

Ludii - The Ludemic General Game System

Éric Piette, Dennis J.N.J. Soemers, Matthew Stephenson, Chiara F. Sironi,
Mark H.M. Winands and Cameron Browne

Department of Data Science and Knowledge Engineering (DKE)

Maastricht University, Bouillonstraat 8-10

Maastricht, 6211 LH, The Netherlands

{eric.piette, dennis.soemers, matthew.stephenson, c.sironi,
m.winands, cameron.browne}@maastrichtuniversity.nl

Abstract

While current General Game Playing (GGP) systems facilitate useful research in Artificial Intelligence (AI) for game-playing, they are often somewhat specialized and computationally inefficient. In this paper, we describe an initial version of a “ludemic” general game system called Ludii, which has the potential to provide an efficient tool for AI researchers as well game designers, historians, educators and practitioners in related fields. Ludii defines games as structures of ludemes, i.e. high-level, easily understandable game concepts. We establish the foundations of Ludii by outlining its main benefits: generality, extensibility, understandability and efficiency. Experimentally, Ludii outperforms one of the most efficient Game Description Language (GDL) reasoners, based on a propositional network, for all available games in the Tilt-yard GGP repository.

1 Introduction

The goal of *General Game Playing* (GGP) is to develop artificial agents capable of playing a wide variety of games [Pitrat, 1968]. Several different software systems for modelling games, commonly called General Game Systems, currently exist for different types of games, including: deterministic perfect-information games [Genesereth *et al.*, 2005], combinatorial games [Browne, 2009], puzzle games [Shaker *et al.*, 2013], strategy games [Mahlmann *et al.*, 2011], card games [Font *et al.*, 2013] and video games [Schaul, 2014].

Since 2005, the General Game System GGP-BASE¹ using the *Game Description Language* (GDL) [Love *et al.*, 2008] has become the standard for academic research in GGP. GDL is a set of first-order logical clauses describing games in terms of simple instructions. While it is designed for deterministic games with perfect information, an extension named “GDL-II” [Schiffel and Thielscher, 2014] has been developed for games with hidden information, and another extension named “GDL-III” [Thielscher, 2017] has been developed for epistemic games.

¹GGP-BASE <https://github.com/ggp-org/ggp-base>

1.1 GDL Background

The generality of GDL provides a high level of algorithmic challenge and has led to important research contributions [Björnsson and Schiffel, 2016], especially in *Monte Carlo tree search* (MCTS) enhancements [Finnsson and Björnsson, 2008; Finnsson and Björnsson, 2010], with some original algorithms combining constraint programming, MCTS, and symmetry detection [Koriche *et al.*, 2017]. Unfortunately, the key structural aspects of games – such as the board or card deck, and arithmetic operators – must be defined explicitly from scratch for each game definition. GDL is also limited in terms of potential applications outside of game AI.

Game descriptions can be time consuming to write and debug, and difficult to decipher for those unfamiliar with first order logic. The equipment and rules are typically interconnected to such an extent that changing any aspect of the game would require significant code rewriting. For example, changing the board size from 3×3 to 4×4 in the Tic-Tac-Toe description would require many lines of code to be added or modified. GDL game descriptions are verbose and difficult for humans to understand, and do not encapsulate the key game-related concepts that human designers typically use when thinking about games. Processing such descriptions is also computationally expensive as it requires logic resolution, making the language difficult to integrate with other external applications. Some complex games can be difficult and time consuming to model (e.g., *Go*), or are rendered unplayable due to computational costs (e.g., *Chess*). The main GGP/GDL repository [Schreiber, 2016] is only extended with a few games every year.

1.2 The Digital Ludeme Project

The Digital Ludeme Project (DLP)² is a five-year research project, recently launched at Maastricht University, which aims to model the world’s traditional strategy games in a single, playable digital database. This database will be used to find relationships between games and their components, in order to develop a model for the evolution of games throughout recorded human history and to chart their spread across cultures worldwide. This project will establish a new field of research called *Digital Archæoludology* [Browne, 2017].

²Digital Ludeme Project: <http://ludeme.eu/>

The DLP aims to model the 1,000 most influential traditional games throughout history, each of which may have multiple interpretations and require hundreds of variant rule sets to be tested. This is therefore not just a mathematical/computational challenge, but also a logistical one requiring a new kind of General Game System. The DLP deals with traditional games of strategy including most board games, card games, dice games, tile games, etc., and may involve non-deterministic elements of chance or hidden information, as long as strategic play is rewarded over random play; we exclude dexterity games, physical games, video games, etc.

In this paper, we formally introduce an initial version of Ludii,³ the first complete Ludemic General Game System able to model and play (by a human or AI) the full range of traditional strategy games. We introduce the notion of ludemes in Section 2, the ludemic approach that we have implemented in Section 3, the Ludii System itself in Section 4, its abilities to provide the necessary applications to the Digital Ludeme Project are highlighted in Section 5, and the underlying efficiency of the Ludii system in terms of reasoning is demonstrated experimentally in Section 6 by a comparison with one of the best GGP-BASE reasoners [Sironi and Winands, 2017].

2 Ludemes

The decomposition of games into their component *ludemes* [Parlett, 2016], i.e. conceptual units of game-related information, allows us to distinguish between a game’s *form* (its rules and equipment) and its *function* (its emergent behaviour through play). This separation provides a clear genotype/phenotype analogy that makes phylogenetic analysis possible, with ludemes making up the “DNA” that defines each game.

This ludemic model of games was successfully demonstrated in earlier work to evolve new board games from existing ones [Browne, 2011]. An important benefit of the ludemic approach is that it encapsulates key game concepts, and gives them meaningful labels. This allows for the automatic description of game rule sets, comparisons between games, and potentially the automated explanation of learnt strategies in human-comprehensible terms. Recent work shows how this model can be enhanced for greater generality and extensibility, to allow any ludeme that can be computationally modelled to be defined using a so-called *class grammar* approach, which derives the game description language directly from the class hierarchy of the underlying source code library [Browne, 2016].

This approach provides the potential for a single AI software tool that is able to model, play, and analyse almost any traditional game of strategy as a structure of ludemes. It also provides a mechanism for identifying underlying mathematical correspondences between games, to establish probabilistic relationships between them, in lieu of an actual genetic heritage.

Note that in this initial version of Ludii reported here, the game definitions are optimised for performance (rather than clarity, generality or extensibility). This demonstrates the

flexibility of the approach that games can be implemented to required levels of clarity, generality or performance, depending on the need.

3 Ludemic Approach

We now outline the *ludemic approach* used to model games.

3.1 Syntax

Definition 1. A Ludii game state s encodes which player is to move in s (denoted by $mover(s)$), and five vectors each containing data for every possible *location*; *what*, *who*, *count*, *state*, and *hidden*. A more precise description of the locations and the specific data in these vector is given after Definition 3.

Definition 2. A Ludii successor function is given by

$$\mathcal{T} : (S \setminus S_{ter}, \mathcal{A}) \mapsto S,$$

where S is the set of all the Ludii game states, S_{ter} the set of all the terminal states, and \mathcal{A} the set of all possible lists of actions.

Given a current state $s \in S \setminus S_{ter}$, and a list of actions $A = [a_i] \in \mathcal{A}$, \mathcal{T} computes a successor state $s' \in S$. Intuitively, a complete list of actions A can be understood as a single “move” selected by a player, which may have multiple effects on a game state (each implemented by a different primitive action).

Definition 3. A Ludii game is given by a 3-tuple of ludemes $\mathcal{G} = \langle Mode, Equipment, Rules \rangle$ where:

- *Mode* = $\{p_0, p_1, \dots, p_k\}$ is a finite set of $k + 1$ players, where $k \geq 1$. Random game elements (such as rolling dice, flipping a coin, dealing cards, etc.) are provided by p_0 , which denotes *nature*. The first player to move in any game is p_1 , and the current player is referred to as the mover.
- *Equipment* = $\langle C^t, C^p \rangle$ where:
 - C^t denotes a list of containers (boards, player’s hands, etc.). Every container $c^t = \langle V, E \rangle$, where $c^t \in C^t$, is a graph with vertices V and edges E . Every vertex $v_i \in V$ corresponds to a playable site (e.g. a square in Chess, or an intersection in Go), while each edge $e_i \in E$ represents that two sites are adjacent.
 - C^p denotes a list of components (pieces, cards, tiles, dice, etc.), some of which may be placed on sites of the containers in C^t . We use the convention that the component $c_0^p \in C^p$ is placed on all “empty” sites.
- *Rules* defines the operations of the game, including:
 - *Start* = $[a_0, a_1, \dots, a_k]$ denotes a list of starting actions. The starting actions are sequentially applied to an “empty” state (state with c_0 on all sites of all containers) to model the initial state s_0 .
 - *Play* : $S \mapsto \mathcal{P}(\mathcal{A})$, where $\mathcal{P}(\mathcal{A})$ denotes the powerset of the set \mathcal{A} of all possible lists of legal actions. This is a function that, given a state $s \in S$, returns a set of lists of actions.

³Ludii is named after its predecessor LUDI [Browne, 2009]

- $End = (Cond_0(s), \vec{S}_0) \cup \dots \cup (Cond_e(s), \vec{S}_e)$ denotes a set of conditions $Cond_i(s)$ under which a given state s is considered to be terminal. Each termination condition $Cond_i(s)$ leads to a score vector \vec{S}_i .

Ludii provides some predefined vectors: Win, Loss, Draw, Tie, and Abort. Moreover, if the *mover* has no legal moves then they are in a (temporary) *Stalemate* and must perform the special action `pass`, unless the *End* rules dictate otherwise. States in which all players were forced to `pass` for the last complete round of play are abandoned as a *Draw*.

We specify *locations* $loc = \langle c, v_i, l_i \rangle$ by their container $c = \langle V, E \rangle$, a vertex $v_i \in V$, and a level $l_i \geq 0$. Every location specifies a specific site in a specific container at a specific level, where most games only use $l_i = 0$ but stacking games may use more levels. For every such location, a game state s encodes multiple pieces of data, as described in Definition 1. The index of a component located at loc in s is given by $what(s, loc)$, the owner (player index) by $who(s, loc)$, the number of components by $count(s, loc)$, the internal state of a component (direction, side, promotion status, etc.) by $state(s, loc)$. If the state of a location loc is hidden information for a certain player p_i , that is given by $hidden(s, loc, p_i)$.

3.2 Ludii Example

Following Definition 3, Ludii provides a variety of ludemes corresponding to simple operations that can be used to define players, equipment or rules. For example, `(line 3)` is a Boolean ludeme that returns `True` if there is a line of three pieces in a container, and `(empty)` is a ludeme that returns a list containing all empty sites of a container.

In Figure 1, a complete description of the game Tic-Tac-Toe is given according to the EBNF-style grammar generated by Ludii. The `mode` ludeme describes the mode of play; a game with alternating turns played between two players, named `P1` and `P2`. The first subset of the `equipment` ludeme describes the main board as a square 3×3 tiling, with the second subset listing the components as a disc piece type named “O” for player 1 and a cross piece type named “X” for player 2. Each turn, the mover plays a piece of their colour at any empty cell, which is implemented by `(to Mover (empty))`. The winning condition for the mover is to create a line of three pieces. Tic-Tac-Toe does not require any *Start* rules.

```
(game "Tic-Tac-Toe"
  (mode {(player "P1") (player "P2")})
  (equipment
    {(board "Board" (square 3))}
    {(disc "O" 1) (cross "X" 2)})
  )
  (rules
    (play (to Mover (empty)))
    (end (line 3) (result Mover win))
  )
)
```

Figure 1: The game of Tic-Tac-Toe modelled in Ludii.

If the board fills before either player wins, then game defaults to a *Draw* after both players are forced to pass. Note that judicious use of default settings for common game behaviours allows succinct game descriptions.

4 Ludii System

The next section introduces the Ludii system itself, describing both the grammar approach and the core of the system.

4.1 Class Grammar

Ludii is a complete general game system [Browne *et al.*, 2014] that uses a *class grammar* approach, in which the game description language is automatically generated from the constructors in the class hierarchy of the Ludii source code [Browne, 2016]. Game descriptions expressed in the grammar are automatically instantiated back into the corresponding library code for compilation, giving a guaranteed 1:1 mapping between the source code and the grammar.

Schaul *et al.* [2011] points out that “*any programming language constitutes a game description language, as would a universal Turing machine*”. Ludii effectively makes its programming language (Java) the game description language. It can theoretically support any rule, equipment or behaviour that can be programmed in Java. The implementation details are hidden from the user, who only sees the simplified grammar which summarises the code to be called.

4.2 The Core of Ludii

The core of Ludii is a ludeme library implemented in Java 8, consisting of a number of classes, each implementing a specific ludeme. A Ludii game \mathcal{G} defining all relevant ludemes (players, equipment, rules) is stored as a single immutable `Game` object. In the context of General Game Playing, displaying any game automatically is important for understanding strategies by AI players. To this end, all equipment in Ludii implements the `Drawable` interface, which means that each item of equipment is able to draw a default bitmap image for itself at a given resolution, for displaying the board state. Containers C^t are able to draw their current components at the appropriate positions, orientations, states, etc. A `View` object provides the mechanism for showing the current game state on the screen and a `Controller` object provides the mechanism for updating the game state based on user input such as mouse clicks. All games available in the system can be played by both humans and/or AI.

As an example, Figure 2 shows a 2-player game \mathcal{G} with $C^t = \{c_0^t\}$, where c_0^t is a hexagonal container with hexagonal tiles. $C^p = \{c_0^p, c_1^p, c_2^p\}$, where c_0^p is the empty component, c_1^p is the white disc for the player p_1 , and c_2^p the black disc for player p_2 . The system has a graph representation of the board for visualisation, the vertices, edges, and faces of this graph are depicted in blue. The dual of this graph, which is the graph given by c_0^t , is depicted in grey.

The game graph itself can be modified during certain *graph games* (e.g. Dots & Boxes) in which a player’s moves involve operations on the graph (e.g. adding or cutting edges or vertices). Reasoning efficiency can be optimized by pre-generating data such as corners of the board, exterior vertices, vertices along the top side of the board, etc. within




Figure 2: A hex board container with hexagonal tiling and the dual of its graph (which is itself a graph).

the `Graph` class, and vertex neighbours indexed by direction, vertices reached by turtle steps, etc. within the `Vertex` class.

The data vectors *what*, *who*, etc. of a state s are implemented in a collection of `ContainerState` objects. Different representations are implemented in order to minimise the memory footprint and to optimise the time needed to access necessary data for reasoning on any game:

- Uniform pieces per player (e.g. Tic-Tac-Toe).
- Distinguished pieces per player (e.g. Chess).
- Piece state or direction per site (e.g. Reversi, Ploy).
- Piece count per site (e.g. Mancala).
- Piece stacking (e.g. Laska).
- No fixed board (e.g. Dominoes).
- Hidden information (e.g. Stratego, card games).

Ludii automatically selects the appropriate state type from the rules before creating the `Game` object, to ensure the most suitable representation is used.

Container states are defined using a custom `BitSet` class, called `ChunkSet`, that compresses the required state information into a minimal memory footprint, based on each game’s definition. For example, if a game involves no more equipment than a board and uniform pieces in N colours, then the game state is described by a `ChunkSet` subdivided into chunks of B bits per board cell, where B is the lowest power of 2 that provides enough bits to represent every possible state per cell (including state 0 for the empty cells).⁴

Using the game state and the different ludemes describing the game rules, the system can compute the legal moves for any state. The tree of ludemes is evaluated to return the list of `Action` objects applicable for the mover. Each `Action` object describes one or more atomic actions to be applied to the game state to execute a move. Actions typically include adding or removing components to/from containers, or changing component counts or states within containers.

Definition 4. A trial τ is a sequence of states s_i and action lists A_i :

$$\langle s_0, A_1, s_1, \dots, s_{f-1}, A_f, s_f \rangle$$

⁴Chunk sizes are set to the lowest power of 2 to avoid issues with chunks straddling consecutive long values.

such that $f \geq 0$, and for all $i \in \{1, \dots, f\}$,

- the played action list A_i is legal for the $mover(s_{i-1})$
- states are updated: $s_i = \mathcal{T}(s_{i-1}, A_i)$
- only s_f may be terminal: $\{s_0, \dots, s_{f-1}\} \cap S_{ter} = \emptyset$

τ is stored in a `Trial` object, providing a complete record of a game played from start to end, including the moves made.

Any reasoning on any game can be parallelised using separate trials per thread. All the data members of the `Game` object are constant and can therefore be shared between threads. A thread will be able to use a `Trial` object to compute any playouts from any state. On the system each `AI` object describes the AI implementation chosen for each player, including computational budget/time limits, hints such as features for biasing playouts [Browne *et al.*, 2019], etc.

5 Benefits and Key Properties

Ludii is being designed and implemented primarily to provide answers to the questions raised by the DLP, but will stand alone as a platform for general games research in areas including AI, design, history and education. Ludii provides many advantages over existing GGP systems, as follows:

Simplicity: Simplicity refers to the ease with which game descriptions can be created and modified, and can be estimated by the number of tokens required to define games. Describing a game with the ludemic approach is typically much simpler compared to a logic-based approach (e.g. LUDII requires only 29 tokens for Tic-Tac-Toe and 298 for Chess, whereas GDL requires 381 and 4,932 tokens respectively) ludemic game descriptions can also be easily modified to test different sizes, geometries or rules. For example, changing the size or shape of a board (e.g. Figure 3) can be accomplished by modifying a single parameter, while the same change in GDL requires many lines of code to be added or modified.

Clarity: Clarity refers to the degree to which game descriptions would be self-explanatory to non-specialist readers. The logic-based game descriptions of GDL are often difficult for humans to interpret. In Ludii, the Java classes that define each ludeme are named using meaningful English labels, providing convenient definitions for the concepts involved. This




Figure 3: The game of Breakthrough on a square (left) and rhombus (right) board shapes.

becomes especially useful for games in which more complex mathematical concepts (geometry, algebra, arithmetic, etc.) are encapsulated within their component ludemes.

Generality: Generality refers to the scope of games covered by the system without the need for extensions. As Ludii uses the class grammar approach to describe the ludemes, it can theoretically support any game that can be programmed in Java. The initial version of Ludii described in this paper already includes many different game types beyond those that can be implemented in GDL. The final version of Ludii will be even more general and facilitate many additional game types.

Extensibility. Extensibility refers to the ease with which new functionality can be added to the system. The initial version of Ludii currently provides 68 different games (and 179 variants) using 259 different ludemes, which test the system’s various options and capabilities. Extending Ludii simply involves adding new classes to the ludeme library, which are then automatically subsumed into the grammar, making extensibility very open-ended. Extending GDL involves significant modifications to the core model and program.

Efficiency: Since the Ludii programmer has complete control of the underlying code – within the constraints of the API and programming guidelines – it is possible to optimise ludemes at any desired level. There is of course a trade-off between game optimisation and description detail. The more optimised a game is, the shorter its description tends to be and the less detail we know about it. This has profound implications for the DLP, in which the ability to reliably compare games for similarity is a key requirement, hence we distinguish between *descriptive* and *optimised* versions of games.

Evolvability: This refers to the likelihood that randomly evolving game descriptions will produce viable children that resemble their parents. GDL game descriptions tend to involve complex chains of logical operations that must be crafted with great care. Randomly applying crossovers and mutations between GDL descriptions is extremely unlikely to produce correct (i.e. playable) results, let alone improve on the parents. Conversely, the ludemic approach is ideally suited to evolutionary approaches such as *genetic programming* [Koza, 1992] and has already proven successful in evolving new high quality games [Browne, 2009].

Cultural Application: Aside from its GGP benefits, Ludii also has several applications as a tool for the new domain of Digital Archæoludology [Browne, 2017]. The Ludii system will eventually be linked to a server and database that stores relevant cultural and historical information about the games. This information will not only provide additional real-world context, but will allow us to reconstruct viable and historically authentic rule-sets for games with incomplete information, develop a “family tree” of traditional games, and help map the spread of games throughout history.

Universality: While Ludii supports a wide range of games, including nondeterministic and hidden information games, we cannot prove the universality of its full grammar within the scope of this paper. We instead show that Ludii is universal for finite deterministic perfect information games:

Theorem 1. *Ludii is universal for the class of finite deterministic games with perfect information.*

In Ludii, this class of games does not require the *nature* player p_0 , and *hidden*(s, loc, p_i) returns false for any state s , any location loc , and any player p_i . The proof⁵ is structured in a similar way to a proof of universality for GDL and GDL-II [Thielscher, 2011]. Based on the definition of extensive-form games [Rasmusen, 2007], we formalise deterministic games and prove that Ludii can define an arbitrary finite game tree. Consequently, all games that can possibly be modelled in GDL can be also described in Ludii.

6 Experiments

The Ludii System – as with most other GGP systems – uses MCTS as its core method for AI move planning, which has proven to be a superior approach for general games in the absence of domain specific knowledge [Finnsson and Björnsson, 2010]. MCTS playouts require fast reasoning engines to achieve the desired number of simulations. Hence, we use *flat Monte Carlo* playouts (i.e., trials τ where $s_f \in S_{ter}$) as the metric for comparing the efficiency of Ludii to other GGP systems.

6.1 Experimental Design

In the following comparison, we compare GGP-BASE and Ludii based on the number of playouts obtained per second. For GGP-BASE, we used the fastest available game implementation⁶ and tested one of the most efficient GGP-BASE reasoners based on *propositional networks* or “propnets” [Sironi and Winands, 2017]. Propnets speed up the reasoning process with respect to custom made or Prolog-based reasoners by translating the GDL rules into a directed graph that resembles a logic circuit, whose nodes correspond to either logic gates or GDL propositions that represent the state, players’ moves and other aspects of the game. Information about the current state can be computed by setting the truth value of the propositions that correspond to the state and propagating these values through the graph. Setting and propagating the truth values of the propositions that correspond to the players’ actions allows us to compute the next state.

All experiments were conducted on a single core of an Intel(R) XEON(R) CPU E5-2680 v2 at 2.80 GHz with 2GB RAM, spending 10 minutes per test.

6.2 Results

The results of our experiments for a selection of games available in GDL, are shown in Table 1. The top section of the table is dedicated to single player games and the bottom section for multiplayer games. The rightmost column shows the speedup factor which Ludii achieves over GGP-BASE.

Table 2 highlights our results for a selection of games, including several historical games, that have no GDL equivalent. The fact that no existing GGP system or game description language supported the full range of games required for the DLP was a driving motivation in developing Ludii.

⁵The complete proof is in the appendix

⁶Taken from [Schreiber, 2016]

Game	Ludii	GDL	Rate
Single player games			
8 Puzzle	26,958	4,113	6.55
8 Queens	963,443	1,946	495.10
16 Queens	785,973	522	575.85
Futoshiki (4×4)	294,772	23,797	12.39
Futoshiki (5×5)	273,348	11,494	23.78
Futoshiki (6×6)	182,245	5,838	31.56
Knight’s Tour (8×8)	91,156	75,000	1.21
Lights Out (5×5)	32,889	11,799	2.79
Nonogram (5×5)	279,242	59,044	4.73
Nonogram (10×10)	45,817	6,883	6.65
Peg Solitaire	7,713	3,172	2.43
Sudoku	88,565	635	139.47
Towers of Hanoi 3	264,345	35,847	7.37
Multi-player games			
Amazons (10×10)	2,933	185	15.85
Breakthrough (6×6)	19,022	3,923	4.85
Breakthrough (8×8)	6,632	1,123	5.91
Chess	1,075	0.06	17,916.67
Connect 4 (6×7)	124,349	13,664	9.10
Connect 4 (12×9)	99,697	9,856	10.12
English Draughts	1,914	872	2.19
Gomoku (15×15)	2,514	927	2.71
Hex (9×9)	21,244	195	108.94
Knightthrough (8×8)	5,085	1,869	2.72
Reversi (8×8)	2,085	203	10.27
Skirmish (8×8)	2,278	124	18.37
Tic-Tac-Toe (3×3)	641,445	85,319	7.52
Tic-Tac-Toe (5×5)	197,972	11,453	17.29
Tron (10×10)	494,999	121,989	4.06
Wolf and Sheep (8×8)	22,305	5,532	4.03

Table 1: The average number of playouts per second for games available on the Tiltyard GGP server.

Game	Ludii	Game	Ludii
Alquerque	796	Mu-Torere	7,719
Ashtapada	12,755	Nine men’s morris	5,121
Connect 6 (19×19)	21,873	Oware	3,275
Dara	5,408	Ploy	819
Fanorona	1,906	Tant Fant	51,549
Hnefatafl (11×11)	326	Three men’s morris	100,174
MineSweeper (8×8)	68,233	Yavalath	189,335

Table 2: The average number of playouts per second for games unavailable in GDL.

6.3 Discussion

Ludii outperforms GGP-BASE in terms of efficiency for all games tested. The performance improvement for single player games (e.g. solitaire puzzles) increased with the size of the puzzle (e.g. Futoshiki from 12.39 to 31.56). For some puzzles with optimised GDL descriptions (e.g. Knight’s Tour) Ludii achieves similar performances, but for others (e.g. N Queens or Sudoku) Ludii achieves significant improvements from 100 to almost 600 times faster.

Ludii is at least 2 times faster for all multi-player games tested. For simpler games, board size is highly correlated

with speed improvement; Ludii is almost 7 times faster for the standard 3×3 Tic-Tac-Toe but almost 17 times faster for the larger 5×5 game. For more complex games – such as Amazons and Skirmish (a variant of Chess where each player has to capture all the enemies pieces) – Ludii is once again more efficient than GDL (over 15 times faster in these cases).

The greatest speed disparity is for Chess, with an improvement rate of almost 18,000. The GDL description of Chess cannot be translated to a propnet because its size exceeds the available memory, therefore we had to use the GDL Prover implementation in the GGP-BASE for comparison. The GGP-BASE Prover is generally slower for more complex games with respect to the propnet, explaining the low number of playouts for Chess (0.06).

The DLP requires the ability to model a broad range of traditional strategy games. In Table 2, Ludii demonstrates its capacity for reasoning in a variety of different games: race games (e.g. Ashtapada), imperfect-information games (e.g. card games), Mancala games (e.g. Oware) games with large boards (e.g. Connect 6), etc. However, for Hnefatafl, the most famous Tafl game, the current version of Ludii obtains a low number of playouts per second. This is likely because Hnefatafl playouts can be lengthy – often over 1,000 moves per game – as the winning condition for each player is unlikely to be reached through strictly random play. Work is in progress to address this for the imminent release version of Ludii.

Kowalski *et al.* [2019] recently proposed the *Regular BoardGames* (RBG) language, based on regular languages, which provides better expressiveness, efficiency, and clarity than GDL. Initial investigations into this language reveal it to be concise but limited to deterministic board games whose geometry can be described in plain ASCII format. It also appears to be less efficient than Ludii’s class grammar approach; e.g. Hex 9×9 achieves 3,425 playouts/second in RBG but 21,244 in Ludii. Future work will include a deeper comparison between the Ludii and RBG approaches.

7 Conclusion

The proposed ludemic General Game System Ludii outperforms GGP-BASE – the current standard for academic AI research into GGP – in terms of reasoning efficiency. It also has advantages in terms of simplicity, clarity, generality, extensibility and evolvability, and has been designed to be applicable to other research fields in addition to game AI.

The potential benefits of this new GGP approach presents several opportunities for future AI work. For example, features discovered by reinforcement learning could be automatically visualised for any game to possibly reveal useful strategies relevant to that game, or provided as human-understandable descriptions based on ludemes with meaningful plain English labels. Another work in progress includes improving AI playing strength by biasing MCTS with features automatically learnt through self-play.

We will continue refining this initial experimental version of Ludii into a stable release version, with improvements in several aspects including further efficiency, for its imminent public release.

Acknowledgements

This research is part of the European Research Council-funded Digital Ludeme Project (ERC Consolidator Grant #771292) being run by Cameron Browne. This work is partially supported by The Netherlands Organisation for Scientific Research (NWO) in the framework of the GoGeneral Project (Grant No. 612.001.121)

References

- [Björnsson and Schiffel, 2016] Yngvi Björnsson and Stephan Schiffel. General game playing. In *Handbook of Digital Games and Entertainment Technologies*, pages 1–23, Singapore, 2016. Springer Singapore.
- [Browne *et al.*, 2014] Cameron Browne, Julian Togelius, and Nathan Sturtevant. Guest editorial: General games. *IEEE Transactions on Computational Intelligence and AI in Games*, 6(4):317–319, Dec 2014.
- [Browne *et al.*, 2019] Cameron Browne, Dennis J.N.J. Soemers, and Eric Piette. Strategic features for general games. In *Proceedings of the 2nd Workshop on Knowledge Extraction from Games co-located with 33rd AAAI Conference on Artificial Intelligence June 27, 2019, Honolulu, Hawaii, USA.*, pages 70–75, 2019.
- [Browne, 2009] Cameron B. Browne. *Automatic generation and evaluation of recombination games*. PhD thesis, Queensland University of Technology, 2009.
- [Browne, 2011] Cameron B. Browne, editor. *Evolutionary Game Design*. Springer, 2011.
- [Browne, 2016] Cameron B. Browne. A class grammar for general games. In *Advances in Computer Games*, volume 10068 of *LNCS*, pages 167–182, Leiden, 2016.
- [Browne, 2017] Cameron B. Browne. Back to the past: Ancient games as a new AI frontier. In *AAAI 2017*, San Francisco, 2017.
- [Finnsson and Björnsson, 2008] Hilmar Finnsson and Yngvi Björnsson. Simulation-based approach to general game playing. In *The Twenty-Third AAAI Conference on Artificial Intelligence*, pages 259–264. AAAI Press, 2008.
- [Finnsson and Björnsson, 2010] Hilmar Finnsson and Yngvi Björnsson. Learning simulation control in general game-playing agents. In *The Twenty-Fourth AAAI Conference on Artificial Intelligence*, pages 954–959. AAAI Press, 2010.
- [Font *et al.*, 2013] Jose M. Font, Tobias Mahlmann, Daniel Manrique, and Julian Togelius. A card game description language. In *Applications of Evolutionary Computation - 16th European Conference, EvoApplications 2013, Proceedings*, volume 7835 *LNCS of Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, pages 254–263, 2013.
- [Genesereth *et al.*, 2005] Michael R. Genesereth, Nathaniel Love, and Barney Pell. General game playing: Overview of the AAAI competition. *AI Magazine*, 26(2):62–72, 2005.
- [Koriche *et al.*, 2017] Frédéric Koriche, Sylvain Lagrue, Éric Piette, and Sébastien Tabary. Constraint-based symmetry detection in general game playing. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 280–287, 2017.
- [Kowalski *et al.*, 2019] Jakub Kowalski, Mika Maksymilian, Jakub Sutowicz, and Marek Szykula. Regular boardgames. In *The Thirty-Third AAAI Conference on Artificial Intelligence*. AAAI Press, 2019.
- [Koza, 1992] John Koza. *Genetic Programming*. MIT Press, Massachusetts, 1992.
- [Love *et al.*, 2008] Nathaniel Love, Timothy Hinrichs, David Haley, Eric Schkufza, and Michael Genesereth. General game playing: Game description language specification, 2008.
- [Mahlmann *et al.*, 2011] Tobias Mahlmann, Julian Togelius, and Georgios N. Yannakakis. Modelling and evaluation of complex scenarios with the strategy game description language. In *2011 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 174–181, 2011.
- [Parlett, 2016] David Parlett. What’s a ludeme? In *Game Puzzle Design*, volume vol. 2, pages 83–86, 2016.
- [Pitrat, 1968] Jacques Pitrat. Realization of a general game-playing program. In *IFIP Congress (2)*, pages 1570–1574, 1968.
- [Rasmusen, 2007] Eric Rasmusen. *Games and Information: An Introduction to Game Theory, 4th ed.* B. Blackwell, 2007.
- [Schaul *et al.*, 2011] Tom Schaul, Julian Togelius, and Jürgen Schmidhuber. Measuring intelligence through games. *CoRR*, abs/1109.1314, 2011.
- [Schaul, 2014] Tom Schaul. An extensible description language for video games. *IEEE Transactions on Computational Intelligence and AI in Games*, 6(4):325–331, Dec 2014.
- [Schiffel and Thielscher, 2014] Stephan Schiffel and Michael Thielscher. Representing and reasoning about the rules of general games with imperfect information. *Journal of Artificial Intelligence Research*, 49:171–206, 2014.
- [Schreiber, 2016] Sam Schreiber. Games-base repository. <http://games.ggp.org/base/>, 2016.
- [Shaker *et al.*, 2013] Mohammad Shaker, Mhd Hasan Sarhan, Ola Al Naameh, Noor Shaker, and Julian Togelius. Automatic generation and analysis of physics-based puzzle games. *2013 IEEE Conference on Computational Intelligence in Games (CIG)*, pages 241–248, 2013.
- [Sironi and Winands, 2017] Chiara F. Sironi and Mark H.M. Winands. Optimizing propositional networks. In *Computer Games. Springer.*, pages 133–151, 2017.
- [Thielscher, 2011] Michael Thielscher. The general game playing description language is universal. In *Proceedings of the Twenty-second International Joint Conference on Artificial Intelligence, IJCAI-11*, pages 1107–1112, 2011.

[Thielscher, 2017] Michael Thielscher. GDL-III: A description language for epistemic general game playing. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 1276–1282, 2017.

A Proof for Theorem 1

Similar to Kowalski *et al.* [2019], we formalise a finite, deterministic, k -player game with perfect information as a tuple (k, T, ι, v) , where:

- $k \in \mathbb{N}$ indicates the number of players.
- T is a finite tree with:
 - Nodes S (also referred to as game states).
 - An initial state $s_0 \in S$ (the root node of T).
 - Terminal states $S_{ter} \subseteq S$ (leaf nodes of T).
 - A predecessor function $f : (S \setminus \{s_0\}) \mapsto S$, such that $f(s)$ denotes the parent of s in T .
- $\iota : (S \setminus S_{ter}) \mapsto \{0, \dots, k\}$ indicating which player has the control in a given state.
- $v : S_{ter} \mapsto \mathbb{R}^k$, such that $v(s)$ denotes the vector of payoffs for k players for any terminal state $s \in S_{ter}$.

This is equivalent to the formalisation of k -player extensive form games by Rasmusen [2007], excluding elements required only for non-determinism or imperfect information.

We prove that, given any arbitrary finite, deterministic game with perfect information as defined above, a Ludii game can be constructed such that there is a one-to-one mapping between states and state transitions between the original game and the Ludii game. The intuition of our proof is to construct a Ludii game where the game board is represented by a graph with an identical structure to the full game tree T . The Ludii game is played by moving a single token, placed on the “root node” in the initial game state, along the graph until a leaf node is reached. For any state z in the original game, there is a corresponding state s in the Ludii game such that the token is located on the vertex corresponding to the position of z in T . Note that explicitly enumerating the complete game tree as a graph is unlikely to be the most optimal representation for most games, but it demonstrates that Ludii is capable of representing all such games.

Definition 5. Let $\mathcal{D} = (k, T, \iota, v)$ denote a finite, deterministic, k -player game with perfect information as formalised above. We define a related Ludii game $\mathcal{G}(\mathcal{D}) = \langle Mode, Equipment, Rules \rangle$, where $Rules = \langle Start, Play, End \rangle$, such that:

- $Mode = \{p_0, p_1, \dots, p_k\}$, where all p_i for $i \geq 1$ correspond to the k different players. The nature player p_0 will remain unused in deterministic games.
- $Equipment = \{\langle c_0^t \rangle, \{\langle c_0^p, c_1^p \rangle\}\}$. The only container $c_0^t = \langle V, E \rangle$ is a graph with a structure identical to the tree T of the original game \mathcal{D} . Due to the structure of c_0^t being identical to the structure of T , we can uniquely identify a vertex $v(z)$ for any state $z \in T$ from the original game. For any such vertex – except for $v(s_0)$ – we can also uniquely identify an adjacent “parent” vertex

$p(v(z))$, such that $p(v(z)) = v(f(z))$; the parent of a vertex corresponds to the predecessor of the corresponding state in the original tree T .

- The *Start* rules are given by a list containing only a single action. This action creates the initial game state by placing the c_1^p token on the site $v(s_0)$ of c_0^t that corresponds to the root node of T .
- Let s denote any non-terminal Ludii game state, such that there is exactly one site $v(z)$ for which $what(s, \langle c_0^t, v(z), 0 \rangle) = c_1^p$. Let z denote the state in the original game that corresponds to the site $v(z)$. Let $g(z)$ denote the children of z in T . Given s , we define $Play(s)$ to return a set $\{A_i\}$ of lists of actions A_i , with one list of actions for every child node $z' \in g(z)$. Each of those lists contains two primitive actions; one that takes the token c_1^p away from $v(z)$ (replacing it with the “empty” token c_0^p), and a second action that places a new token c_1^p on the site $v(z')$ of c_0^t that corresponds to the child $z' \in T$.
- The end rules are given by $End = \{\langle Cond_i(\cdot), \vec{S}_i \rangle\}$. For any terminal game state $z_i \in S_{ter}$, let $v(z_i)$ denote the site in the graph c_0^t that corresponds to the position of z_i in T . We add a tuple $(Cond_i(\cdot), \vec{S}_i)$ to End such that $Cond_i(s)$ returns true if and only if $what(s, loc) = c_1^p$ for $loc = \langle c_0^t, v, 0 \rangle$, and $\vec{S}_i = v(z_i)$. Intuitively, we use a separate end condition for every possible terminal state $z_i \in S_{ter}$ in the original game \mathcal{D} , which checks specifically for that state by making sure the c_1^p token is placed on the matching vertex $v(z_i)$.
- Let s denote any non-terminal Ludii game state, such that there is exactly one site $v(z)$ for which $what(s, \langle c_0^t, v(z), 0 \rangle) = c_1^p$. Let z denote the state in the original game that corresponds to the site $v(z)$. Then, we define $mover(s) = \iota(z)$.

Lemma 1. Let $\mathcal{G}(\mathcal{D})$ denote a Ludii game constructed as in Definition 5. Every game state s that can be reached through legal gameplay in such a game has exactly one vertex $v \in c_0^t$ such that $what(s, \langle c_0^t, v, 0 \rangle) = c_1^p$, and $what(s, \langle c_0^t, u, 0 \rangle) = c_0^p$ for all other vertices $u \neq v$.

Intuitively, this lemma states that every game state reachable through legal gameplay has the c_1^p token located on exactly one vertex, and that all other vertices are always empty (indicated by c_0^p).

Proof. Let s_0 denote the initial game state. The *Start* rules are defined to place a single c_1^p token on $v(z_0)$, where z_0 denotes the initial state in the \mathcal{D} game, which means that the lemma holds for s_0 .

Let s denote any non-terminal game state for which the lemma holds. Then, the assumptions in Definition 5 for an adequate definition of $Play(s)$ are satisfied, which means that $\{A_i\} = Play(s)$ is a non-empty set of lists of actions, one of which must be selected by $mover(s)$. Every A_i is defined to take away the token c_1^p from the vertex it is currently at, and to place it on exactly one new vertex. This means that the lemma also holds for any successor state $\mathcal{T}(s, A_i)$, which proves the lemma by induction. \square

We are now ready to prove the main theorem from the paper:

Theorem 2. *Ludii is universal for the class of finite deterministic games with perfect information.*

Proof. Let \mathcal{D} denote any arbitrary game as formalised above, with a tree T . Let $\mathcal{G}(\mathcal{D})$ denote a Ludii game constructed as described in Definition 5. We demonstrate that for any arbitrary traversal through T , from s_0 to some terminal state $z_{ter} \in S_{ter}$, there exists an equivalent trial τ , as defined in Definition 4, in $\mathcal{G}(\mathcal{D})$. By “equivalent” trial, we mean that the sequence of states traversed is equally long, the order in which players are in control is equal, and the payoff vectors at the end are equal.

Let z_0, z_1, \dots, z_f denote any arbitrary line of play in the original game \mathcal{D} , such that z_0 is the initial game state, and $z_f \in S_{ter}$. By construction, the initial game state s_0 of $\mathcal{G}(\mathcal{D})$ has the token c_1^p placed on the vertex $v(z_0)$ corresponding to the root node of T . This means that we have a one-to-one mapping from z_0 to s_0 , where $what(s_0, \langle c_0^t, v(z_0), 0 \rangle) = c_1^p$.

Let z_i denote some non-terminal state in the sequence z_0, z_1, \dots, z_{f-1} , such that we already have uniquely mapped z_i to a Ludii state s_i where $what(s_i, \langle c_0^t, v(z_i), 0 \rangle) = c_1^p$. Lemma 1 guarantees that the assumptions required for Definition 5 to adequately define $Play(s_i)$ are satisfied. Furthermore, we know that $\iota(z_i) = mover(s_i)$, which means that the same player is in control. The definition of $Play(s_i)$ ensures that there is exactly one legal list of actions A_i such that $\mathcal{T}(s_i, A_i) = s_{i+1}$, where $what(s_{i+1}, \langle c_0^t, v(z_{i+1}), 0 \rangle) = c_1^p$ (note that z_{i+1} must be a successor of z_i in T). We pick this s_{i+1} to uniquely map to z_{i+1} .

By induction, this completes the unique mapping between sequences of states z_0, z_1, \dots, z_f and s_0, s_1, \dots, s_f , uniquely specifies the lists of actions A_i that must be selected along the way, and ensures that z_i is always mapped to a state s_i such that the c_1^p token is placed on $v(z_i)$. This last observation ensures that one of the *End* conditions in $\mathcal{G}(\mathcal{D})$ triggers for s_f , and that the correct payoff vector $\vec{S} = v(z_f)$ is selected. \square