

# On the Usefulness of Interactive Computer Game Logs for Agent Modelling

Matthew Sheehan and Ian Watson

Dept. of Computer Science, University of Auckland  
Private Bag 92019, Auckland, New Zealand.  
mshe065@ec.auckland.ac.nz, ian@cs.auckland.ac.nz

**Abstract.** Interactive computer game logs show the potential for use as replacement for time-consuming supervisory learning processes for embodied, situated agents. However, due to the inherent nature of the data in the logs themselves, for the time being this promise cannot be fulfilled. An unsuccessful attempt to use largely rule-based data mining processes to learn behaviours from game logs led to finding the inherently top-down nature of such processes was fundamentally at odds with the unsupervised bottom-up learning requirement of the problem. Innate issues with game logs toward the goal of unsupervised agent learning are discussed. Possible approaches to the problem subsuming successful applications of various methods in narrower fields are presented for both symbolic and sub-symbolic advocates.

## 1 Introduction

This paper primarily addresses the use of interactive computer game logs for embodied, task-generalised, goal oriented agent learning. In particular, in our view, is the potential of game logs to supplement or even replace time-consuming supervisory learning methods. While our own research was largely unsuccessful in achieving this goal, what was learnt is considered useful for future researchers and is presented here.

A computer game log can be recognised as a file that is used to store captured (usually live) game data, often in a sequential or time-stamped manner. A game log file can be human text readable, or require a parsing application (for an example see [1]). Whether it is the sequential moves taken in a turn based game or a FPS ‘demo’ file that captures all server-client network traffic down to timeframes of a tenth of a second, the defining and encompassing attribute of the term’s use in this paper is the storage of game activity in an abridged manner.

## 2 Motivation and Related Work

The inherent opportunities and motivation for artificial intelligence research in computer game genres where the player (and or opponents) are situated within

the game environment is well documented by Laird and van Lent [2]. The First Person Shooter (FPS) genre in particular has been popular with researchers for its complexities and availability of software. A number of internet sites house large, freely accessible game log repositories built using game logs uploaded by game enthusiasts.

Some of the best known uses of game logs to model the behaviour of simulated agents is the Robocup Coach and Simulated League Competitions (e.g. [3], [4]). Robocup examples can be differentiated to the current research by the general level of complexity of data within the game logs themselves, due to the relatively simple nature of the agents, game rules and two-dimensional environment the Robocup games are situated in.

More pertinently, Gorman et al [5] have used a three dimensional, FPS environment to implement a two tier (reactive/strategic) situated agent using Quake II game demo files. Tactical tier implementations were investigated by the two research groups separately but were never integrated.

The majority of other implementations of situated/embodied agents in FPS environments use a largely symbolic, top-down approach [6, 7]. However, none of the symbolic examples listed above use game logs as a means of learning new rules. The closest examples are perhaps the annotated expert decisions that have been used by Könik and Laird [8] in an adventure game setting and the expert traces used by Nejati et al [9] to learn hierarchical tasks, currently a feature of Choi et al's FPS agent implementation [7].

### 3 Approach

Quake II demos, though not able to be directly read using a text-editor, given the correct map the in-game console are able to be used to playback the entire game. With the API, it is possible to extract game information/attributes from the world state at each tenth of a second frame. Because of the inherent rule-based nature of games and previous successful rule-based agent implementations, a largely rule-based data mining approach was taken, forming flat files from extracted information and running them on data mining software, using both algorithmic and brute force techniques. However, it was found that this approach was not only largely unsuccessful, but more interestingly, perhaps deeply flawed, discussed in the next section.

### 4 The Usefulness of Interactive Game Logs in Agent Modeling

The symbol grounding problem is a key issue for symbolic approaches to agent modelling with FPS demo game logs in particular due to the sub-symbolic (i.e. quantitative) nature of the data. Using an example from our own research in attempting to find aiming rules, one time-frame's captured world state might

contain the positional coordinates and angles of player and opponent but no labelled links to what these positions and angles represent. Unlike [8] it is infeasible to elucidate every action in a FPS setting.

In attempting to model item-collection behaviour from the game logs it was found that some actions are hierarchical and interruptible. For example, should an agent be walking towards an ammunition pack when an opponent appears, an action to protect itself against attack should (generally) override the collection of items.

Though the quality of behavioural data seems evident when watched as a video, as current interactive game log attributes are recorded in a global positioning system, these behaviours cannot be explicitly extracted. For instance, aiming information that is agent-relative for learning general rather than situation specific behaviours can be calculated from these attributes, but it requires the use of complex trigonometric processes to be applied to the data (perspective projection).

Deriving basic lessons by watching experts means the ability to break down and slow down complex maneuvers into their constituent smaller manoeuvres, an exercise that effectively requires expert knowledge to begin with. An expert manoeuvre has implicit within it tactical weapon control (which weapon to use, when), modified circle/strafing behaviour, map knowledge and ambushing behaviour, with each constituent part seamlessly blended into the others parts, unlikely to be repeated exactly again, even in the same part of the map.

The game environment is three dimensional, but environment items and agent movements and decisions make the problem space of a much higher dimensionality. A means of filtering and lowering the dimensionality of this data (addressed in the next section) could help mitigate this issue.

## 5 Possible Approaches

This section shows possible approaches to implementing a means of using interactive game logs for embodied agent learning by suggesting general means of approach coupled with techniques that have overcome some of the inherent difficulties outlined above in more specific domains.

### 5.1 Symbolic/Rule-Based Approach

The largest difficulty in learning declarative rules automatically from interactive game logs is in the sub-symbolic and procedural nature of the game log data. One answer (not without its own difficulties, admittedly) to the problem of the symbol grounding problem would be to make the contents of the game logs declarative, rather than procedural. For this to occur, it would be necessary that the game itself to work in declarative terms or at least have a declarative layer built into the game engine (something akin to the ‘Symbolic Environmental Representation’ in [8]). From this game layer symbolic states and transitions

could be collected for logs, to be used by symbolic cognitive architectures or integrated for implementation with artificial agent game interfaces such as Berndt and Watson’s [10] Open AI Standard Interface Specification (OASIS).

This layer might be built using a standardized game ontology perhaps using qualitative reasoning techniques, with particular regard to qualitative spatial reasoning [11–13] (using abstractions to model physical space) and qualitative physics [12] (modelling the physics of a world using rules rather than exact quantitative methods). Zagal et al [14] have already begun to outline a game ontological language for game analysis that could possibly be extended to the standardised qualitative programming of some types of game, while Zhou and Ting [15] have used qualitative methods in a three-dimensional FPS environment to model moveable objects. While these two examples, in themselves, do not come close to a complete solution to modelling a physical three dimensional simulation symbolically, they are indicative of approaches that may be successful in the future.

Nejati et al’s [9] hierarchical learning approach is a beginning to confronting hierarchical learning processes. However, like the annotated expert decisions in [8], their top down approach means that it requires positive (i.e. successful) examples with a set goal and start point are required for each behaviour to be learnt, not possible from recorded game demos without large scale trainer intervention.

It would seem that using data mining techniques would be useful, especially in regard to the large amounts of data that are present, plus current techniques providing means to bring quantitative, inherently time-based data to a symbolic level using knowledge based temporal abstraction [16]. Unfortunately, rule learning from even qualitative data time series seems to still require *top-down* methods of knowing the patterns and relationships you are looking for, before being able to find the sort of relationships between each sort of pattern [17]. Our own discouraging results in using data mining techniques to find rules in demo log data emphasise the inherent fundamental discord between the necessity of a bottom-up learning approach and one of the golden rules of data mining being already knowing what you are looking for [18]. Even when using brute force techniques, rules found must be recognised as being valid.

It is believed a bottom up approach could work better, bottom up approaches also useful in not requiring agent sophistication before learning begins.

## 5.2 Sub-Symbolic Approaches

A sub-symbolic approach to modelling an agent would seem to be the best way to approach the problem given the nature of the data in the current game logs. Perhaps more efforts have not been made in this area by sub-symbolic and nouvelle AI advocates due to a generally held belief championed by Brooks [19] condemning simulated environments.

It is interesting to note that after beginning with completely reactive, strictly bottom up approaches using neural nets and self organising maps to model

agents' movement [20], later strategic layered approaches took on more top-down attributes including setting nodal positional points (formed by clustering the entire set of player positions) which the agent would travel between and the making of goals in the form of item pickup points [5].

To be forced to create a sub-cognitive means to deal with data that is much closer to the kind of data a real-world agent would face could be seen as a effort to model not only a real world environment, but the kind of sensory inputs that would be received also. However, a large amount of preprocessing of the data is still required, even if specific logs are created (instead of using expert logs from the online repositories) before it can be used for learning. While it could be noted that the important part of the logs is the behavioural data they contain, not the form they are in, the ability to 'cheat' (whether conciously or unconciously) to produce circumstances favourable to the agent's learning circumstances is possible. Means to extract important features or bring levels of dimensionality down with large amounts of data include self-organised maps, principle component analysis (both used in methods by Thureau et al [21]) and independent component analysis. The success of Floyd et al [4] in largely automating the learning process, even in a less complex environment and producing effectively reactive agents, shows CBR methods could provide a means to sidestep this process.

## 6 Conclusions

Interactive game logs show promise as a source for expert behaviour to supplement or even replace time-consuming supervised expert tracing schemes. However, there are still a number of issues that need to be addressed before their full potential can be utilised. Both the content and means to extract that content automatically has a large bearing on their usefulness: the larger the supervisory aspect required by the method of agent training, the less useful they become, as it is easier to produce single logs and elucidate every action separately in (basic) movements for agents to learn than to extract behaviours from the thousands of game logs created by expert players by hand. Recommendations to increase this usefulness can be generalised into two parts: firstly, the improvement of means of information representation in the logs themselves; and secondly, the choice of methods by the researcher that lead to largely automatic information extraction.

## References

1. Gorman, B., Fredriksson, M., Humphrys, M.: Qase - an integrated api for imitation and general ai research in commercial computer games. In: Proceedings of the IEEE 7th International Conference on Computer Games, CGAMES (2005) 207–214
2. Laird, J.E., van Lent, M.: Human-level ais killer application: Interactive computer games. In: Proceedings of the AAAI Fall Symposium, AAAI Press (2000) 80–97
3. Kuhlmann, G., Knox, W.B., Stone, P.: Know thine enemy: A champion robocup coach agent. In: In Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI 06), AAAI (2006) 146368

4. Floyd, M.W., Esfandiari, B., Lam, K.: A case based reasoning approach to imitating robocup players. In: Proceedings of the Twenty-First International FLAIRS Conference, AAAI Press (2008)
5. Gorman, B., Thureau, C., Bauckhage, C., Humphrys, M.: Bayesian imitation of human behavior in interactive computer games. In: Proceedings of the 18th International Conference on Pattern Recognition. Volume 1., ICPR (2006) 1244–1247
6. Laird, J.E.: It knows what you're going to do: adding anticipation to a quakebot. In: AGENTS '01: Proceedings of the fifth international conference on Autonomous agents, New York, NY, USA, ACM (2001) 385–392
7. Choi, D., Konik, T., Nejati, N., Park, C., Langley, P.: A believable agent for first-person perspective games. In: Proceedings of the Third Artificial Intelligence and Interactive Digital Entertainment International Conference (AIIDE 2007), AAAI Press (2007)
8. Konik, T., Laird, J.: Learning goal hierarchies from structured observations and expert annotations. In: Proceedings of the Fourteenth International Conference on Inductive Logic Programming, Springer (2004) 198–215
9. Nejati, N., Langley, P., Konik, T.: Learning hierarchical task networks by observation. In: Proceedings of the 23 rd International Conference on Machine Learning. Volume 148., ACM Press (2006) 665–672
10. Berndt, C.N., Watson, I., Guesgen, H.: Oasis: An open ai standard interface specification to support reasoning, representation and learning in computer games. In Aha, D., Muñoz-Avila, H., van Lent, M., eds.: Proceedings of the 2005 IJCAI Workshop on Reasoning, Representation, and Learning in Computer Games. (2005) 19–24
11. Freksa, C.: Qualitative spatial reasoning. In Mark, D.M., Frank, A., eds.: Cognitive and Linguistic Aspects of Geographic Space. Kluwer Academic Publishers, Dordrecht (1991) 361–372
12. Forbus, K.D.: Qualitative reasoning. In Tucker, J.A., ed.: The Computer Science and Engineering Handbook. CRC Press, Boca Raton, FL (1997) 715–733
13. Forbus, K.D., Mahoney, J.V., Dill, K.: How qualitative spatial reasoning can improve strategy game ais. *IEEE Intelligent Systems* **17** (2002) 25–30
14. Zagal, J., Mateas, M., Fernandez-Vara, C., Hochhalter, B., Lichti, N.: Towards an ontological language for game analysis. In: Proceedings of the Digital Interactive Games Research Association Conference (DiGRA 2005). (2005)
15. Zhou, S., Ting, S.P.: Qualitative physics for movable objects in mout. In: ANSS '06: Proceedings of the 39th annual Symposium on Simulation, Washington, DC, USA, IEEE Computer Society (2006) 320–325
16. Shahar, Y.: A framework for knowledge-based temporal abstraction. *Artificial Intelligence* **90** (1997) 79–133
17. Sacchi, L., Larizza, C., Combi, C., Bellazzi, R.: Data mining with temporal abstractions: learning rules from time series. *Data Mining Knowledge Discovery* **15** (2007) 217–247
18. Pyle, D.: *Data Preparation for Data Mining*. Morgan Kaufmann, San Francisco, CA (1999)
19. Brooks, R.A.: Elephants don't play chess. *Robotics and Autonomous Systems* **6** (1990) 3–15
20. Thureau, C., Bauckhage, C., Sagerer, G.: Combining self organizing maps and multilayer perceptrons to learn bot-behavior for a commercial game. In: Proceedings of GAME-ON03. (2003) 119–123
21. Thureau, C., Paczian, T., Bauckhage, C.: Is bayesian imitation learning the route to believable gamebots? In: GAME-ON North America. (2005) 3–9