

Constraint-Based Symmetry Detection in General Game Playing

Frédéric Koriche, Sylvain Lagrue, Éric Piette and Sébastien Tabary

CRIL, CNRS UMR 8188, Université d'Artois, France

{koriche, lagrue, epiette, tabary}@cril.fr

Abstract

Symmetry detection is a promising approach for reducing the search tree of games. In General Game Playing (GGP), where any game is compactly represented by a set of rules in the Game Description Language (GDL), the state-of-the-art methods for symmetry detection rely on a *rule graph* associated with the GDL description of the game. Though such rule-based symmetry detection methods can be applied to various tree search algorithms, they cover only a limited number of symmetries which are apparent in the GDL description. In this paper, we develop an alternative approach to symmetry detection in stochastic games that exploits constraint programming techniques. The minimax optimization problem in a GDL game is cast as a stochastic constraint satisfaction problem (SCSP), which can be viewed as a sequence of *one-stage* SCSPs. Minimax symmetries are inferred according to the microstructure complement of these one-stage constraint networks. Based on a theoretical analysis of this approach, we experimentally show on various games that the recent stochastic constraint solver MAC-UCB, coupled with constraint-based symmetry detection, significantly outperforms the standard Monte Carlo Tree Search algorithms, coupled with rule-based symmetry detection. This constraint-driven approach is also validated by the excellent results obtained by our player during the last GGP competition.

1 Introduction

The topic of *General Game Playing* (GGP) is to develop artificial agents capable of playing a wide variety of games, without any human intervention [Genesereth *et al.*, 2005]. In a GGP tournament, the rules of an unknown game are supplied to the agent and, after a short period of deliberation, the agent must be able to play this game effectively against other players. The game rules are represented in the *Game Description Language* (GDL) [Love *et al.*, 2008]. Basically, a GDL program is a set of first-order logical clauses specifying the initial state of the game, the legal moves of each player, the effects of these moves, and the scores obtained at terminal states.

Recent versions of GDL [Schiffel and Thielscher, 2014; Thielscher, 2016] are expressive enough to cover various classes of finite-horizon games including, in particular, *stochastic games* in which a distinguished “chance” player is used to simulate the probabilistic effects of joint actions.

Conceptually, any turn-based two-player zero-sum (deterministic or stochastic) game has an optimal, *minimax value function* that specifies the expected outcome of the game, for every possible state, under perfect play by all players. In theory, such games may be solved by recursively computing the value function in a search tree of size b^d , where b is the number of legal moves per state, and d is the number of moves needed to reach a terminal state. Yet, even for games of moderate size, exhaustive search is unfeasible in GGP tournaments, due to the short deliberation time allowed to players.

To this point, *symmetry detection* is a popular inference method for reducing exploration in a search space, by transferring knowledge between equivalent regions of this space. As games typically involve many equivalent states and moves, such similarities can be exploited for transferring the value function between nodes of the search tree. Furthermore, since most symmetry detection methods are implemented using graph automorphism algorithms [McKay and Piperno, 2014], the main component for inferring game symmetries is a *graphical structure* from which a permutation group is induced. Ideally, this permutation group should be as close as possible to the group of *minimax symmetries*, which preserve the optimal strategies of the game.

Most approaches to symmetry detection in GGP rely on some *rule graph* associated with the GDL representation [Kuhlmann and Stone, 2007; Schiffel, 2010; Zhang *et al.*, 2015]. Basically, a rule graph contains nodes for the terms, atoms, and rules of the GDL representation, together with edges connecting these symbols. Any automorphism of this rule graph is a symmetry that preserves the game semantics. However, only a restricted subset of game symmetries can be recognized from such automorphisms, due to the *syntactical bias* imposed by the GDL representation.

In this paper, we propose an alternative approach to symmetry detection in stochastic games, inspired from constraint programming techniques. A *Constraint Satisfaction Problem* (CSP) consists in a set of variables, each associated with a domain of values, and a set of constraints specifying allowed combinations of values for subsets of

variables. A CSP solution is an assignment of variables to values that satisfies all constraints. A common approach for detecting solution symmetries in CSPs is to use the *microstructure complement* of the network [Jégou, 1993; Cohen *et al.*, 2006]. This structure is a hypergraph, whose nodes are associated with variable-value pairs, and whose hyperedges joint sets of nodes which are “disallowed” by some constraint. Any automorphism of this hypergraph, called *constraint symmetry*, is a solution symmetry.

Since stochastic games are *sequential* decision problems, they generally cannot be reduced to CSPs. Instead, our approach is to encode a GDL program and its minimax objective function into a *Stochastic Constraint Satisfaction Problem* (SCSP) [Walsh, 2002; Tarim *et al.*, 2006; Hnich *et al.*, 2012]. This problem involves variables, domains and constraints, but the variables are partitioned into *decision* variables, encoding state descriptions and players’ actions, and *stochastic* variables, encoding the actions of the chance player. Any solution of the SCSP is a minimax strategy of the GDL game. By construction, such a network can be viewed as a sequence of *one-stage* SCSPs, each specifying a game turn. Our symmetry detection method uses the microstructure complement of these one-stage SCSPs, from which constraint symmetries form a subgroup of the group of minimax symmetries. Importantly, these microstructure complements can also be exploited for detecting symmetries between *approximate* minimax strategies, which are found by combining constraint propagation and bandit exploration.

We provide an experimental evaluation of this approach using the MAC-UCB algorithm for solving SCSPs [Koriche *et al.*, 2016]. We show on various games that MAC-UCB, coupled with constraint-based symmetry detection, significantly outperforms Monte Carlo Tree Search algorithms, coupled with rule-based symmetry detection. We also emphasize that this approach was implemented in the “Woodstock” player, which has won the 2016 International GGP competition.

2 Games and Symmetries

2.1 Stochastic Games

The problems under consideration in this study are finite-horizon combinatorial games that incorporate probabilistic choices. A *game signature* consists in a tuple (I, F, A) , where $I = \{1, \dots, k\}$ is the set of *players*, F is a finite set of *fluents* used to describe game states, and A is a finite set of *actions* or *moves*. A *stochastic game* over the signature (I, F, A) is a tuple $\mathcal{G} = (s_1, S_\dagger, L, P, u)$, where:

- $s_1 \in 2^F$ is the *initial state*, and $S_\dagger \subseteq 2^F$ is the set of terminal states,
- $L : I \times 2^F \rightarrow 2^A$ maps each player i and each state s to the set $L_i(s)$ of *legal moves* of i at s ,
- $P : 2^F \times A^k \times 2^F \rightarrow [0, 1]$ maps each state s and joint action $\mathbf{a} \in \mathbf{L}(s)$ to a probability distribution $P(\cdot | s, \mathbf{a})$ over 2^F , where $\mathbf{L}(s) = L_1(s) \times \dots \times L_k(s)$,
- $u : I \times S_\dagger \rightarrow \mathbb{R}$ maps each player i and each terminal state s to the utility $u_i(s)$ of i at s .

A *history* (or *finite play*) of length T in \mathcal{G} is a sequence of the form $(s_1, \mathbf{a}_1, s_2, \dots, s_{T-1}, \mathbf{a}_T, s_T)$ where $\mathbf{a}_t \in \mathbf{L}(s_t)$

and $P(s_{t+1} | s_t, \mathbf{a}_t) > 0$, for $t \in \{1, \dots, T-1\}$. \mathcal{G} is called a *T-horizon game* if every history of length T in \mathcal{G} includes a terminal state in S_\dagger , and conversely, every terminal state in S_\dagger is included in some history of length T in \mathcal{G} .

As usual, a pure stationary strategy is a map from states to actions. Of particular interest in this study is the *minimax* strategy for which the value function (for player i) is

$$V_i^*(s) = \begin{cases} \max_{a_i} \min_{a_{-i}} Q_i^*(s, \mathbf{a}) & \text{if } s \notin S_\dagger \\ u_i(s) & \text{otherwise} \end{cases}$$

$$\text{where } Q_i^*(s, \mathbf{a}) = \sum_{s' \in S} P(s' | s, \mathbf{a}) V_i^*(s') \quad (1)$$

and where a_i ranges over $L_i(s)$, and a_{-i} ranges over the projection of $\mathbf{L}(s)$ onto $I \setminus \{i\}$. This strategy is optimal for the important class of *turn-based two-player zero-sum* stochastic games [Condon, 1992], given by $k = 2$, $u_1(s) = -u_2(s)$ for $s \in S_\dagger$, and $L_1(s) = \{\top\}$ or $L_2(s) = \{\top\}$ for $s \in S_G$, where \top is a distinguished “noop” action with no effects.

Evaluating the minimax value function in combinatorial games is generally intractable. A natural approach to alleviate this issue is to approximate V_i^* by a “depth-bounded” minimax value that limits the recursion to some fixed depth d , and applies a heuristic evaluation function when d is reached [Lanctot *et al.*, 2013]. Given a map $\hat{V}_i : S \rightarrow \mathbb{R}$ such that $\hat{V}_i(s) = u_i(s)$ when $s \in S_\dagger$, the *depth-d minimax* strategy is

$$V_i^d(s) = \begin{cases} \max_{a_i} \min_{a_{-i}} Q_i^d(s, \mathbf{a}) & \text{if } d > 0 \text{ and } s \notin S_\dagger \\ \hat{V}_i(s) & \text{otherwise} \end{cases}$$

$$\text{where } Q_i^d(s, \mathbf{a}) = \sum_{s' \in S} P(s' | s, \mathbf{a}) V_i^{d-1}(s'). \quad (2)$$

2.2 Game Symmetries

Intuitively, a symmetry in a set O of objects is a permutation over O that leaves some property of those objects unchanged. The particular group of symmetries that we obtain depends on the property we choose to preserve. In this paper, two main types of game symmetries are distinguished: *structural* symmetries that preserve the game structure, and *minimax* symmetries that preserve the minimax strategies of the game. Given a stochastic game \mathcal{G} with signature (I, F, A) , a structural symmetry of \mathcal{G} is a bijective function σ on $I \cup F \cup A$ such that for all $i \in I$, $f \in F$, $s, s' \in S_G$, and $\mathbf{a} \in A^k$,

$$\sigma(i) = i, \sigma(a) \in A, \sigma(f) \in F \quad (3a)$$

$$\sigma(f) \in s_1 \text{ iff } f \in s_1 \quad (3b)$$

$$\sigma(s) \in S_\dagger \text{ iff } s \in S_\dagger \quad (3c)$$

$$\sigma(a) \in L_i(\sigma(s)) \text{ iff } a \in L_i(s) \quad (3d)$$

$$P(\sigma(s') | \sigma(s), \sigma(\mathbf{a})) = P(s' | s, \mathbf{a}) \quad (3e)$$

$$u_i(\sigma(s)) = u_i(s) \quad (3f)$$

where $\sigma(s) = \{\sigma(f) | f \in s\}$, and $\sigma(\mathbf{a}) = (\sigma(a_1), \dots, \sigma(a_k))$. A minimax symmetry is a bijective function σ over $I \cup F \cup A$ satisfying conditions (3a-3c) together with

$$V_i^*(\sigma(s)) = V_i^*(s) \quad (4)$$

for all $i \in I$ and $s \in S_G$. By extension, *depth-d minimax symmetries* are defined by replacing (4) with

$$V_i^d(\sigma(s)) = V_i^d(s) \quad (5)$$

Any structural symmetry is a minimax symmetry, since the conditions (3d-3f) imply that two states with equivalent subgames must have the same minimax value. Yet, structural symmetries and depth- d minimax symmetries are generally incomparable, excepted for $d = T$ in T -horizon games.

3 Rule-Based Symmetry Detection

3.1 GDL Representations

In the GGP setting, games are represented by first-order stratified logic programs, using the GDL language that includes several keywords specifying the game dynamics. Stochastic games are captured by an extension of GDL that describes the chance player random [Schiffel and Thielscher, 2014].

Example 1. In the standard Tic-Tac-Toe game, two players 1 and 2 take turns by marking the cells in a grid until a row, column or diagonal is filled by a player. We focus here on a randomized variant of this game for the 2×2 grid; the players 1 and 2 can only choose a row or a column, and the exact position in the chosen line is drawn at random by the chance player. For instance, the rules

$$\begin{aligned} \text{next}(\text{row}(X, I)) &\leftarrow \text{does}(I, \text{chooseRow}(X)) \\ \text{next}(\text{col}(Y, I)) &\leftarrow \text{does}(I, \text{chooseCol}(Y)) \\ \text{next}(\text{cell}(X, Y, I)) &\leftarrow \text{does}(\text{random}, \text{mark}(X, Y, I)) \end{aligned}$$

specify the effects of players' actions. Namely, if a player I selects a row X (resp. a column Y) in the current state, then the fluent $\text{row}(X, I)$ (resp. $\text{col}(Y, I)$) is true in the next state, and if the chance player random marks a position (X, Y) with player I in the current state, then the fluent $\text{cell}(X, Y, I)$ is true in the next state. Alternatively, the rules:

$$\begin{aligned} \text{legal}(\text{random}, \text{mark}(X, Y, I)) &\leftarrow \\ &\text{true}(\text{row}(X, I), \text{true}(\text{cell}(X, Y, *))) \end{aligned} \quad (6a)$$

$$\begin{aligned} \text{legal}(\text{random}, \text{mark}(X, Y, I)) &\leftarrow \\ &\text{true}(\text{col}(Y, I), \text{true}(\text{cell}(X, Y, *))) \end{aligned} \quad (6b)$$

capture the legal moves of the chance player: marking a cell (X, Y) with I is possible if the position (X, Y) is available (*), and the row X or column Y is chosen by player I .

For a set R of GDL rules, the *dependency graph* of R associates a node with each literal L in R , and an arc (L, L') if L occurs in the body of some rule with head L' .

Because GDL is Turing complete [Saffidine, 2014], even for "valid" programs [Love *et al.*, 2008], we focus here on a *propositional* fragment of this language (including random), denoted P-GDL, satisfying two properties:

(7a) the number of ground instances of functional terms and the size of clauses are *bounded*, and

(7b) the dependency graph of the program is *acyclic*.

Any P-GDL program R can be rewritten into an equivalent ground instance R_g of size polynomial in $|R|$. The ground terms in R_g are partitioned into a set $I^+ = \{0, 1, \dots, k\}$ of players, where 0 denotes the chance player, a set F of fluent terms, a set A of action terms, and a set U of utility

constants. Similarly, the rules in R_g can be partitioned into six categories:

$$\begin{aligned} &\text{role}(i) \\ &\text{init}(f) \\ &\text{terminal} \leftarrow L_1, \dots, L_q \\ &\text{legal}(i, a) \leftarrow L_1, \dots, L_{q'} \\ &\text{next}(f) \leftarrow \text{does}(0, a_0), \dots, \text{does}(k, a_k) \\ &\text{goal}(i, \nu) \leftarrow L_1, \dots, L_{q''} \end{aligned}$$

where $i \in I^+$, $\nu \in U$, $f \in F$ and $a, a_0, \dots, a_k \in A$. Furthermore, each L_j is an atom A_j or its negation $\text{not } A_j$, where A_j is an expression of the form $\text{true}(f)$ or $f_1 \neq f_2$.

The *Clark's completion* [Clark, 1978] of R_g is the set of propositional formulas R_c , where (i) each symbol not in R_g is replaced with \neg , (ii) each atom A occurring as a head of some rule is associated with its set B_A of bodies, i.e.

$$B_A = \{L_1 \wedge \dots \wedge L_q \mid A \leftarrow L_1, \dots, L_q \in R_g\}$$

and (iii) for each atom A , all rules with head A are replaced with a single formula $A \leftrightarrow \bigvee B_A$. For an atom A , we use $R_g \vdash A$ to denote that A is true in the (unique) stable model of R_g , and we use $R_c \models A$ to denote that $R_c \wedge \neg A$ is unsatisfiable. By Corollary 4.21 in [Ben-Eliyahu and Dechter, 1994], if R_g satisfies (7b), we must have $R_g \vdash A$ iff $R_c \models A$, that is, R_g and R_c are *logically equivalent*.

Let $\mathcal{G}(R) = (s_1, S_\dagger, L, P, u)$ be the stochastic game over the signature (I, F, A) , where $I = I^+ \setminus \{0\}$, and such that

$$\begin{aligned} f \in s_1 &\text{ iff } R_g \vdash \text{init}(f), \\ s \in S_\dagger &\text{ iff } R_g \cup \text{true}(s) \vdash \text{terminal}, \\ a \in L_i(s) &\text{ iff } R_g \cup \text{true}(s) \vdash \text{legal}(i, a), \\ u_i(s) = \nu &\text{ iff } R_g \cup \text{true}(s) \vdash \text{goal}(i, \nu), \text{ and} \\ P(s' \mid s, \mathbf{a}) &= \frac{1}{|L_0(s)|} \text{ if } s' \in \rho(s, \mathbf{a}), \text{ and } 0 \text{ otherwise.} \end{aligned}$$

Here, $\text{true}(s) = \{\text{true}(f) \mid f \in s\}$, and $\rho(s, \mathbf{a})$ is the set of states s' for which there is an action $a_0 \in L_0(s)$, such that for all $f' \in s'$, $R_g \cup \text{true}(s) \cup \{\text{does}(i, a_i)\}_{i=0}^k \vdash \text{next}(f')$. Thus, $P(\cdot \mid s, \mathbf{a})$ is a uniform distribution over the states updated from s , using the moves (a_0, \mathbf{a}) coupling the action profile \mathbf{a} of players, and all possible actions $a_0 \in L_0(s)$.

Example 2. In the ground instance R_g of the GDL program R describing the randomized Tic-Tac-Toe game for the 2×2 grid, the rules (6a) and (6b) are each replaced with 8 ground formulas, obtained by instantiating the variables $X \in \{x_1, x_2\}$, $Y \in \{y_1, y_2\}$, and $I \in \{1, 2\}$. Notably, in the ground formulas

$$\begin{aligned} \text{legal}(\text{random}, \text{mark}(x_1, y_2, 1)) &\leftarrow \\ &\text{true}(\text{row}(x_1, 1), \text{true}(\text{cell}(x_1, y_2, *))) \end{aligned} \quad (8a)$$

$$\begin{aligned} \text{legal}(\text{random}, \text{mark}(x_1, y_2, 1)) &\leftarrow \\ &\text{true}(\text{col}(y_2, 1), \text{true}(\text{cell}(x_1, y_2, *))) \end{aligned} \quad (8b)$$

$\text{row}(x_1, 1)$, $\text{col}(y_2, 1)$ and $\text{cell}(x_1, y_2, *)$ are fluent terms, and $\text{mark}(x_1, y_2, 1)$ is an action term. If we assume that in the current state s , the first player has chosen row x_1 , and the number of free positions in x_1 is y , then the probability distribution $P(\cdot \mid s, \mathbf{a})$ over the next states using the action profile $\mathbf{a} = (\text{mark}(x_1, y_2, 1), \top, \top)$ is simply $1/y$.




Figure 1: The ground rule graph of formulas (8a) and (8b).

3.2 Rule Graphs

For the propositional fragment of GDL examined in this study, we use a simplified version of the (enhanced) rule graph specified in [Schiffel, 2010]. Formally, the *ground rule graph* of a P-GDL program R is the colored graph $G = (N, E, \chi)$ over R_g , such that

- each ground term $t \in I^+ \cup F \cup A \cup U$ is mapped to a distinct node n_t ; if t is a player $i \in I^+$, then $\chi(n_t) = i$, if t is a utility value $v \in U$, then $\chi(n_t) = v$, if t is a fluent term in F , then $\chi(n_t) = \text{fluent}$, and if t is an action term in A , then $\chi(n_t) = \text{action}$;
- each atom A of the form $p(t_1, \dots, t_q)$ is mapped to a distinct node n_A with color $\chi(n_A) = p$, and a clique over $\{n_A, n_{t_1}, \dots, n_{t_q}\}$;
- each literal L of the form $\text{not } A$ is mapped to a distinct node n_L with $\chi(n_L) = \text{not}$, and an edge (n_L, n_p) ;
- each rule r of the form $A \leftarrow L_1, \dots, L_q$ is mapped to a distinct node n_r with color $\chi(n_r) = \text{rule}$, and a clique over $\{n_r, n_A, n_{L_1}, \dots, n_{L_q}\}$.

Proposition 1. *Let R be a P-GDL program and G be its ground rule graph. Then, any automorphism of G is a structural symmetry of $\mathcal{G}(R)$.*

Example 3. As depicted in Figure 1, the ground rule graph of the ground formulas (8a) and (8b) includes two similar four-cliques. An automorphism of this graph can be obtained by permuting the nodes r_1 , true_2 and $\text{row}(x_1, 1)$, with r_2 , true_3 and $\text{col}(y_2, 1)$, respectively.

Note that the group of automorphisms of G depends on how the game is described by R_g . For example, suppose that two actions a and a' have the same preconditions and effects, but R_g uses two rules $\text{legal}(i, a) \leftarrow A$ and $\text{legal}(i, a) \leftarrow \text{not } A$ for action a , and one fact $\text{legal}(i, a')$ for a' . Then, no automorphism in G can permute a and a' because the neighbors of n_a and $n_{a'}$ are not equivalent.

4 Constraint-Based Symmetry Detection

The key idea of this study is to translate GDL programs into stochastic constraint networks, from which the microstructure complement can be used for detecting various types of

symmetries. We focus on a slight generalization of the original SCSP model [Walsh, 2002] that incorporates conditional probability distributions over stochastic variables.

Formally, a *Stochastic Constraint Satisfaction Problem* is a tuple $\mathcal{N} = (X, Y, D, C, P, \theta)$, where

- $X = (x_1, \dots, x_n)$ is an ordered finite set of variables,
- $Y \subseteq X$ is the set of *stochastic variables*,
- $D = \{D(x_i)\}_{i=1}^n$ is a set of *finite domains* each associated with a variable in X ,
- C is a set of *constraints*,
- $P = \{P_y\}_{y \in Y}$ is a set of *conditional probability tables* each associated with a variable in Y , and
- $\theta \in [0, 1]$ is a threshold value.

Any variable in $X \setminus Y$ is called a *decision variable*. Given an ordered subset $Z = (z_1, \dots, z_m)$ of variables in X , we denote by $D(Z)$ the relation $D(z_1) \times \dots \times D(z_m)$. An *instantiation* on Z is a tuple $\mathbf{v} \in D(Z)$, also written as a set of pairs variable-value $\{(z_1, v_1), \dots, (z_m, v_m)\}$. The instantiation \mathbf{v} is *complete* if $Z = X$. For a subset $Z' \subseteq Z$, we use $\mathbf{v}|_{Z'}$ to denote the projection $\{(z_i, v_i) : z_i \in Z'\}$ of \mathbf{v} onto Z' .

As usual, each constraint $c \in C$ is defined over a set of variables $\text{scp}_c \subseteq X$, called the *scope* of c , and consists in a mapping from $D(\text{scp}_c)$ to $\{0, 1\}$. Any instantiation \mathbf{v} on scp_c for which $c(\mathbf{v}) = 1$ (resp. $c(\mathbf{v}) = 0$) is called *allowed* (resp. *disallowed*). Each conditional probability table P_y is defined over a set of variables $\text{scp}_y \subseteq X$ occurring *before* y in X , and maps every instantiation \mathbf{v} on scp_y to a probability distribution $P_y(\cdot | \mathbf{v})$ over $D(y)$.

The probability $P(\mathbf{v})$ and the consistency $C(\mathbf{v})$ of a complete instantiation \mathbf{v} are respectively given by

$$P(\mathbf{v}) = \prod_{y \in Y} P_y(\mathbf{v}|_{\{y\}} | \mathbf{v}|_{\text{scp}_y}) \text{ and } C(\mathbf{v}) = \prod_{c \in C} c(\mathbf{v}|_{\text{scp}_c})$$

An instantiation \mathbf{v} is *globally consistent* (GC) if it can be extended to a complete instantiation \mathbf{v}' such that $C(\mathbf{v}') = 1$.

A *policy* π for \mathcal{N} is a tree, in which nodes are labeled according to the ordering X ; the root is labeled with x_1 , and each child of a node x_i is labeled with x_{i+1} . Each outgoing edge of x_i is labeled with a value in $D(x_i)$; decision nodes have a unique child, and stochastic nodes y_i have $|D(y_i)|$ children. Each leaf is labeled with the value $C(\mathbf{v})$ of the complete instantiation \mathbf{v} connecting that leaf to the root. The *expected consistency* of a policy π is given by

$$C(\pi) = \sum_{\mathbf{v}} P(\mathbf{v})C(\mathbf{v})$$

where \mathbf{v} ranges over the root to leaf paths of π . A *solution* of \mathcal{N} is a policy π such that $C(\pi) \geq \theta$. Note that in the particular case where $\theta = 1$, π is a solution of \mathcal{N} iff every instantiation in π is GC.

Finally, \mathcal{N} is *T-stage SCSP* if X is partitioned into T stages, i.e. $X = (Z_1, Y_1, \dots, Z_T, Y_T)$, where (Z_1, \dots, Z_T) is a partition of $X \setminus Y$, and (Y_1, \dots, Y_T) is a partition of Y .

4.1 From GDL to SCSP

Based on the framework suggested in [Koriche *et al.*, 2016], any P-GDL program R can be encoded into a multi-stage SCSP, where each stage simulates a game turn. Formally, let R_g be the ground instance of R , where $I^+ = \{0, 1, \dots, k\}$ is the set of players, A is the set of actions, F is the set of fluents, and U is the set of utility values. Let R_c be the Clark's completion of R_g , and T be a positive integer. Then, the T -stage SCSP of R is the tuple $\mathcal{N}_{R,T} = (X, Y, D, C, P, \theta)$ such that $\theta = 1$, and

- $X = (Z_1, Y_1, \dots, Z_T, Y_T)$, where each Z_t includes a Boolean variable ω_t indicating the presence of a terminal state, a set of Boolean variables $\{f_{j,t}\}_{j=1}^n$, each associated with a distinct fluent in F , a set of variables $\{a_{i,t}\}_{i=1}^k$ with domain A , each associated with a distinct player i , and a set of variables $\{u_{i,t}\}_{i=1}^k$ with domain U each associated with a player i . Each Y_t includes a single variable $a_{0,t}$ with domain A , associated with random;
- $C = C_1 \cup \dots \cup C_T$, where C_1 includes constraints encoding the unary clauses over `init` in R_c , and for every stage t , C_t includes a constraint $c_{A,t}$ encoding the rule $A \leftrightarrow \bigvee B_A$ in R_c . Constraints over `legal` for the same player i are merged into a single constraint $\text{legal}_{i,t}$;
- $P = \{P_1, \dots, P_T\}$, where P_t encodes the conditional probability distribution of $a_{t,0}$ at stage t . The scope of P_t is given by the scope $\{f_{j,t}\}_{j=1}^m$ of the constraint $\text{legal}_{0,t}$ (minus $a_{0,t}$), and for each instantiation v on this scope, $P_t(\cdot | v)$ is the uniform distribution over the set $\text{legal moves } \{a \in A \mid \text{legal}_{0,t}(a, v) = 1\}$.

Based on this construction, we can observe that for each stage $X_t = (Z_t, Y_t)$, the scopes of the constraints terminal_t , $\text{legal}_{i,t}$ and $\text{goal}_{i,t}$, together with the scope of the table P_t , are all covered by the variables in X_t . Thus, the only component that connects a stage and its successor is the constraint next_t , which determines the values of fluents at turn $t + 1$, given the actions simultaneously played in turn t .

Example 4. Consider again the randomized Tic-Tac-Toe game for the 2×2 grid. By Clark's completion, the ground formulas (8a) and (8b) are merged into a single formula:

$$\begin{aligned} & \text{legal}(\text{random}, \text{mark}(x_1, y_2, 1)) \leftrightarrow \\ & \text{true}(\text{row}(x_1, 1)), \text{true}(\text{cell}(x_1, y_2, *)) \vee \\ & \text{true}(\text{col}(y_2, 1)), \text{true}(\text{cell}(x_1, y_2, *)) \end{aligned}$$

At each stage t , this formula is encoded into a $\text{legal}_{0,t}$ constraint with scope $\text{mark}_{0,t}$, $\text{row}(x_1, 1)_t$, $\text{col}(y_2, 1)_t$ and $\text{cell}(x_1, y_2, *)_t$. The relation of this constraint states that the value $(x_1, y_2, 1)$ is allowed for the action variable $\text{mark}_{0,t}$ iff the fluent variable $\text{cell}(x_1, y_2, *)_t$ is true, and at least one of the fluent variables $\text{row}(x_1, 1)_t$ and $\text{col}(y_2, 1)_t$ is true. Similarly, by Clark's completion, the ground formula

$$\text{next}(\text{cell}(x_1, y_2, 1)) \leftarrow \text{does}(\text{random}, \text{mark}(x_1, y_2, 1))$$

is replaced with

$$\text{next}(\text{cell}(x_1, y_2, 1)) \leftrightarrow \text{does}(\text{random}, \text{mark}(x_1, y_2, 1))$$

and then encoded at each stage t into a $\text{next}_{0,t}$ constraint which states the fluent variable $\text{cell}(x_1, y_2, 1)_{t+1}$ is true at turn $t + 1$ iff the value of $\text{mark}_{0,t}$ at turn t is $(x_1, y_2, 1)$.

4.2 Microstructure complements

We have now all ingredients in hand to extend the notion of *microstructure complement* [Cohen *et al.*, 2006] to stochastic constraint networks. Given an SCSP \mathcal{N} , the microstructure complement of \mathcal{N} is the hypergraph \mathcal{H} , where the set of nodes is the collection of all variable-value pairs $\{(x, v) \mid x \in X, v \in D(x)\}$, and the set of hyperedges is the union over all constraints $c \in C$ of the instantiations on scp_c which are disallowed by c . Thus, $v = \{(z_1, v_1), \dots, (z_k, v_k)\}$ is a hyperedge of \mathcal{H} iff $\{z_1, \dots, z_k\}$ is the scope of some constraint $c \in C$ for which $C(v) = 0$.

By extension, the microstructure complement of $\mathcal{N}_{R,T}$ is the colored hypergraph $\mathcal{H}_{R,T}$, wherein nodes and hyperedges are defined as above, and the coloring function χ of $\mathcal{H}_{R,T}$ maps each node (x, v) to a triplet, given as follows:

- if x is a fluent variable $f_{j,t}$, then $\chi(x, v) = (t, \star, \star)$,
- if x is a terminal variable ω_t , then $\chi(x, v) = (t, \star, v)$,
- if x is an action variable $a_{i,t}$ then $\chi(x, v) = (t, i, \star)$, and
- if x is a utility variable $u_{i,t}$ then $\chi(x, v) = (t, i, v)$,

where \star is an arbitrary symbol disjoint from $I^+ \cup U$. By construction, different colors are assigned to variables of different type, and to variables occurring at different stages. Any color-preserving automorphism of $\mathcal{H}_{R,T}$ is called a *constraint symmetry* of $\mathcal{N}_{R,T}$.

Proposition 2. *Let R be the P-GDL program of a T -horizon game. Then, any constraint symmetry of $\mathcal{N}_{R,T}$ is a structural symmetry of $\mathcal{G}(R)$.*

It is important to keep in mind that any solution policy of $\mathcal{N}_{R,T}$ is just a "legal strategy" that satisfies all game rules. In order to encode minimax policies, let $\mathcal{N}_{R,T}^*$ be the extension of $\mathcal{N}_{R,T}$ to the constraints $C^* = (C_1^*, \dots, C_T^*)$, where $C_t^* = C_t \cup \{V_{i,t}^*, Q_{i,t}^*\}_{i=1}^k$. Here, $V_{i,t}^*$ and $Q_{i,t}^*$ are used to encode the minimax strategy (1). Namely, the scopes of $V_{i,t}^*$ and $Q_{i,t}^*$ are respectively given by $\{f_{j,t}\}_{j=1}^n \cup \{u_{i,t}\}$ and $\{f_{j,t}\}_{j=1}^n \cup \{a_{i,t}\}_{i=1}^k \cup \{u_{i,t}\}$, where $u_{i,t}$ is the utility variable of player i . The state-value constraint $V_{i,t}^*$ maps any state s to

$$u_{i,t} = \begin{cases} \max_{a_i} \min_{a_{-i}} Q_{i,t}^*(s, \mathbf{a}) & \text{if } \omega_t = 0, \\ \text{goal}_{i,t}(s) & \text{otherwise.} \end{cases}$$

Correspondingly, the action-value constraint $Q_{i,t}^*$ maps any state s and action profile \mathbf{a} to

$$u_{i,t} = \sum_{a_0} P_t(a_0 \mid s, \mathbf{a}) V_{i,t+1}^*(s)$$

where a_0 ranges over all values in $D(a_{0,t})$ such that $\text{legal}_{0,t}(a_0, s) = 1$. By enforcing global consistency over $\mathcal{N}_{R,T}^*$, we must have $Q_{i,t}^*(s, \mathbf{a}) = V_{i,t}^*(s)$, which in turn implies that only optimal actions for player i are kept in the instantiations (s, \mathbf{a}) which are GC. The network $\mathcal{N}_{R,T}^d$ is defined analogously using the d -depth minimax strategy (2).

Proposition 3. *Let R be the P-GDL program of a T -horizon game. Then, any constraint symmetry of $\mathcal{N}_{R,T}^*$ (resp. $\mathcal{N}_{R,T}^d$) is a minimax (resp. d -depth minimax) symmetry of $\mathcal{G}(R)$.*




Figure 2: A sub-hypergraph of the micro-structure complement associated with the one-stage SCSP of Example 5.

Though theoretically interesting, the above result can be difficult to implement due to the large size of the microstructure complement of a T -stage SCSP. Fortunately, $\mathcal{N}_{R,T}^*$ can be viewed as a sequence of *one-stage* SCSPs, each associated with a restricted microstructure complement. Concretely, the t th one-stage SCSP of $\mathcal{N}_{R,T}^*$ is the restriction $N_{R,t}^*$ of $\mathcal{N}_{R,T}^*$ to the stage X_t , that is, $N_{R,t}^* = (X_t, Y_t, D, C_t^*, P_t, \theta)$. The corresponding microstructure complement is the projection $H_{R,t}^*$ of $\mathcal{H}_{R,T}$ onto the variables X_t , satisfying the following *backward consistency* condition: if v_t is a tuple on X_t which does not include any hyperedge in $H_{R,t}^*$, then there is an extension v of v_t on (X_1, \dots, X_t) which does not include any hyperedge in $H_{R,1}^* \cup \dots \cup H_{R,t}^*$. For the stochastic constraint network $\mathcal{N}_{R,T}^d$, the t th one-stage SCSP $N_{R,t}^d$ and its microstructure complement $H_{R,t}^d$ are defined in a similar way.

Proposition 4. *Let R be the P-GDL program of a T -horizon game. Then, for every $t \in \{1, \dots, T\}$, any constraint symmetry of $N_{R,t}^*$ (resp. $N_{R,t}^d$) preserves the minimax (resp. d -depth minimax) strategies of $\mathcal{G}(R)$.*

From a practical viewpoint, constraint symmetries of one-stage SCSPs can be exploited for reducing the search of minimax strategies: once an instantiation v_t on the stage X_t has been proved GC, automorphisms of $H_{R,t}^d$ can be applied to derive symmetric GC instantiations at this stage. Equivalently, if any extension of v_t has reached a dead end, then all symmetric instantiations of v_t at stage X_t can be ignored.

Example 5. Consider the randomized 2×2 Tic-Tac-Toe game after a couple of turns. During the first stage, player 1 selected `chooseRow(x1)`, and the chance player marked (x_1, y_1) with 1. Next, player 2 selected `chooseCol(y2)`, and the chance player marked (x_2, y_2) with 2. Figure 2 depicts the microstructure complement of the third one-stage SCSP. For the sake of clarity, only action variables of players 0 (random) and 1 are represented. We can here observe that all legal moves for player 1 are symmetrical. Indeed, the automorphism obtained by permuting the symbols “1” and “2” in the values of the action variables, yields a symmetry between rows x_1 and x_2 , and a symmetry between columns y_1 and y_2 . Analogously, the automorphism given by permuting the symbols “x” and “y” in those values, gives rise to a symmetry between row x_1 (resp. x_2) and column y_1 (resp. y_2).

5 Experiments

Based on our framework, we now present a series of experiments conducted on a cluster of Intel Xeon E5-2643 CPU 3.3 GHz with 64 GB of RAM and four threads under Linux. All problems used in experiments are turn-based two-player zero-sum games. Specifically, we focus on 25 representative games: 20 deterministic GDL games, and 5 stochastic GDL games (with random). For all games, we used the 2015 Tiltyard Open (International GGP Competition) setup: 180s for the *start clock* (deliberation time before the first turn) and 15s for the *play clock* (deliberation time per turn).

MAC-UCB-SYM. In order to implement our constraint-based symmetry detection method, we used (a variant of) the MAC-UCB algorithm developed in [Koriche *et al.*, 2016]. Any GDL description R is first translated into the Clark’s completion of its ground instance R_c , which is then encoded into a T -stage SCSP $\mathcal{N}_{R,T}$. The horizon T was fixed to 200. The goal of MAC-UCB is to find d -minimax strategies, that is, solution policies for the SCSP $\mathcal{N}_{R,T}^d$. The MAC (Maintaining Arc Consistency) technique searches solution policies up to depth d , and the UCB (Upper Confidence Bounds) stochastic bandit method estimates the value $\hat{V}(s)$ of states s at depth d . In what follows, MAC-UCB-SYM is the upgrading of MAC-UCB that incorporates symmetry detection using the microstructure complement of one-stage SCSPs. The choice of the depth d was derived by partitioning the deliberation time into three ratios $(r_{\text{MAC}}, r_{\text{UCB}}, r_{\text{SYM}})$, which correspond to the proportions of runtime allocated for MAC, FLAT-UCB, and symmetry detection, respectively. Based on a sensitivity analysis of MAC-UCB-SYM, we used the ratios (45%, 30%, 25%).

MAC-UCB-SYM satisfies the condition of backward consistency by iteratively expanding the partial instantiations of a policy tree rooted at the initial state s_1 . For a partial instantiation v on (X_1, \dots, X_{t-1}) which is consistent with all constraints in $(C_1^d, \dots, C_{t-1}^d)$, the algorithm stores in $H_{R,t}^d$ all hyperedges e_t such that $v \cup e_t$ is inconsistent with $N_{R,t}^d$. So, for any independent set v_t of $H_{R,t}^d$, the instantiation $v \cup v_t$ is guaranteed to be consistent with (C_1^d, \dots, C_t^d) .

We used the NAUTY algorithm [McKay and Piperno, 2014] for detecting automorphisms of microstructure complements. Concretely, MAC-UCB-SYM uses a hash table of 32Gb with a hash function specified in [Zobrist, 1990]. For each one-stage SCSP $N_{R,t}^d$, we store in the hash table the tuples v_t on X_t for which the value $Q^d(v_t)$ has been estimated by MAC-UCB.¹ When the algorithm expands a partial instantiation v on (X_1, \dots, X_{t-1}) with an independent set v_t of $H_{R,t}^d$, we detect whether v_t is symmetric to some tuple v'_t in the hash-table. If so, the value $Q^d(v'_t)$ is directly mapped to $Q^d(v_t)$ without exploring the whole subtree up to depth d .

GGP Competitors. MAC-UCB-SYM was compared to several General Game Players: the first player is the (multiagent version) of the UCT algorithm [Sturtevant, 2008], which is the state-of-the-art method for deterministic games. The

¹Here, $Q^d(v_t) = Q_{i,t}^d(s, \mathbf{a})$ where (s, \mathbf{a}) is the projection of v_t onto the fluent and action variables $\{\mathbf{f}_{j,t}\}_{j=1}^n \cup \{\mathbf{a}_{i,t}\}_{i=1}^k$ of X_t .

Game	Deterministic GDL games					
	MAC-UCB	UCT	UCT-SYM	GRAVE	GRAVE-SYM	SANCHO
Amazons torus 10×10	84.2 ($\pm 1.2\%$)	96.0 ($\pm 2.2\%$)	98.1 ($\pm 1.7\%$)	78.4 ($\pm 3.4\%$)	86.7 ($\pm 2.7\%$)	86.2 ($\pm 3.1\%$)
Breakthrough suicide	93.0 ($\pm 2.3\%$)	93.1 ($\pm 1.6\%$)	81.9 ($\pm 3.7\%$)	59.4 ($\pm 0.9\%$)	73.2 ($\pm 2.9\%$)	77.8 ($\pm 4.0\%$)
Chess	76.4 ($\pm 2.5\%$)	95.8 ($\pm 1.8\%$)	95.3 ($\pm 2.1\%$)	88.1 ($\pm 4.2\%$)	95.4 ($\pm 2.5\%$)	87.9 ($\pm 2.1\%$)
Connect Four 20×20	87.5 ($\pm 3.5\%$)	97.0 ($\pm 1.2\%$)	100.0 ($\pm 0.0\%$)	65.1 ($\pm 3.1\%$)	88.5 ($\pm 2.2\%$)	96.0 ($\pm 0.9\%$)
Connect Four Simultaneous	73.7 ($\pm 2.8\%$)	87.8 ($\pm 1.7\%$)	96.1 ($\pm 0.9\%$)	77.2 ($\pm 3.4\%$)	93.2 ($\pm 3.6\%$)	82.0 ($\pm 2.6\%$)
Copolymer with pie	73.9 ($\pm 1.5\%$)	87.2 ($\pm 2.3\%$)	93.3 ($\pm 0.5\%$)	80.2 ($\pm 1.6\%$)	91.6 ($\pm 1.8\%$)	77.9 ($\pm 3.6\%$)
Dots and boxes suicide	65.4 ($\pm 1.7\%$)	95.2 ($\pm 0.9\%$)	83.7 ($\pm 2.4\%$)	80.5 ($\pm 1.6\%$)	70.3 ($\pm 2.7\%$)	88.0 ($\pm 1.5\%$)
English Draughts	85.1 ($\pm 2.8\%$)	97.9 ($\pm 1.3\%$)	97.4 ($\pm 1.3\%$)	70.1 ($\pm 4.1\%$)	71.2 ($\pm 3.1\%$)	59.3 ($\pm 1.5\%$)
Free For All 2P	53.4 ($\pm 0.7\%$)	81.5 ($\pm 2.4\%$)	84.8 ($\pm 1.9\%$)	61.2 ($\pm 0.7\%$)	72.3 ($\pm 1.6\%$)	71.2 ($\pm 2.3\%$)
Hex	84.0 ($\pm 1.4\%$)	100.0 ($\pm 0.0\%$)	100.0 ($\pm 0.0\%$)	74.2 ($\pm 3.2\%$)	89.8 ($\pm 2.9\%$)	78.1 ($\pm 1.5\%$)
Knight Through	81.2 ($\pm 2.4\%$)	96.0 ($\pm 1.3\%$)	90.6 ($\pm 2.6\%$)	82.1 ($\pm 3.5\%$)	81.2 ($\pm 2.3\%$)	88.1 ($\pm 2.6\%$)
Majorities	84.4 ($\pm 2.6\%$)	95.7 ($\pm 1.6\%$)	96.8 ($\pm 1.4\%$)	81.2 ($\pm 2.6\%$)	92.3 ($\pm 2.3\%$)	87.2 ($\pm 3.2\%$)
Pentago	53.1 ($\pm 1.5\%$)	84.8 ($\pm 3.1\%$)	66.2 ($\pm 2.8\%$)	72.1 ($\pm 2.3\%$)	58.4 ($\pm 2.8\%$)	54.3 ($\pm 0.9\%$)
Quarto	54.9 ($\pm 1.6\%$)	87.1 ($\pm 3.4\%$)	65.6 ($\pm 2.6\%$)	65.2 ($\pm 3.2\%$)	63.8 ($\pm 1.6\%$)	57.9 ($\pm 2.3\%$)
Sheep and Wolf	74.8 ($\pm 3.2\%$)	92.5 ($\pm 2.4\%$)	94.6 ($\pm 0.9\%$)	64.4 ($\pm 3.7\%$)	63.2 ($\pm 3.6\%$)	62.1 ($\pm 1.5\%$)
Shmup	58.0 ($\pm 1.7\%$)	88.2 ($\pm 2.6\%$)	63.7 ($\pm 2.2\%$)	55.3 ($\pm 1.5\%$)	52.1 ($\pm 0.2\%$)	53.0 ($\pm 0.6\%$)
Skirmish zero-sum	85.4 ($\pm 2.5\%$)	98.7 ($\pm 0.4\%$)	100.0 ($\pm 0.0\%$)	77.7 ($\pm 1.6\%$)	95.8 ($\pm 2.3\%$)	67.4 ($\pm 4.2\%$)
TicTac Chess 2P	94.9 ($\pm 3.4\%$)	97.6 ($\pm 0.7\%$)	96.5 ($\pm 0.4\%$)	76.3 ($\pm 1.3\%$)	93.2 ($\pm 2.3\%$)	86.1 ($\pm 3.3\%$)
TTCC4 2P	84.4 ($\pm 2.3\%$)	97.3 ($\pm 1.2\%$)	97.2 ($\pm 2.1\%$)	62.4 ($\pm 2.3\%$)	85.7 ($\pm 3.1\%$)	65.8 ($\pm 4.1\%$)
Reversi Suicide	72.2 ($\pm 3.2\%$)	100.0 ($\pm 0.0\%$)	100.0 ($\pm 0.0\%$)	66.1 ($\pm 2.7\%$)	78.7 ($\pm 2.2\%$)	58.2 ($\pm 2.2\%$)
Stochastic GDL games						
Backgammon	92.1 ($\pm 2.7\%$)	98.0 ($\pm 0.6\%$)	96.1 ($\pm 1.4\%$)	70.2 ($\pm 4.5\%$)	86.8 ($\pm 3.9\%$)	100.0 ($\pm 0.0\%$)
Can't Stop	88.2 ($\pm 1.7\%$)	94.1 ($\pm 2.3\%$)	96.8 ($\pm 1.7\%$)	93.4 ($\pm 1.3\%$)	93.7 ($\pm 3.2\%$)	100.0 ($\pm 0.0\%$)
Kaseklau	73.5 ($\pm 3.6\%$)	73.4 ($\pm 1.3\%$)	72.1 ($\pm 0.9\%$)	67.4 ($\pm 2.8\%$)	60.2 ($\pm 3.2\%$)	88.1 ($\pm 2.6\%$)
Pickomino	75.4 ($\pm 1.8\%$)	74.3 ($\pm 2.4\%$)	82.4 ($\pm 2.8\%$)	87.1 ($\pm 3.0\%$)	95.6 ($\pm 1.0\%$)	92.1 ($\pm 2.9\%$)
Yahtzee	87.4 ($\pm 1.6\%$)	88.2 ($\pm 2.3\%$)	83.1 ($\pm 3.3\%$)	64.1 ($\pm 3.2\%$)	60.9 ($\pm 2.5\%$)	91.8 ($\pm 3.3\%$)

Table 1: Results of MAC-UCB-SYM against each GGP player.

second player is Cazenave’s GRAVE algorithm (2015), that implements the Generalized Rapid Action Value Estimation technique, a generalization of RAVE [Gelly and Silver, 2007]. The last player is SANCHO (version 1.61c), a Monte Carlo Tree Search method elaborated by S. Draper and A. Rose, which has won the 2014 GGP Competition. Rule-based symmetry detection methods were implemented for UCT and GRAVE, yielding the UCT-SYM player and GRAVE-SYM player, respectively. By analogy with MAC-UCB-SYM, UCT-SYM and GRAVE-SYM used NAUTY for detecting automorphisms in the ground rule graph, together with a hash table for exploiting symmetries. Based on a sensitivity analysis of these players, UCT-SYM (resp. GRAVE-SYM) used 20% (resp. 25%) of its deliberation time for symmetry detection.

Experiments were conducted on 66,000 game contests, namely, 300 matches for each deterministic game, and 1000 matches for each stochastic game, in order to consider the probabilistic effect of outcomes. For the sake of fairness, the role of players was exchanged during each match.

Results. In Table 1, each column reports the proportion of wins of MAC-UCB-SYM against the selected adversary. For example, the third row of the second column indicates that, for the Chess game, MAC-UCB-SYM wins 95.8% of contests against UCT with a standard deviation of 1.8%.

In light of these results, the following observations can be made: (i) MAC-UCB-SYM outperforms MAC-UCB especially for large games (ex: Amazons torus 10×10), which indicates that constraint-based symmetry detection is an effective method for reducing exploration in large search trees; (ii) MAC-UCB-SYM also dominates UCT, GRAVE, and SANCHO, even for deterministic games which is the class targeted by these algorithms; and (iii) MAC-UCB-SYM is even better against UCT-SYM and GRAVE-SYM for most games. The last point clearly indicates that constraint-based symmetry methods, used to recognize (approximate) minimax symmetries, are more effective than rule-based symmetry detection methods, used to recognize structural symmetries.

6 Conclusion

In this paper, we have shown that constraint-based symmetry detection can prove effective for solving GDL games. Beyond the conclusive results obtained by coupling the stochastic constraint solver MAC-UCB with constraint-based symmetry detection, our framework emerges as a flexible approach for recognizing various types of game symmetries. Notably, a key perspective of further research is to detect *equilibrium symmetries* that preserve (mixed) Nash equilibria, or correlated equilibria, using a constraint-based approach for modeling these solution concepts.

References

- [Ben-Eliyahu and Dechter, 1994] Rachel Ben-Eliyahu and Rina Dechter. Propositional semantics for disjunctive logic programs. *Annals of Mathematics and Artificial Intelligence*, 12(1-2):53–87, 1994.
- [Cazenave, 2015] Tristan Cazenave. Generalized rapid action value estimation. In *Proceedings of IJCAI*, pages 754–760, 2015.
- [Clark, 1978] Keith L. Clark. Negation as failure. In *Logic and Databases*, pages 293–322. Plenum Press, 1978.
- [Cohen *et al.*, 2006] David Cohen, Peter Jeavons, Christopher Jefferson, Karen E. Petrie, and Barbara M. Smith. Symmetry definitions for constraint satisfaction problems. *Constraints*, 11(2-3):115–137, 2006.
- [Condon, 1992] Anne Condon. The complexity of stochastic games. *Information and Computation*, 96:203–224, 1992.
- [Gelly and Silver, 2007] Sylvain Gelly and David Silver. Combining online and offline knowledge in UCT. In *Proceedings of ICML*, pages 273–280, 2007.
- [Genesereth *et al.*, 2005] Michael Genesereth, Nathaniel Love, and Barney Pell. General game playing: Overview of the AAI competition. *AAAI Magazine*, 26(2):62–72, 2005.
- [Hnich *et al.*, 2012] Brahim Hnich, Roberto Rossi, S. Armagan Tarim, and Steven Prestwich. Filtering algorithms for global chance constraints. *Artificial Intelligence*, 189:69–94, 2012.
- [Jégou, 1993] Philippe Jégou. Decomposition of domains based on the micro-structure of finite constraint-satisfaction problems. In *Proceedings of AAI*, pages 731–736, 1993.
- [Koriche *et al.*, 2016] Frédéric Koriche, Sylvain Lagrue, Éric Piette, and Sébastien Tabary. General game playing with stochastic CSP. *Constraints*, 21(1):95–114, 2016.
- [Kuhlmann and Stone, 2007] Gregory Kuhlmann and Peter Stone. Graph-based domain mapping for transfer learning in general games. In *Proceedings of ECML*, pages 188–200, 2007.
- [Lanctot *et al.*, 2013] Marc Lanctot, Abdallah Saffidine, Joel Veness, Christopher Archibald, and Mark H.M. Winands. Monte Carlo *-minimax search. In *Proceedings of IJCAI*, pages 580–586, 2013.
- [Love *et al.*, 2008] Nathaniel Love, Timothy Hinrichs, David Haley, Eric Schkufza, and Michael Genesereth. General game playing: Game description language specification. Technical Report LG-2006-01, Stanford University, 2008.
- [McKay and Piperno, 2014] Brendan D. McKay and Adolfo Piperno. Practical graph isomorphism, II. *Journal of Symbolic Computation*, 60(0):94 – 112, 2014.
- [Saffidine, 2014] Abdallah Saffidine. The game description language is Turing complete. *IEEE Transactions on Computational Intelligence and AI in Games*, 6(4):320–324, 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.
- [Schiffel, 2010] Stephan Schiffel. Symmetry detection in general game playing. In *Proceedings of AAI*, 2010.
- [Sturtevant, 2008] Nathan R. Sturtevant. An analysis of UCT in multi-player games. *International Computer Games Association Journal*, 31(4):195–208, 2008.
- [Tarim *et al.*, 2006] S. Armagan Tarim, Suresh Manandhar, and Toby Walsh. Stochastic constraint programming: A scenario-based approach. *Constraints*, 11(1):53–80, 2006.
- [Thielscher, 2016] Michael Thielscher. GDL-III: A proposal to extend the game description language to general epistemic games. In *Proceedings of ECAI*, pages 1630–1631, 2016.
- [Walsh, 2002] Toby Walsh. Stochastic Constraint Programming. In *Proceedings of ECAI*, pages 111–115, 2002.
- [Zhang *et al.*, 2015] Haifeng Zhang, Dangyi Liu, and Wenxin Li. Space-consistent game equivalence detection in General Game Playing. In *Proceedings of CGW*, pages 165–177, 2015.
- [Zobrist, 1990] Albert L. Zobrist. A new hashing method with application for game playing. *International Computer Games Association Journal*, 13(2):69–73, 1990.