies for Interactive Digical Storytelling and Entertainment.
Nelson, M.; Roberts, D.; Isbell, C.; and Mateas, M. 2006. Reinforcement learning for declarative optimization-based drama management. In AAMAS 2006, 775–782.
Rabin, S. 2002. AI Game Programming Wisdom. Charles River Media.

.