

LUDOCORE: A Logical Game Engine for Modeling Videogames

Paper #477

Abstract

LUDOCORE is a logical “game engine”, linking game rules as reasoned about by game designers to the formal logic used by automated reasoning tools in AI. A key challenge in designing this bridge is engineering a concise, safe, and flexible representation that is compatible with the semantics of the games that logical models created with our engine intend to represent.

Building on the event calculus, a formalism for reasoning about state and events over time, and a set of common structures and idioms used in modeling games, we present a tool that is capable of generating gameplay traces that illustrate the game’s dynamic behavior. It supports incremental modeling of player and non-player entities in the game world, modification of game rules without extensive non-local changes, and exploratory temporal and structural queries. In addition, its logical models can support play as real-time, graphical games with minimal user-interface description.

Introduction

While the term *videogame* brings to mind graphics, sounds, and story worlds, at the core of every game is a formal rule system. We are interested in declaratively modeling these games so that the emergent properties of their rule systems can be understood. The tools of logicist AI hold promise for bringing such properties to light, but have not been used for design or analysis of videogames, in part because of a mismatch between how designers think of their rule systems and the way logical specifications are usually written. We propose the use of a logical “game engine” to ease and accelerate the modeling of game worlds in formal logic. Our engine, LUDOCORE, provides a set of primitives and abstractions that link game-level concepts to the first-order logic understood by AI reasoning tools. In addition, our engine supports generating gameplay traces, integrated player modeling, and both temporal and structural queries. Finally, games modeled in LUDOCORE can be directly read in, without compilation or translation, to a custom game engine (OMITTED 2009a), allowing them to be played as real-time, graphical games.

An existing way to represent purely abstract games in logic is GDL, the game description language (Love et al.

2006). GDL was designed to specify games for the general gameplaying competition (Genesereth 2005), in which computers competitively play unfamiliar, turn-based games. Thus, the language itself was designed primarily to write infrequently-edited specifications to be read and understood by computer players as state-transition systems. In contrast, when modeling videogames, designers are much more concerned with the ease with which they can incrementally build up a game’s rule system through a series of experimental tweaks (Fullerton 2008, p. 14). LUDOCORE supports this kind of modeling, at least for technically savvy designer-programmers.

Our iterative-design motivation also contrasts with the goals of formal verification. While verification attempts to prove that a set of desired properties hold, it is not clear during a game’s prototyping phase which properties are even desirable. In fact, coming to understand what properties a game has, and which ones would be desirable, is the main challenge.

Our bridge from games to formal logic is based on the event calculus, a logical formalism commonly associated with commonsense reasoning. In previous research (OMITTED 2008), we proposed that the event calculus is an attractive basis on which to build informative, queryable prototypes. Since then, in building games with the event calculus, we encountered various design patterns and idioms, common to all games; it is this experience which prompted the design of LUDOCORE.

Building Games on the Event Calculus

Standard game-design texts discuss a game’s mechanics in terms of objects with dynamic properties (the state of the game), and triggered behaviors (events and conditional state updates) (Adams 2009, p. 295; Fullerton 2008, p. 112). While state and events can be modeled with several formalisms, or even ad-hoc logical encodings, a common challenge is the well known frame problem, commonly solved by adding frame axioms. In a language like GDL, frame axioms explicitly declare when state doesn’t change from turn to turn. An alternate proposed solution is based on the commonsense law of inertia (Shanahan 1997b), which reads that states retain their values until an event changes them. This mirrors the usual game-programming assumption that variables stay set to the same value until changed by an active

behavior.

Given our desire to model state and events, and avoid frame axioms that are tedious to maintain, the event calculus (EC) is a natural choice. Additionally, there is a large body of literature discussing the use of EC as a practical knowledge representation, *e.g.* Mueller’s (2006) book on applying it to commonsense reasoning problems involving interaction of objects over time, for which inertial state is the common case.

The discrete event calculus (Mueller 2004) is based on fluents (predicates whose truth values vary over time) and events, which happen at particular integer-valued timepoints and can change the truth values of fluents. Its key predicates are: *happens*(E, T), which says that an event happens at a timepoint; *holds_at*(F, T), which says that a fluent is true at a timepoint; and *initiates*(E, F, T) and *terminates*(E, F, T), which map event occurrences to changes in fluent truth values. In addition to state that does not change without cause (inertial fluents), the circumscription used in the event calculus implies that events have no effects besides those that can be derived. Together, these two kinds of default reasoning give the event calculus elaboration tolerance (McCarthy 1998), the ability to modify a knowledge representation without re-engineering it, because new assertions override defaults.

To realize EC in a computational setting, we use answer-set programming (ASP), an approach proposed in unpublished notes accompanying Mueller’s (2006) book,¹ which Kim, Lee, and Palla (2009) proved preserves the expected EC semantics. Answer sets are sets of literals that represent acceptable beliefs in an abstract world (Lifschitz 2008). When applied to programs in the discrete event calculus, these answer sets amount to assertions about what was true and what happened at each timepoint. Further, for games modeled in EC, the narrative of events amounts to what we call a gameplay trace. While there are more direct ways to extract a single trace from a game, ASP does not forward-simulate a game, but rather reasons abstractly about the space of possible executions, which is far more flexible.

The combination of EC+ASP is an interesting tool for game-related AI because ASP provides fast inference to models of a game’s execution, and EC is a solid knowledge representation for abstract worlds. Together they facilitate generating traces from concise declarative descriptions.

The Logical Game Engine

There are some drawbacks to modeling games directly in EC+ASP. In particular, nearly all of the T variables are superfluous because, outside of the event calculus axioms, the logical rules of a game tend to refer only to the current time. Secondly, crafting more complex games directly in the event calculus formalism leads to duplication of common preconditions for an event in each of its initiates/terminates clauses. This duplication creates a maintenance challenge for the game’s author, who might want to make simple

changes to the conditions for a particular event. Finally, expressing the fact that some set of game events conflict with each other, despite being independently possible, requires mapping a complex idiom (of threading special conditions through most game logic) over each new game design. Mutually conflicting events are a common occurrence even in simple turn-based games, but describing this constraint directly is, in our experience, error-prone. These annoyances are resolved in our engine while retaining the advantages of EC+ASP.

Programming computer games is an immensely difficult task, and only through the continued application of software engineering practices has a sequence of increasingly powerful game engines made possible the rich games we expect today (Blow 2004). Game engines attempt to provide standard solutions for game programming problems without prescribing particular rules or setting for a game. These solutions are exposed by a set of APIs that games can be programmed against. Some engines, such as Torque² and Unreal,³ go as far to provide custom programming languages to ease the integration of a game’s specification with the engine’s services. LUDOCORE is modeled after this richer variant of game engines: it provides not only APIs that encapsulate solutions to difficult logical modeling problems (such as conflicting events) but effectively provides a new language that is tailored to the application of specifying logical game worlds. Games produced with our engine are not only smaller than their equivalent specification using only the event calculus, but also easier to maintain throughout meaningful design changes, and easier for the game’s author (or even automated tools) to analyze.

Our game engine is essentially a background theory for logical game descriptions. In the rest of this section we will describe the logical predicates of our game engine and how individual games can leverage them.

State, events, and consequences

Since state and events are natural elements of a game’s definition, we expose the event calculus’s fluents and events in our game engine with the predicates *game_state* and *game_event*, respectively. A game state assertion in LUDOCORE looks like this:

$$\text{game_state}(\text{at}(A, R)) \leftarrow \text{agent}(A), \text{room}(R).$$

This assertion reads that *at* is conceptually a table that records the relation between agents and rooms (answering, “is agent A in room R ?”). Elements of this table stay set until changed otherwise, a property inherited from the event calculus’s commonsense law of inertia.

For state that should be updated dynamically as a function of other state, we provide the *state_helper* mechanism, which provides a computed view on inertial state. State helpers are, semantically, event-calculus fluents with inertia disabled, which cannot be directly initiated or terminated. This stratification into primary and derived fluents that we enforce syntactically is one safe idiom for avoiding the ramification problem in the event calculus (Shanahan 1999).

¹<http://decreasoner.sourceforge.net/csr/ecas/>

²<http://www.torquepowered.com>

³<http://www.unrealtechnology.com>

The *state_helper* assertion below provides a convenient view on the *at* state for checking when a particular agent is at their starting location (*starts_in*) without having to name that location. This example also illustrates how rules can be conditioned on the current game state using the *holds* predicate. This state helper rule implies that *home* is a dynamic (time-varying) property of the game world:

$state_helper(home(A)) \leftarrow holds(at(A, R)), starts_in(A, R).$

Game-event declarations are relatively straightforward:

$game_event(heals(A)) \leftarrow agent(A).$

However, game events can have significantly more structure in LUDOCORE than in the event calculus alone. For example, the conditions under which a game event is possible (for example, the legality of moves in a board game) are specified using the *possible* predicate:

$possible(heals(A)) \leftarrow holds(home(A)).$

Possibility assertions are used not only for defining the rules of the game, but also to control event selection when generating gameplay traces. Importantly, possibility is deeply intertwined with the event conflict mechanism. An assertion like the one below can keep two otherwise *possible* events from happening together:

$conflicts(heals(A), moves(A, R)).$

By providing *possible* and pairwise *conflicts* conditions for game events in a game model, the required event-selection logic can be implemented and debugged once in the game engine, instead of by each game author. By default, actions that are not marked to conflict are safe to co-occur (whereas GDL enforces exactly one player action per turn, which is unsuitable for modeling general videogames).

Game events can be tagged by the game’s author with whether they are direct player actions or spontaneous actions that can only be caused by a non-player entity, thought of as “nature” or “the game master”. The *player_event* and *nature_event* predicates signify this tagging. Although these annotations do not correspond to what we normally think of as game rules, they play an important role in scoping the applicability of player-modeling rules, described with the *player_asserts* and *player_forbids* predicates (and corresponding nature predicates), discussed later.

Linking game state to game events, the *initiates* and *terminates* predicates from the event calculus are exposed to game authors with minimal change. The only modification our engine uses is that these predicates are defined in a time-invariant fashion, always referring to the current game state. The example below asserts that the *moves* event causes an update to the *at* game state, and that the *hits* event causes the target of the hit to no longer be *alive*, under certain conditions:

$initiates(moves(A, R), at(A, R)).$

$terminates(hits(A_1, A_2), alive(A_2)) \leftarrow$
 $holds(armed(A_1)), not\ holds(armed(A_2)).$

Finally, the initial state of the game world is given by, simply enough, the *initially* predicate, which is often conditioned on world configuration. Even in the examples above,

we have suggested the existence of *room* and *starts_in* predicates as background knowledge used to control our state-and-event logic. Many predicates of this kind can be thought of as specifying a static world configuration: time-invariant facts about the game world, such as level geometry, item properties, and tables of weapon effects.

Player model interface

In addition to tagging the subset of game events that are player actions with *player_event*, we build an additional interface for specifying models for players’ behavior.

Simply allowing a certain set of events to be considered by the player, by tagging them as player events, does not immediately yield an expressive modeling interface. The *player_asserts* and *player_forbids* predicates can be used to express stronger preferences that an action be taken or should never be taken under any circumstances (*forbids* beats *asserts* in our engine). The example below illustrates these assertions in the context of a hypothetical game in which picking up objects is a player action:

$player_asserts(pickup(O)) \leftarrow kind(O, gold).$

$player_forbids(pickup(O_1)) \leftarrow$
 $kind(O_1, K), holds(carrying(O_2)), kind(O_2, K).$

This model describes a player who never misses a chance to pick up a gold object, but never picks up duplicates (even of gold objects).

The combination of such assertions allows for the specification of a complex, nondeterministic player action policy that automatically respects the game’s mechanics. The default player model is maximally permissive (it reads that the player considers all player events) which allows meaningful play traces to be extracted from games even when no effort has been put into player modeling. Every clause of *player_forbids* or *player_asserts* serves to pare down the space of gameplay traces to those that are more reasonable to expect, given the provided knowledge.

Identical in structure to the player modeling interface, we provide *nature_asserts* and *nature_forbids* predicates to operate on the events marked with *nature_event*. Whether the nature model is used to model an opponent, a game master, a collection of non-player characters, or even used at all, is largely a game author’s choice. Many of the rules that end up in a nature model could be pushed into the game’s core rules, resulting in an equivalent logical model. Leaving certain elements of the game in the realm of the nature model, however, makes it easy to experiment with modification to these policies without modifying the accepted base rules. Having such flexibility is crucial when modeling videogames with a live world, full of enemies, moving platforms, and other active, non-player entities.

Relation to the general event calculus

In comparison to the general event calculus, we disallow *holds* and *happens* from ever being used at the head of a rule. In event-calculus terms, that means we disallow state constraints and triggers.

In general, state constraints are not game mechanics. For example, the state constraint that the hero is never at the bottom of a pit while alive might be true of a bit of design, but it is not itself a game rule. There may be rules that produce this state constraint: for example, saying that the hero dies when he hits the bottom of pits (a *terminates* clause); and so would a rule that prevented the hero from ever falling off ledges into pits (a *possible* clause). The result is that every state is either only changed by *initiates* and *terminates*, or is known to be completely passive, in the case of state helpers. Disallowing general state constraints ensures that games have well-defined, consistent, procedural meanings (which is required for them to be human-playable).

By disallowing a game’s author from directly specifying when an event *happens*, we can ensure that our engine has complete control over event management, allowing us to implement *possible*, *conflicts*, and the player and nature models. The logic required to implement the required semantics of these predicates is highly circular and counterintuitive. Since such a mechanism is required for almost any interesting videogame model, by implementing it once and for all in LUDOCORE, a game author doesn’t have to create an ad-hoc reimplementations of similar constructs for each game modeled.

While we disable rules with *happens* at the head, our engine permits a conceptual equivalent to the use of triggers (rules that specify that an event happens whenever a particular combination of state holds) through the nature model. For example, to model a triggered collision event, the nature model could include a line that says *nature_asserts* the collision event between two objects if their position is identical.

In summary, our engine provides a modified view of the event calculus that is designed to make game-level assertions easy to express while minimizing the possibility of a game author (modeler) accidentally introducing purely logical bugs such as deductive loops and contradictions. This lets designers focus on whether their set of game rules is an interesting game, and how it functions, as opposed to spending much time worrying about whether their set of logical assertions specifies a game at all.

LUDOCORE in Action

Having defined the major ideas in our game engine, we now use it to generate interesting gameplay traces, and to perform incremental rule modification, world configuration changes, and player modeling.

Gameplay Trace Inference

The simplest use of our engine is to simply generate any gameplay trace compatible with a game’s definition. One of the difficulties for game designers is understanding the potential consequences of rule interactions. When initially crafting game models, these unconstrained traces provide quick access to possible interactions within the game, quickly revealing undesirable behaviors resulting from game-design bugs. The output of our tools provides not only a log of game events, but the complete game state at all times, to ease diagnosis.

The most common use of LUDOCORE is combining a game with a set of hypothetical or speculative assumptions (SAs). SAs are commonly used to assert that certain events do happen or never happen, or that a certain condition is met by a certain timepoint (e.g., victory within 20 turns).

The mechanism of specifying speculative assumptions (based on integrity constraints) is very powerful and provides a simple, modular querying interface for asking questions about a game. “Is my game winnable?” requires only one SA to specify (asserting that *victory* happens at some timepoint). To scope the question further, asking “is my game winnable with avatar health never dropping below 5?”, requires only one more SA.

Here, the logical inference approach taken in EC+ASP shines. Had we used randomized forward search in the game’s state space, constraints on happenings of the final timepoint would be difficult to encode and inefficient to compute, requiring exhaustive search to prove nonexistence. ASP lets us “run the game backwards” or even sideways⁴ as needed to satisfy trace constraints or quickly prove them impossible.

While gameplay traces can be highly informative (because each includes a complete narrative of events), even the absence of traces can be informative. If adding a particular SA yields no traces for a game that previously admitted many, then it has been proven (by contradiction) that the SA is incompatible with the game’s rules, i.e. that the assumption is false in the game’s abstract world.

Modifying rules and configuration

Recall that game design is often an iterative process, adding, removing, and tweaking rules as a design progresses. Rules, in LUDOCORE, are represented by the *game_event*, *game_state*, *possible*, *conflicts*, and *initiates/terminates* predicates. A game’s author might make changes to these rules (and examine the traces for the new game) as part of a major intentional design change, or simply as a way of quickly testing the implications of alternative formulations of a particular mechanic in the game. Often it is useful to disable (by simply commenting-out) several rules to focus on sub-parts of a game in isolation without distracting interactions, e.g. disabling ammunition consumption on weapons when examining health-point management. In a language without elaboration tolerance, this would require more complex editing than simply commenting out a rule. Creating a variant of chess without castling in GDL, for example, cannot be done simply by removing the castling rule, but requires other edits as well.

We have used LUDOCORE to model a popular computer game, *Motherload*,⁵ in which a mining robot recovers ore and treasures from a network of caverns without running out of fuel. In experimentation with this model, we used gameplay trace inference to gain insight into how the game plays, and the effects of design decisions that could be made. To

⁴Inferring which state is compatible with a particular narrative of events

⁵<http://www.xgenstudios.com/play/motherload>

illustrate rule modifications, we draw on two examples from this game model.

If we wanted to add a fixed inventory capacity to our mining robot (constraining an existing event on the basis of existing game state), we would modify the *possible* conditions for our game’s *mine* event to depend on a count of the number of items for which the *bagged* state for rocks held. This single change would not only stop traces that violate the new rule from being generated but reject traces when even an externally assumed narrative included too many mining events:

```
state_helper(have_space) ←
  count(R, holds(bagged(R)), N), N < 10.
possible(mine(R)) ←
  holds(present(R)), touching(R), fueled, have_space.
```

To modify the game to charge players energy points for mining instead of moving, we would drop the *initiates* and *terminates* rules linking the moving event to the energy state and create two new ones linking mining to moving (one to terminate the previous energy value and another to initiate the diminished value). This re-wiring task incurs only a four-line change to the game source when using our engine, due to the conciseness of expressions enabled by our game engine. With access to gameplay trace inference, we can quickly verify the effectiveness of the rule modification without the need for human play testers.

This ease of rule modifications directly derives from the elaboration tolerance afforded by the underlying EC formalism. McCarthy (1998) describes elaboration tolerance as the ability (of a knowledge representation) to accept changes to known facts without having to “start over”, which is exactly what we realize for the space of games. Further, in other research applying EC to games, we have shown how entire modules (corresponding to game mechanics and vocabularies) can be swapped in and out without trouble (OMITTED 2008).

In addition to swapping rules, it is quite easy to pair a fixed game with different sets of configuration such as map layout, static item properties, and tweakable constants. For example, a simple map can be used for early testing, and a more complex one used later for detailed player modeling. In addition to manually specified game configurations, it is possible to let the ASP solver reason across possible configurations for the game. In our model of the game *Minesweeper*, we are able to reason over possible placements of mines that are consistent with player observations. In a simple dungeon crawl game, we allowed the solver to manipulate the presence and connectivity of rooms on a map. The ability to make structural queries of this sort is a novel feature for game engines.

Using player and nature models

While we think of speculative assumptions as convenient but throwaway constraints, specifying complex assumptions can be tedious. The player modeling interface provides a more straightforward way of building larger player models that will be retained and modified, rather than used in only a few queries. Sets of *player_asserts* and *player_forbids*

clauses are collected to carve out the space of behavior that can be reasonably expected of the kinds of player being investigated.

A custom nature model can similarly be used to carve the behavior of other in-game entities. In the dungeon crawl game, we used the nature model to control monster wandering behavior. We had tagged the monster movement event as a nature event, implying that it should always be considered in trace finding when using the default nature model. To effect patrolling behavior, we used the following assertion, which yields monsters who wander only within their own marked areas:

```
nature_forbids(move_to(M, R)) ← not patrolled_by(M, R).
```

Applications Enabled

LUDOCORE enables a number of broader applications, some of which have already been investigated in other forms.

Game design

Designing games in a logical game engine, or at least using one to prototype game ideas, can provide designers with insight about design pros and cons, and suggest possible improvements. Many logical formalisms can allow verification that desired properties hold, and undesired properties don’t. However, designers often work in an iterative, exploratory manner, and find exact yes-or-no queries somewhat difficult to formulate (OMITTED 2009b). The experimentation with gameplay traces that our engine supports can help designers understand the possibilities of their design by iteratively “zooming in” on specific kinds of traces using SAs, and observing how different player models shape emergent properties of the game. We have already begun pursuing this application, creating a game prototyping system (OMITTED 2009a) in which human-playable, board-game-like prototypes can be built, using logical gameplay rules. These prototypes support two kinds of playtesting: traditional observations of human players, and a new mode of machine playtesting, in which observation of machine-generated gameplay traces is the focus.

Logical game engines may also form a core component of emerging research in automated videogame design. In a non-logical approach, Togelius and Schmidhuber (2008) generates *Pac-Man* variants from a parameterized space and uses reinforcement learning to demonstrate gameplay in the resulting games. Games implemented in LUDOCORE can be varied in a much more open-ended and incremental manner than with a parameterized space of variation, and its elaboration tolerance address the problem of brittleness that drove researchers to parameterized representations in the first place.

Furthermore, given sets of traces, inductive logic programming (ILP) can be used to induce models for the style of gameplay exhibited in those traces. The player modeling interface in LUDOCORE makes this more feasible by having player models built from only two key predicates. Many popular ILP systems, such as Progol (Muggleton 1995) can only learn single predicates at a time. In our player modeling interface, single-predicate ILP systems can learn the

player.asserts predicate from example gameplay traces. Some research has even been done on learning EC rules themselves (Moyle and Muggleton 1997), which, for games, would open up the possibility of inducing new game rules from a collection of desired gameplay traces.

Procedural level design can be done by specifying and solving constraints for what constitutes a good level, sometimes as part of a larger search process (e.g. Smith et al. 2009). Particularly relevant here, a community-driven project recently added procedurally generated maps and base layouts to the real-time strategy game Warzone 2100, using ASP to specify and solve constraints on the maps' layouts.⁶ Designing a good level, however, often means designing a level that supports good gameplay, and these sorts of static constraints only indirectly speak to gameplay. Other approaches, such as that of Togelius et al. (2007), generate levels and score them with a fitness function that predicts whether they support interesting gameplay. For games written using our logical game engine, a constraint-based approach can easily include these kinds of constraints on gameplay in addition to constraints on static level properties, because the level-generation search and game-tree search are unified into the same query mechanism.

Crafting game playing agents

Despite our motivating focus on analyzing game designs, our logical game engine has applications to playing games as well. Because gameplay trace inference in our engine corresponds exactly to what is commonly known as EC planning (Shanahan 1997a), it is possible to use our engine directly in a general game player. Thielscher (2009) already argued for the applicability of ASP with temporal-logic models for the contemplation phase of GGP competitions.

The player modeling interface we provide is designed to accept incremental additions of knowledge about how a player (or their opponents) make choices. Though this interface cannot express a minimax-like policy (that includes quantification over models at every timepoint), it does coalesce facts into a unified move-set selection policy which would allow minimax to operate in the more restricted space in which the modeled opponent is really playing.

The ease of adding and removing rules in our games has another benefit for those wishing to craft general game players. The elaboration tolerance of the representation makes syntactic construction of novel games much easier. By building all combinations of a fixed set of add-on mechanics, a generative space of games can be realized, providing a much wider selections of games for agents to be tested on.

Conclusion

By carefully engineering the abstractions in LUDOCORE, we hope to build a bridge between the concerns of game design and logicist AI. For game designers, who are continually thinking about formal rule systems, but almost always in the context of procedural programming environments, we propose a logical game engine as a link between their needs

and the standard tools of AI reasoning. On the AI side, we hope our games-specific logical background theory providing base support for the building blocks of videogames will encourage the use of automated reasoning in an exciting new domain.

References

- Adams, E. 2009. *Fundamentals of Game Design*. New Riders.
- Blow, J. 2004. Game development: Harder than you think. *ACM Queue* 1(10):28–37.
- Fullerton, T. 2008. *Game Design Workshop*. Morgan Kaufmann.
- Genesereth, M. 2005. General game playing: Overview of the AAAI competition. *AI Magazine* 26(2):62–72.
- Kim, T.-W.; Lee, J.; and Palla, R. 2009. Circumscriptive event calculus as answer set programming. In *Proc. 21st Intl. Joint Conference on Artificial Intelligence*, 823–829.
- Lifschitz, V. 2008. What is answer set programming? In *Proc. 23rd Ntl. Conference on Artificial Intelligence (AAAI)*, 1594–1597.
- Love, N.; Hinrichs, T.; Haley, D.; Schkufza, E.; and Genesereth, M. 2006. General game playing: Game description language specification. Technical Report LG-2006-01, Stanford Logic Group.
- McCarthy, J. 1998. Elaboration tolerance. In *Proc. 4th Symposium on Logical Formalizations of Commonsense Reasoning*.
- Moyle, S., and Muggleton, S. 1997. Learning programs in the event calculus. In *Proc. 7th Intl. Workshop on Inductive Logic Programming*, 205–212.
- Mueller, E. T. 2004. Event calculus reasoning through satisfiability. *Journal of Logic and Computation* 14(5):703–730.
- Mueller, E. T. 2006. *Commonsense Reasoning*. Morgan Kaufmann.
- Muggleton, S. 1995. Inverse entailment and Progol. *New Generation Computing Journal* 13:245–286.
- OMITTED. 2008. Omitted for blind review.
- OMITTED. 2009a. Omitted for blind review.
- OMITTED. 2009b. Omitted for blind review.
- Shanahan, M. 1997a. Event calculus planning revisited. In *Proc. 4th European Conference on Planning*, 390–402.
- Shanahan, M. 1997b. *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*. MIT Press.
- Shanahan, M. 1999. The ramification problem in the event calculus. In *Proc. 16th Intl. Joint Conference on Artificial Intelligence*, 140–146.
- Smith, G.; Treanor, M.; Whitehead, J.; and Mateas, M. 2009. Rhythm-based level generation for 2D platformers. In *Proc. 4th Intl. Conference on the Foundations of Digital Games*, 175–182.
- Thielscher, M. 2009. Answer set programming for single-player games in general game playing. In *Proc. 25th International Conference on Logic Programming*, 327–341.
- Togelius, J., and Schmidhuber, J. 2008. An experiment in automatic game design. In *Proc. 2008 IEEE Symposium on Computational Intelligence and Games*, 111–118.
- Togelius, J.; De Nardi, R.; and Lucas, S. M. 2007. Towards automatic personalised content creation for racing games. In *Proc. 2007 IEEE Symposium on Computational Intelligence and Games*, 252–259.

⁶<http://warzone2100.org.uk/manual-diorama>