

# Combinatorial Games with Incomplete Information\*

Junkang Li<sup>1,2</sup>, Bruno Zanuttini<sup>2</sup> and Véronique Ventos<sup>1</sup>

<sup>1</sup>NukkAI, Paris, France

<sup>2</sup>Normandie Univ.; UNICAEN, ENSICAEN, CNRS, GREYC, 14 000 Caen, France  
 junkang.li@nukk.ai, bruno.zanuttini@unicaen.fr, vventos@nukk.ai

## Abstract

Games with incomplete information model multi-agent interaction in which players do not have common knowledge of the game they play. We propose a minimal generalisation of combinatorial games to incorporate incomplete information, called combinatorial game with incomplete information (CGII). The most important feature of CGIIs is that all actions are public, which allows better visualisation of each player’s knowledge and incomplete information. To further motivate the study of this new formalism, we show that computing optimal strategies for CGIIs has the same computational complexity as for general extensive-form games.

## 1 Introduction

Game theory is a mathematical framework for studying multi-agent interactions. We focus on *extensive-form games* (EFG), in which the interaction between agents takes place sequentially, i.e. every agent takes turns to make a move. Prominent examples of such games are Chess and Go.

Of particular interest to us is the notion of games with *incomplete information*, which are games in which agents do not have common knowledge of the game they play. For instance, an agent does not know the number of participants in an auction, or how much these participants value the object to be sold; a Poker player does not see the cards in their opponent’s hidden hands, hence cannot know for sure the exact consequence (i.e. payoff) of calling and raising bets; a Bridge or Hearts player does not know the cards that their opponent can play during a trick since this depends on their hidden hand; etc.

The notion of (in)complete information is frequently confused with the one of (in)perfect information. Complete information describes situations in which the whole structure of a game (the number of players, the game tree, the information sets of each player, the owner of each node, the payoff for each player at each leaf node, etc.) is common knowledge among all the players of the game. On the other hand, perfect information is a more stringent requirement than complete

information. Not only the structure of the game is common knowledge, but all players have full observability and perfect recall of the history (which is essentially a record of every decision made by every player so far). In other words, players always know their exact position in the game tree when asked to make the next decision. To summarise, incomplete information is an example of imperfect information; see Faliszewski *et al.* [2016, Sec. 2.4.2].

We propose a new and minimal formalism for EFGs with incomplete information that we call combinatorial games with incomplete information (CGIIs). In such a game, Nature picks a world from a universe according to some common prior; each player may have different observability of this world. Then, the game proceeds sequentially, during which there is no chance factor and all moves by the players are publicly observable. This formalism is designed to be a minimal generalisation of the notion of combinatorial games (which are Boolean games of no chance and with perfect information; see Siegel [2013]) and to closely capture the epistemic aspect of games with incomplete information.

For such games, we are interested in knowing how much reward an agent or a team of agents can guarantee for themselves; this corresponds to the notion of maxmin value, well known in optimisation under uncertainty, in which we aim to ensure that the worst possible outcome is not too bad.

By design, our new formalism seems particularly restrictive when compared to general EFGs, where hidden actions and arbitrary chance nodes are allowed. However, we show that the computational complexity of computing optimal strategies (with respect to the maxmin value) for CGIIs is as hard as for EFGs, which allows concluding that the difficulty of playing games comes from incomplete information/knowledge alone, not from hidden actions or mid-game chance factors. This also justifies that restricting algorithmic studies to CGIIs is without loss of generality. We also give a construction to enforce coordination between players in CGIIs under the constraint of public actions, which allows modelling situations similar to concurrent actions.

## 2 Related Work

**Game theory.** The study of games with incomplete information was pioneered by Harsanyi [1967; 1968a; 1968b], who proposes a formalism to model games of incomplete

\*A long version with proofs of all claims is available at <https://hal.science/hal-04568854>.

information as EFGs with imperfect information. This formalism, called the Harsanyi model of incomplete information, introduces *types* of players, or equivalently, a universe of worlds for which each player has a potentially different partial observability (also called the Aumann model of incomplete information). For detailed and formal definitions, see textbooks on game theory, e.g. Maschler *et al.* [2020, Chapter 9].

Combinatorial games, the inspiration of our formalism of CGII, are studied in the field of combinatorial game theory, established in the 70s by two books by Conway [2000] and Berlekamp *et al.* [2001; 2003a; 2003b; 2004]. For recent advances in this field, see Nowakowski [1996; 2002; 2015], Albert and Nowakowski [2009], and Larsson [2019].

Our formalism is also inspired by Frank and Basin [1998], who, in order to model the card play phase of the card game Bridge, propose a game with public actions and one-sided incomplete information in which the opponent has complete information. Frank and Basin [2001] show that finding optimal pure strategies for these games is NP-complete. Ginsberg [2001] proposes the first exact algorithm for these games, and implements it for Bridge robots. Parallely, Chu and Halpern [2001] study a model of games with incomplete information with common payoffs, and only one round of concurrent interaction after Nature picks the world; they show that it is NP-complete to play such games optimally.

Like us, Kovarik *et al.* [2022] highlight the distinction between public and private actions. They also argue that this distinction, essential for recent search algorithms, is partially lost when we model sequential multi-agent interaction with EFGs, which do not *explicitly* tell whether an action is public or not. They propose an alternative model for stochastic games that makes this distinction prominent, and show how to transform such models to augmented EFGs and vice versa.

**Complexity of games.** Most work in the literature on the computational complexity of games concerns the complexity of finding Nash equilibria, especially for normal-form games [Gilboa and Zemel, 1989; Daskalakis *et al.*, 2009]. For more references, see Conitzer and Sandholm [2008], who also show that it is NP-complete to decide whether Nash Equilibria with certain natural properties exist.

Koller and Megiddo [1992], Koller *et al.* [1996], and von Stengel [1996] make seminal contributions to understanding the complexity of two-player zero-sum EFGs. They also give polynomial-time algorithms for computing behaviour maxmin strategies of EFGs with perfect recall, based on linear programming.

Maxmin for a team of players with common payoffs is called team maxmin equilibrium (TME) in the literature, and was first proposed by von Stengel and Koller [1997]. Basilio *et al.* [2017] and Celli and Gatti [2018] propose another notion called TMECor (“Cor” stands for “correlation”), which allows agents in the same team to access a correlation device in order to coordinate their mixed strategies. Building on these works, Gimbert *et al.* [2020] and Zhang *et al.* [2023] study the complexity of TME and TMECor, thereby yielding a relatively complete picture of the complexity of behaviour and mixed maxmin for two-team EFGs.

The complexity of other models of decision making have also been extensively studied, e.g. Markov decision process [Mundhenk *et al.*, 2000; Bernstein *et al.*, 2002; Goldsmith and Mundhenk, 2007], propositional planning [Rintanen, 2004], graph games [Chatterjee and Henzinger, 2012; Chatterjee *et al.*, 2013]. Similar to these works, we confirm the intuition that partial observability and multi-agent coordination increases the difficulty of optimal decision making.

## 3 Combinatorial Games with Incomplete Information

### 3.1 Definitions

Combinatorial games are EFGs of no chance and with perfect information. To generalise this formalism minimally to allow incomplete information, we propose the following definition.

**Definition 3.1** (CGII). *A combinatorial game with incomplete information (CGII) is a tuple of the following elements:*

- *An Aumann model  $\langle U, A, (\mathcal{R}_i)_{i \in A}, \rho \rangle$ , where  $U$  is a finite set of worlds called universe,  $A$  is a set of agents,  $\mathcal{R}_i$  is an equivalence relation over  $U$  for each agent  $i \in A$ , and  $\rho \in \Delta(U)$  is a probability distribution over the universe called common prior;*
- *A tree  $T$  called public tree, the nodes of which  $(N(T))$  are partitioned into  $\{N_i(T)\}_{i \in A} \cup L(T)$ , where  $L(T)$  is the set of all leaves;*
- *A reward function  $u_i : L(T) \times U \rightarrow \mathbb{R}$  for each  $i \in A$ .*

Note that the children of a node (available actions at that node) do not depend on the real (and partially observable) world  $\omega$ ; only the rewards depend on  $\omega$ .

A CGII is said to be *Boolean* if all its reward functions have values in  $\mathbb{B}$ .<sup>1</sup> The Aumann model of a CGII defines each agent’s observability over the universe, which characterises their incomplete information.

**Pure strategies in a CGII.** A CGII as an EFG with incomplete information proceeds as follows. First, Nature picks the *real world*  $\omega \in U$  according to  $\rho$ . Then the *state game* in  $\omega$  proceeds from the root of the public tree  $T$ ; agents take turns to pick a child of the current node, depending on their equivalence class of the real world. This continues until a leaf  $l$  is reached, and agent  $i$  receives a payoff  $u_i(l, \omega)$ .

**Definition 3.2** (Pure strategy). *A pure strategy of an agent  $i \in A$  is a mapping  $s_i : N_i(T) \times U \rightarrow N(T)$  such that for all  $v \in N_i(T)$ :*

- *For all  $\omega \in U$ ,  $s_i(v, \omega)$  is a child of  $v$ ;*
- *$\forall \omega, \omega' \in U, \omega \mathcal{R}_i \omega' \implies s_i(v, \omega) = s_i(v, \omega')$ .*<sup>2</sup>

*The set of all pure strategies of agent  $i$  is denoted by  $\Sigma_i^P$ .*

From the definition of a strategy, one can see that the actions of every agent are indeed public: when making a decision at a node, an agent knows perfectly where the node is in

<sup>1</sup>In Boolean games, the rewards 0 and 1 are interpreted as a loss and a win, respectively.

<sup>2</sup>This means agent  $i$  must pick the same child for a node in any two worlds indistinguishable by them.




Figure 1: The public tree of a CGII and the game tree of its corresponding EFG, both with rewards omitted.

the public tree, which in particular means they observe and remember the actions picked by every agent in the past, starting from the root of the public tree. In addition, compared to general games with incomplete information, the state games of a CGII have the particularity that they share the same game tree  $T$ , which does not have chance nodes. These three features (public actions, unique game tree, and no chance) are the defining features of our formalism CGII.

Each CGII describes an EFG in which Nature picks the real world at the root, and information sets are determined by the players' observability of the world.

**Example.** Consider the following CGII: the public tree is shown in Figure 1 (on the left); the universe reads  $\{\omega_0, \omega_1\}$ ; agent 2 can distinguish these two worlds while agent 1 cannot. This CGII models a variant of Matching Pennies with incomplete information. The game tree of its corresponding EFG is also shown in Figure 1 (on the right).

On the right, since agent 1 cannot observe the real world, they must play H in both state games, or T in both. This constraint is respected by the notion of strategies in a CGII: on the left, agent 1 only has two pure strategies H and T since  $\omega_0$  and  $\omega_1$  are indistinguishable by agent 1.

Similarly, on the right, agent 2 can pick between heads or tails, depending on the choices of Nature and agent 1. On the left, agent 2 can again pick between heads or tails, depending on agent 1's choice and the real world, since agent 2 can distinguish between  $\omega_0$  and  $\omega_1$ ; note that the latter point is not reflected by the public tree, but by the Aumann model.

**Teams and Information in a CGII.** In a CGII, agents  $i$  and  $j$  are said to be in the same team if  $u_i = u_j$ .

**Definition 3.3 (Team).** A team is an inclusion-wise maximal group of agents with the same reward function.

We now define a team's degree of incomplete information.

- **Multi-agent incomplete information (MA-II):** an arbitrary team.
- **Single-agent incomplete information (SA-II):** a team of agents with the same equivalence relation (i.e.  $\mathcal{R}_i = \mathcal{R}_j$  for all agents  $i$  and  $j$  in the team).
- **Complete information (CI):** a team whose agents all have the finest equivalence relation (i.e.  $\mathcal{R}_i = \{(\omega, \omega) \mid \omega \in U\}$  for all agents  $i$  in the team).

In particular, CI implies SA-II, which implies MA-II. Intuitively, a team is a group of decentralised agents with shared

interests working cooperatively. In a CGII, a team with SA-II can be regarded as one single agent since every agent in this team has the same information and all actions are public.

**Example.** In the CGII in Figure 1, if the two agents have the same reward function, then they are in a team with MA-II; otherwise, each is a (single-agent) team with SA-II.<sup>3</sup>

Due to the public actions property, there is a close link between the degree of incomplete information of a team in a CGII and the degree of imperfect information of the corresponding team in the EFG defined by the CGII:

- a team with CI in the CGII is a player with perfect information in the EFG;
- a team with SA-II in the CGII can be seen as a single player with perfect recall in the EFG;
- a team with MA-II in the CGII is a team of players who all have perfect recall in the EFG.

This correspondence will allow us to establish upper bounds on the complexity of solving CGIIs.

**Team maxmin in a CGII.** Let  $(s_1, \dots, s_n) \in \Sigma_1^P \times \dots \times \Sigma_n^P$ , where  $n = |A|$ , be a pure strategy profile. We write  $(s_1, \dots, s_n)(\omega)$  for the unique leaf reached under this profile when the real world is  $\omega$ .

**Definition 3.4 (Expected utility).** The expected utility for an agent  $i \in A$  under a pure strategy profile  $(s_1, \dots, s_n)$  is defined to be:

$$U_i(s_1, \dots, s_n) = \sum_{\omega \in U} \rho(\omega) u_i((s_1, \dots, s_n)(\omega), \omega).$$

Let  $\mathcal{T} \subseteq A$  be a team. Notice that all agents in a team share the same expected utility function, which we denote by  $\mathcal{U}_{\mathcal{T}}$ . A pure strategy of the team is uniquely defined by the pure strategy of each of its players. In particular, the set of pure strategies of a team  $\mathcal{T}$ , denoted by  $\Sigma_{\mathcal{T}}^P$ , is in bijection with  $\prod_{i \in A} \Sigma_i^P$ . In the following, we also write  $\Sigma_{-\mathcal{T}}^P = \prod_{i \notin \mathcal{T}} \Sigma_i^P$ , the set of pure strategy profiles of the players not in  $\mathcal{T}$ .

**Definition 3.5 (Pure maxmin for a team).** The pure maxmin value for a team  $\mathcal{T} \subseteq A$  is defined to be

$$\underline{v}_{\mathcal{T}} := \max_{s_{\mathcal{T}} \in \Sigma_{\mathcal{T}}^P} \min_{s_{-\mathcal{T}} \in \Sigma_{-\mathcal{T}}^P} \mathcal{U}_{\mathcal{T}}(s_{\mathcal{T}}, s_{-\mathcal{T}}).$$

Intuitively, this value is the largest expected reward that a team can guarantee to get by playing a pure strategy.

The notion of behaviour/mixed strategy can be defined similarly to the one for EFGs: a mixed strategy of an agent  $i$  is a probability mixture of pure strategies of  $i$ ; a behaviour strategy of  $i$  picks, at each node and for each equivalence class of  $\mathcal{R}_i$ , a probability mixture of children (instead of just a child as for pure strategies). Hence, expected utility with respect to behaviour/mixed strategy profiles and behaviour/mixed maxmin for a team can be similarly defined.<sup>4</sup>

In the following, we focus on zero-sum two-team CGIIs. We call the two teams *player MAX* and *player MIN*, and denote them by  $+$  and  $-$ , respectively.

<sup>3</sup>The team of agent 2 even has CI.

<sup>4</sup>Behaviour maxmin and mixed maxmin are commonly known as TME and TMECor in the literature [Celli and Gatti, 2018].




Figure 2: Two EFGs for Matching Pennies.

### 3.2 Motivation for CGIIs

Our motivations for introducing CGII as a subclass of games with incomplete information are multiple. First and foremost, the formalism of CGII aims to be a minimal generalisation of combinatorial games to allow incomplete information. Indeed, it is clear that a CGII with a singleton universe is a combinatorial game. This new formalism allows modelling a number of card games, notably Bridge.<sup>5</sup>

But more importantly, the formalism of CGII also aims to minimally capture the notion of knowledge and incomplete information. Due to the public actions property, the only source of the imperfect (in particular, incomplete) information of every agent comes from their partial observability of the real world, drawn at the beginning of a game.

In contrast, we argue that the distinction between perfect and imperfect information does not completely capture the essence of players’ knowledge. For example, in the game Matching Pennies, MAX and MIN pick a side of a coin concurrently; this can be modelled by two different EFGs, shown in Figure 2. In the EFG on the left, MAX has perfect information while MIN has imperfect information but perfect recall; and in the EFG on the right, the situation is reversed. However, the roles of MAX and MIN are symmetric in Matching Pennies; MAX also has exactly the same information/knowledge in both EFGs when they need to choose an action. Hence, considering CGIIs allows one to focus on an unambiguous notion of knowledge of the players, as captured by the Aumann model and the initial drawing of a world.

**Expressiveness.** At first sight, the requirements of public actions and no chance seem particularly restrictive: many popular tabletop games with incomplete information allow private actions (e.g. concealed Kong in Mahjong, pass in Hearts) or have randomness and chance factors besides the initial drawing (e.g. dice rolls during a game). One may worry that, due to these restrictions, CGII is not expressive enough to be conceptually or algorithmically interesting. However, we argue that this impression is not correct.

First, an initial drawing over the universe is actually quite expressive. For example, for the dice rolls we evoke above, if their number and occasions are fixed in advance, then their results can be encoded into the initial drawing of worlds.<sup>6</sup> Another example is given by video games, which typically use a random seed as the sole source of randomness for all procedurally generated levels and random events during a

<sup>5</sup>The card play of Bridge can be described as a CGII in which MAX has SA-II and MIN has MA-II.

<sup>6</sup>This will only enlarge the game tree by a polynomial factor.

	MIN	CI	SA-II	MA-II
MAX		CI	<b>NP-c</b>	<b><math>\Sigma_2^P</math>-c</b>
	CI	P	<b>NP-c</b>	<b><math>\Sigma_2^P</math>-c</b>
	SA-II	NP-c	NP-c	<b><math>\Sigma_2^P</math>-c</b>
	MA-II	NP-c	NP-c	<b><math>\Sigma_2^P</math>-c</b>

Table 1: Complexity of PURE MAXMIN for CGIIs.

playthrough. Similar ideas have been investigated in automated planning [Palacios and Geffner, 2009, Sec. 10].

Second, even with only public actions, we show in Subsection 4.1 that we can still design a game to force a team of players to coordinate their actions. This means that we can essentially encode concurrent actions (as in standard Matching Pennies) using only public actions (and no chance except the initial drawing).

All in all, we suggest that at least as far as computation of optimal strategies is concerned, CGII, rather than EFG, be the right model for studying sequential multi-agent interactions depending on each player’s knowledge. Moreover, as we will show, CGIIs are as hard to solve as EFGs, which confirms our intuition that the difficulty of a game actually comes from the incomplete information of a player, and not from their inability to observe the moves made by the other players.

## 4 Complexity of PURE MAXMIN for CGIIs

The decision problem PURE MAXMIN is defined as follows.

**Definition 4.1** (PURE MAXMIN). *Let  $\mathcal{G}$  be a class of zero-sum CGIIs. Then PURE MAXMIN( $\mathcal{G}$ ) is the following decision problem.*

INPUT A CGII  $G \in \mathcal{G}$  and a rational number  $m$ .

OUTPUT Decide whether the pure maximin value for team MAX in  $G$  satisfies  $v_+(\Sigma_+^P, \Sigma_-^P) \geq m$ .

We study the complexity of PURE MAXMIN for CGIIs depending on the degrees of incomplete information for MAX and MIN: complete information (CI), single-agent incomplete information (SA-II), multi-agent incomplete information (MA-II). For complexity analyses, we consider the parameters  $|T|$  (number of nodes in the public tree),  $|U|$  (number of worlds), and possibly the number of bits to encode the utilities, the common prior, and the threshold  $m$ .

The complexity of PURE MAXMIN is summarised in Table 1. By definition, the complexity of each case is increasingly monotone in both MAX’s and MIN’s degree of incomplete information (CI/SA-II/MA-II). Hence, only a few hardness results have to be proved to establish the table. The results written in bold font are new from this work; the others can be directly deduced from the literature.

The membership results in Table 1 follow from results by Koller and Megiddo [1992, Sec. 3.3]; in particular, memberships in NP and in  $\Sigma_2^P$  follow from the fact that given a strategy of MAX, computing MIN’s best response is a problem in coNP, and even linear time when MIN has perfect recall.

Hence, we focus on hardness results. The following result is by Frank and Basin [2001, Sec. 6].<sup>7</sup>

<sup>7</sup>In their setting, there is no prior over the worlds; they are in-




Figure 3: The public tree of the coordination game.

**Proposition 4.2.** PURE MAXMIN is NP-hard for Boolean CGIIs in which MAX has SA-II and MIN has CI.

The symmetric case does not trivially follow from this result (since the minimax theorem does not hold for pure strategies) and necessitates a proof:

**Proposition 4.3.** PURE MAXMIN is NP-hard for Boolean CGIIs in which MAX has CI and MIN has SA-II.

*Proof sketch.* By a reduction from VERTEX COVER. Given a graph  $(V, E)$ , consider the universe  $U = \{\omega_e \mid e \in E\}$ ; the worlds are observable by MAX but not by MIN. During the game, MAX picks a vertex  $v \in V$ , then MIN picks an edge  $e' \in E$ . In a world  $\omega_e \in U$ , MAX gets a payoff of 1 if  $v$  covers  $e$  and MIN does not correctly guess this edge (i.e.  $e' \neq e$ ); otherwise, MAX gets 0. One can verify that the pure maxmin value of this game is at least  $1 - k/|E|$  if and only if the graph has a vertex cover of size at most  $k$ .  $\square$

#### 4.1 Multi-Agent Coordination in CGIIs

**Coordination game.** Now we turn our attention to CGIIs with multi-agent teams. We first show how to construct CGIIs to impose a perfect coordination between agents from the same team (à la Matching Pennies).

Consider the following Boolean CGII, which we call *coordination game*. This game has two agents of MAX, referred to as MAX 1 and MAX 2, and no agent of MIN; its universe has 4 worlds and reads  $U = \{(b_1, b_2) \mid b_1, b_2 \in \mathbb{B}\}$ ; its Aumann model has the uniform common prior and is such that for  $i = 1, 2$ , MAX  $i$  only observes  $b_i$ ; its public tree is shown in Figure 3; the reward for MAX is 1 if and only if  $a_1 \oplus b_1 = a_2 \oplus b_2$ , where  $\oplus$  is the exclusive or of two bits and  $a_i$  is the action chosen by MAX  $i$ .

We refer to  $b_i$  as the *hidden bit* of MAX  $i$  since it is only observable by MAX  $i$ . The coordination game is designed in such a way that MAX 1 and MAX 2 must perfectly coordinate their answer in order to win. Intuitively, MAX 1 and MAX 2 need to agree on the same answer  $A \in \mathbb{B}$ , then stick to it during the game by playing  $a_i = A \oplus b_i$ . Indeed, if they employ this strategy, then they guarantee a win since

$$a_1 \oplus b_1 = (A \oplus b_1) \oplus b_1 = A = (A \oplus b_2) \oplus b_2 = a_2 \oplus b_2.$$

**Remark.** Under these two winning strategies (one for each value of  $A$ ), both MAX 1 and 2 pick the actions 0 and 1 with equal probability. Indeed, once the common answer  $A$  is fixed, which action to play by MAX  $i$  is dictated by their hidden bit  $b_i$ . Hence, the bits  $b_1$  and  $b_2$  act as the keys of a

terested in the strategies that win in the greatest number of worlds. This is equivalent to finding maxmin strategies with respect to the uniform prior in our setting.

one-time pad to encrypt/mask the intended answer (i.e.  $A$ ) of MAX 1 and 2. This is the key element to ensure that MAX 1 and 2 must cooperate without cheating.

**Proposition 4.4.** In a coordination game, the only winning pure strategies of team MAX are of the following form: for some  $A \in \mathbb{B}$ , MAX 1 plays  $A \oplus b_1$  and MAX 2 plays  $A \oplus b_2$ .

*Proof.* Notice that the pure strategies of MAX 1 can be written in the form  $(a_1^0, a_1^1)$ , which means they choose  $a_1^0$  if  $b_1 = 0$  and  $a_1^1$  if  $b_1 = 1$ . As for MAX 2, they have the right to pick  $a_2$  as a function of  $a_1$  and  $b_2$ . If MAX 1 plays  $(A, A)$  (i.e. they play  $A$  regardless of  $b_1$ ) for some  $A \in \mathbb{B}$ , then MAX 2 has no winning strategy, since the winning condition  $A \oplus b_1 = a_2 \oplus b_2$  cannot be satisfied for both values of  $b_1$ . Now if MAX 1 plays  $(A, A \oplus 1)$  for some  $A \in \mathbb{B}$ , then to satisfy the winning condition, MAX 2 is forced to play  $a_2 = A \oplus b_2$ ; hence  $a_i = A \oplus b_i$ .  $\square$

The same reasoning also shows that these pure strategies are also the only winning behaviour strategies, and that the winning mixed strategies are exactly the mixtures of them.

**Remark.** From this proof, one can see that if MAX 1 cheats by using their hidden bit  $b_1$  incorrectly (i.e. does not use  $b_1$  to encrypt their intended answer and always picks the same action), then MAX 2 cannot cooperate perfectly since they cannot observe the value of  $b_1$ .

In addition, when MAX 1 plays correctly (i.e. chooses a strategy of the form  $(A, A \oplus 1)$ ), then MAX 2 must also pick  $A$  as their intended answer and mask it with their own bit  $b_2$  in order to win. Notice that in this case, the action picked by these two agents are uniformly and independently distributed. This is an important feature since agents (of MAX or MIN) in the following part of the game tree cannot deduce any information about the intended answer of these two agents by observing only their actions.

**Interrogation game.** We now generalise the coordination game to the following situation: we have a finite set of questions  $Q$ , and MAX has a Boolean answer for each question  $\{A_q \in \mathbb{B}\}_{q \in Q}$ , or equivalently a mapping from  $Q$  to  $\mathbb{B}$ . We wish to verify whether MAX's mapping satisfies some given binary constraints  $\{C_{qq'} \subseteq \mathbb{B}^2 \mid q, q' \in Q, q \neq q'\}$ : MAX's mapping is said to be *valid* if it satisfies all the constraints, that is,  $(A_q, A_{q'}) \in C_{qq'}$  for all  $C_{qq'}$ .

**Example.** For cliques of a given graph, the questions are the vertices of this graph; MAX's mapping induces a subgraph (MAX's answer to a vertex corresponds to whether to include this vertex in their intended subgraph); the binary constraints impose the requirement that all vertices in this subgraph be connected. Then MAX's mapping is valid if and only if it describes a clique of the graph.

To model this situation, consider the following Boolean CGII, which we call *interrogation game*: two agents of MAX (MAX 1 and MAX 2), and no agent of MIN; the universe reads  $U = \{(q_1, b_1, q_2, b_2) \mid q_1, q_2 \in Q, b_1, b_2 \in \mathbb{B}\}$ ; the Aumann model is such that for  $i = 1, 2$ , MAX  $i$  only observes  $q_i$  and  $b_i$ ; the common prior is uniform; the public tree is the same one as for the coordination game (i.e. the one

in Figure 3); MAX loses if and only if either (1)  $q_1 = q_2$  but  $a_1 \oplus b_1 \neq a_2 \oplus b_2$  or (2)  $q_1 \neq q_2$  but  $(a_1 \oplus b_1, a_2 \oplus b_2) \notin C_{q_1 q_2}$ .

This CGII has size  $\mathcal{O}(|Q|^2)$ : the universe has size  $\mathcal{O}(|Q|^2)$ , while the public tree has size  $\mathcal{O}(1)$ . Notice that a coordination game is just an interrogation game with only one question (hence no binary constraint). We refer to  $(q_i, b_i)$  as the *hidden information* of MAX  $i$ . Inspired by the coordination game, we propose the following definition.

**Definition 4.5** (Perfect coordination). *In an interrogation game, a perfect coordination of team MAX is a pure strategy of MAX of this form: there is a set  $\{A_q \in \mathbb{B}\}_{q \in Q}$  such that for all  $i$ , MAX  $i$  will play the action  $a_i = A_{q_i} \oplus b_i$  in all worlds in which their hidden information is  $(q_i, b_i)$ . For such a strategy, the set  $\{A_q\}_{q \in Q}$  is called the intended mapping or intended answer of the perfect coordination.*

By a similar argument to the one for the coordination game, the reward condition (1) ensures that MAX 1 and 2 have an incentive to implement a perfect coordination, which is a dominant strategy. In other words, (1) imposes non-adaptivity of MAX’s answers. As for the reward condition (2), it ensures that all binary constraints are satisfied by the intended mapping of a perfect coordination, since by (1) we have  $a_i \oplus b_i = A_{q_i}$  for all  $i$ . In summary, we have established the following result.

**Proposition 4.6.** *In an interrogation game, a pure strategy of team MAX is winning if and only if it is a perfect coordination with a valid intended mapping.*

It is straightforward to construct interrogation games involving team MIN such that if MIN does not cooperate, MAX receives a large reward. Similarly, we can also extend the construction above to allow  $k$ -ary constraints with  $k \geq 2$ . The interrogation game will then involve  $k$  agents of MAX, each with their hidden information  $(q_i, b_i)$ , and has size  $\mathcal{O}(2^k |Q|^k)$ . Such an interrogation game can be used to encode problems such as  $k$ -SAT.<sup>8</sup>

## 4.2 Hardness for Two-Team CGIIs

With the gadgets of interrogation game, it is straightforward to show that PURE MAXMIN is  $\Sigma_2^P$ -hard for CGIIs in which both MAX and MIN are multi-agent teams, for instance by a reduction from the canonical problem  $\exists \forall 3\text{SAT}$ . However, we provide a stronger result:  $\Sigma_2^P$ -hardness holds even when MAX has complete information.

**Proposition 4.7.** *PURE MAXMIN is  $\Sigma_2^P$ -hard for CGIIs in which MAX has CI and MIN has MA-II.*

*Proof sketch.* By a reduction from the  $\Sigma_2^P$ -complete problem SUCCINCT SET COVER [Umans, 1999]: given a collection of 3-DNF formulae and an integer  $k$ , decide whether there is a subset  $S$  of size at most  $k$  the disjunction of which is a tautology.

We design a game in which Nature draws a DNF formula from the collection, 3 variables, and 4 hidden bits, according to the uniform common prior. The DNF is known to MAX, who plays 1 or 0 according to whether it should be in  $S$ . This

<sup>8</sup>Contrastingly, we leave open the problem of constructing an interrogation game in which MAX’s answers are not binary.

	MIN	CI	SA-II	MA-II
MAX				
	CI	<b>P</b>	<b>P</b>	<b>coNP-c</b>
	SA-II	<b>P</b>	<b>P</b>	<b>coNP-c</b>
	MA-II	<b>NP-c</b>	<b>NP-c</b>	$\Sigma_2^P\text{-c}/\Delta_2^P\text{-c}$

Table 2: Complexity of BEHAVIOUR MAXMIN and of MIXED MAXMIN for CGIIs.

answer is masked (to MIN) by the hidden bit of MAX, as in a coordination game. Then MIN chooses either to verify the size of  $S$  or to verify that the disjunction of  $S$  is a tautology. The other 3 hidden bits are used in the latter verification: MIN plays an interrogation game over the 3 variables, with the constraint to falsify the disjunction of  $S$ .

Finally, since MAX is designed to have CI, they know the variables and the hidden bits of MIN. To ensure that MAX does not play a strategy that depends on MIN’s information, we introduce one additional agent of MIN whose role is to punish MAX whenever MAX plays such a strategy.  $\square$

**Remark.** *The construction shows something stronger:  $\Sigma_2^P$ -hardness holds even when MIN has joint complete information (i.e. if the agents of MIN could pool their information, then they would have complete information).*

## 5 Complexity of BEHAVIOUR MAXMIN and MIXED MAXMIN

The decision problems BEHAVIOUR MAXMIN and MIXED MAXMIN can be defined similarly to PURE MAXMIN, the only difference being that MAX can use behaviour/mixed strategies instead of just pure strategies.<sup>9</sup>

The complexity of BEHAVIOUR MAXMIN and MIXED MAXMIN is summarised in Table 2. Again, the complexity is increasingly monotone in both dimensions, and results written in bold font are new. The only case where the complexity differs between behaviour and mixed strategies is the case in which both MAX and MIN have MA-II; in this case, BEHAVIOUR MAXMIN and MIXED MAXMIN are respectively  $\Sigma_2^P$ - and  $\Delta_2^P$ -complete.

The membership results follow from those for EFGs, which are superclasses of CGIIs: the results for **P** are by Koller and Megiddo [1992, Sec. 3.5], and the others by Zhang *et al.* [2023, Appx. C].

Therefore, we only have to establish the hardness results when MAX and/or MIN have MA-II. We first adapt a reduction from 3-SAT by Chu and Halpern [2001].

**Proposition 5.1.** *Both BEHAVIOUR MAXMIN and MIXED MAXMIN are NP-hard for Boolean CGIIs in which MAX has MA-II and MIN has CI.*

<sup>9</sup>In our definition for all these decision problems, MIN only uses pure strategies, which is without loss of generality. Indeed, MIN is a team of agents with perfect recall, hence every MIN’s behaviour strategy has an equivalent mixed strategy [Maschler *et al.*, 2020, Theorem 6.11]. In addition, the best responses in mixed strategies are no better than the best responses in pure strategies due to the linearity of expected utility with respect to mixtures of strategies.

*Proof.* For a 3-CNF with  $N$  clauses, consider the following CGII. The universe consists of the clauses, which are observable by MAX 1 but not by MAX 2, and with the uniform prior. During the game, MAX 1 picks a variable, then MAX 2 observes this variable and picks a truth value. They win if and only if the variable picked by MAX 1 is in the clause picked by Nature, and the truth value picked by MAX 2 for this variable renders this clause true.

Since there is no agent of MIN in this game, playing behaviour or mixed strategies is no better than pure ones. It is also straightforward to verify that MAX can guarantee an expected payoff of 1 if the 3-CNF is satisfiable; otherwise, the maxmin value for MAX is at most  $1 - 1/N$ .  $\square$

Now, coNP-hardness for the symmetric case (when MAX has CI and MIN has MA-II) essentially follows from this result. For mixed strategies, the minimax theorem ensures that when switching the roles of MIN and MAX, and negating the utilities in the game from the proof of Proposition 5.1, the maxmin value for MAX is at least  $-(1 - 1/N)$  if the 3-CNF is unsatisfiable, and  $-1$  otherwise. The hardness for BEHAVIOUR MAXMIN follows from the fact that mixed maxmin and behaviour maxmin have the same value due to MAX's perfect recall.

**Proposition 5.2.** BEHAVIOUR MAXMIN is  $\Sigma_2^P$ -hard for CGIIs in which both MAX and MIN have MA-II.

*Proof sketch.* By a reduction from  $\exists\forall 3SAT$  (for a 3-DNF formula  $\varphi(x, y)$ , decides whether  $\exists x\forall y \varphi(x, y)$  holds) which is known to be  $\Sigma_2^P$ -hard [Schaefer and Umans, 2002]. Given such a formula, we construct a CGII with 3 agents of MAX and 3 agents of MIN. The worlds consist of one existential (resp. universal) variable and one hidden bit for each agent of MAX (resp. of MIN); the common prior is uniform; each agent only observes their variable and hidden bit. During the game, the agents of MAX take turns to choose between 0 and 1, then so do the agents of MIN. The total payoff for MAX is computed as follows: (1) an inconsistency among the agents of MAX (in the sense of an interrogation game) yields  $-N$  for MAX, where  $N$  is a large real number; (2) an inconsistency among the agents of MIN yields  $+N$  for MAX; (3) if at least one term in  $\varphi(x, y)$  is satisfied by the assignment picked by the agents of MAX and MIN, then MAX receives  $+1$ .

By choosing  $N$  large enough, agents of MAX have an incentive to perform a perfect coordination, and the same goes for agents of MIN. In particular, MAX has no incentive to play non-pure behaviour strategies, which would cause inconsistency to happen with a non-zero probability. It is then straightforward to verify that  $\exists x\forall y \varphi(x, y)$  holds if and only if MAX can guarantee an expected utility of at least  $+1/n^3$ , where  $n$  is the maximum between the number of existential variables and the number of universal ones.  $\square$

**Proposition 5.3.** MIXED MAXMIN is  $\Delta_2^P$ -hard for CGIIs in which both MAX and MIN have MA-II.

*Proof sketch.* By a reduction from LAST SAT (for a 3-CNF, decide whether the lexicographically maximum satisfying assignment has value 1 for the last variable), which is  $\Delta_2^P$ -hard [Krentel, 1988]. The construction is very similar to the last

proof. Given a 3-CNF, we write the variables as  $x_1, \dots, x_n$ , and we construct a CGII with 3 agents of MAX and 3 agents of MIN. The worlds consist of one variable and one hidden bit for each agent of MAX or MIN; the common prior is uniform; each agent only observes their variable and hidden bit. During the game, the agents of MAX take turns to choose between 0 and 1, then so do the agents of MIN. The total payoff for MAX is computed as follows: (1) an inconsistency among the agents of MAX or a clause violated by their assignment yields  $-2N$  for MAX, where  $N$  is a large real number; (2) an inconsistency among the agents of MIN or a clause violated by their assignment yields  $+N$  for MAX; (3) for the first agent of MAX (resp. of MIN), if their hidden variable and bit are  $x_k$  and  $b$ , and they pick  $1 \oplus b$ , then MAX receives  $+2^{n-k}$  (resp.  $-2^{n-k}$ ); (4) MAX receives a bonus  $+1$  if the variable  $x_n$  is assigned  $1 \oplus b_1^+$  by the first agent of MAX.

By choosing  $N$  large enough, both MAX and MIN have an incentive to perform a perfect coordination (which can be pure or mixed for MAX) with a satisfying assignment. Let  $x = (x_1, \dots, x_n)$  be the lexicographically maximum satisfying assignment (if there is no such assignment, then MAX is bound to get a large negative expected utility). If  $x_n = 1$ , then MAX can guarantee an expected utility of  $+1/n$  by choosing this assignment for their perfect coordination; the best MIN can do is to choose this assignment. If  $x_n = 0$ , MAX has an expected utility of at most 0 when MIN plays this assignment: MAX gets 0 by playing the same assignment, and possibly less when playing other satisfying (hence lexicographically smaller) assignments with a non-zero probability.  $\square$

## 6 Conclusion

We have proposed a new formalism for extensive-form games with incomplete information that we name combinatorial games with incomplete information. Compared to EFGs, CGIIs only have public actions and one chance node at the beginning of the game, thereby putting better emphasis on the aspect of incomplete information/knowledge of the players.

Apart from the conceptual simplicity, the interests in this new formalism are also justified by the complexity results. Indeed, all the upper bounds for CGIIs are provided by membership results that also hold for EFGs, while all the lower bounds, proven by hardness results, coincide with the upper bounds. In particular, for every degree of observability, CGIIs have the same complexity as EFGs.

We have also shown how to model binary concurrent actions to enforce multi-agent coordination in CGIIs. We leave to future work how to model other types of hidden actions, in particular non-binary concurrent actions. Future work also includes tightening the complexity results to show that hardness holds even for Boolean CGIIs with a minimum number of agents and distributed knowledge of the real world for each team; designing a generic polynomial transformation from an arbitrary two-team EFG into a CGII; and extending the study to general-sum multi-team CGIIs with respect to solution concepts that generalise maxmin (e.g. strategies to commit to [Letchford and Conitzer, 2010]). Algorithmic studies adapted to CGIIs will also be of interest, with the long-term goal to implement better AIs for games such as Bridge.

## References

- [Albert and Nowakowski, 2009] Michael H. Albert and Richard J. Nowakowski, editors. *Games of No Chance 3*, volume 56 of *Mathematical Sciences Research Institute Publications*. Cambridge University Press, 2009.
- [Basilico *et al.*, 2017] Nicola Basilico, Andrea Celli, Giuseppe De Nittis, and Nicola Gatti. Team-maxmin equilibrium: Efficiency bounds and algorithms. In Satinder Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI 2017)*, pages 356–362. AAAI Press, 2017.
- [Berlekamp *et al.*, 2001] Elwyn R Berlekamp, John H Conway, and Richard K Guy. *Winning Ways for Your Mathematical Plays, Volume 1*. CRC Press, 2nd edition, 2001.
- [Berlekamp *et al.*, 2003a] Elwyn R Berlekamp, John H Conway, and Richard K Guy. *Winning Ways for Your Mathematical Plays, Volume 2*. CRC Press, 2nd edition, 2003.
- [Berlekamp *et al.*, 2003b] Elwyn R Berlekamp, John H Conway, and Richard K Guy. *Winning Ways for Your Mathematical Plays, Volume 3*. CRC Press, 2nd edition, 2003.
- [Berlekamp *et al.*, 2004] Elwyn R Berlekamp, John H Conway, and Richard K Guy. *Winning Ways for Your Mathematical Plays, Volume 4*. CRC Press, 2nd edition, 2004.
- [Bernstein *et al.*, 2002] Daniel S. Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. The complexity of decentralized control of markov decision processes. *Math. Oper. Res.*, 27(4):819–840, 2002.
- [Celli and Gatti, 2018] Andrea Celli and Nicola Gatti. Computational results for extensive-form adversarial team games. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI 2018)*, pages 965–972. AAAI Press, 2018.
- [Chatterjee and Henzinger, 2012] Krishnendu Chatterjee and Thomas A. Henzinger. A survey of stochastic  $\omega$ -regular games. *J. Comput. Syst. Sci.*, 78(2):394–413, 2012.
- [Chatterjee *et al.*, 2013] Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. A survey of partial-observation stochastic parity games. *Formal Methods Syst. Des.*, 43(2):268–284, 2013.
- [Chu and Halpern, 2001] Francis C. Chu and Joseph Y. Halpern. On the np-completeness of finding an optimal strategy in games with common payoffs. *Int. J. Game Theory*, 30(1):99–106, 2001.
- [Conitzer and Sandholm, 2008] Vincent Conitzer and Thomas Sandholm. New complexity results about Nash equilibria. *Games Econ. Behav.*, 63(2):621–641, 2008.
- [Conway, 2000] John H. Conway. *On Numbers and Games*. A K Peters/CRC Press, 2nd edition, 2000.
- [Daskalakis *et al.*, 2009] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009.
- [Faliszewski *et al.*, 2016] Piotr Faliszewski, Irene Rothe, and Jörg Rothe. Noncooperative game theory. In Jörg Rothe, editor, *Economics and Computation, An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*, Springer texts in business and economics, pages 41–134. Springer, 2016.
- [Frank and Basin, 1998] Ian Frank and David A. Basin. Search in games with incomplete information: A case study using bridge card play. *Artif. Intell.*, 100(1-2):87–123, 1998.
- [Frank and Basin, 2001] Ian Frank and David A. Basin. A theoretical and empirical investigation of search in imperfect information games. *Theor. Comput. Sci.*, 252(1-2):217–256, 2001.
- [Gilboa and Zemel, 1989] Itzhak Gilboa and Eitan Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- [Gimbert *et al.*, 2020] Hugo Gimbert, Soumyajit Paul, and B. Srivathsan. A bridge between polynomial optimization and games with imperfect recall. In Amal El Fallah Seghrouchni, Gita Sukthankar, Bo An, and Neil Yorke-Smith, editors, *Proceedings of the Nineteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, pages 456–464. International Foundation for Autonomous Agents and Multiagent Systems, 2020.
- [Ginsberg, 2001] Matthew L. Ginsberg. GIB: imperfect information in a computationally challenging game. *J. Artif. Intell. Res.*, 14:303–358, 2001.
- [Goldsmith and Mundhenk, 2007] Judy Goldsmith and Martin Mundhenk. Competition adds complexity. In John C. Platt, Daphne Koller, Yoram Singer, and Sam T. Roweis, editors, *Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems (NIPS 2007)*, pages 561–568. Curran Associates, Inc., 2007.
- [Harsanyi, 1967] John C. Harsanyi. Games with incomplete information played by “Bayesian” players, Part I. The Basic Model. *Management Science*, 14(3):159–182, 1967.
- [Harsanyi, 1968a] John C. Harsanyi. Games with incomplete information played by “Bayesian” players, Part II. Bayesian Equilibrium Points. *Management Science*, 14(5):320–334, 1968.
- [Harsanyi, 1968b] John C. Harsanyi. Games with incomplete information played by “Bayesian” players, Part III. The Basic Probability Distribution of the Game. *Management Science*, 14(7):486–502, 1968.
- [Koller and Megiddo, 1992] Daphne Koller and Nimrod Megiddo. The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior*, 4(4):528–552, 1992.
- [Koller *et al.*, 1996] Daphne Koller, Nimrod Megiddo, and Bernhard von Stengel. Efficient computation of equilibria for extensive two-person games. *Games and Economic Behavior*, 14(2):247–259, 1996.



- [Kovarík *et al.*, 2022] Vojtech Kovarík, Martin Schmid, Neil Burch, Michael Bowling, and Viliam Lisý. Rethinking formal models of partially observable multiagent decision making. *Artif. Intell.*, 303:103645, 2022.
- [Krentel, 1988] Mark W. Krentel. The complexity of optimization problems. *J. Comput. Syst. Sci.*, 36(3):490–509, 1988.
- [Larsson, 2019] Urban Larsson, editor. *Games of No Chance 5*, volume 70 of *Mathematical Sciences Research Institute Publications*. Cambridge University Press, 2019.
- [Letchford and Conitzer, 2010] Joshua Letchford and Vincent Conitzer. Computing optimal strategies to commit to in extensive-form games. In David C. Parkes, Chrysanthos Dellarocas, and Moshe Tennenholtz, editors, *Proceedings of the Eleventh ACM Conference on Electronic Commerce (EC-2010)*, pages 83–92. ACM, 2010.
- [Maschler *et al.*, 2020] Michael Maschler, Eilon Solan, and Shmuel Zamir. *Game Theory*. Cambridge University Press, 2nd edition, 2020.
- [Mundhenk *et al.*, 2000] Martin Mundhenk, Judy Goldsmith, Christopher Lusena, and Eric Allender. Complexity of finite-horizon markov decision process problems. *J. ACM*, 47(4):681–720, 2000.
- [Nowakowski, 1996] Richard J. Nowakowski, editor. *Games of No Chance*, volume 29 of *Mathematical Sciences Research Institute Publications*. Cambridge University Press, 1996.
- [Nowakowski, 2002] Richard J. Nowakowski, editor. *More Games of No Chance*, volume 42 of *Mathematical Sciences Research Institute Publications*. Cambridge University Press, 2002.
- [Nowakowski, 2015] Richard J. Nowakowski, editor. *Games of No Chance 4*, volume 63 of *Mathematical Sciences Research Institute Publications*. Cambridge University Press, 2015.
- [Palacios and Geffner, 2009] Héctor Palacios and Hector Geffner. Compiling uncertainty away in conformant planning problems with bounded width. *J. Artif. Intell. Res.*, 35:623–675, 2009.
- [Rintanen, 2004] Jussi Rintanen. Complexity of planning with partial observability. In Shlomo Zilberstein, Jana Koehler, and Sven Koenig, editors, *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004)*, pages 345–354. AAAI, 2004.
- [Schaefer and Umans, 2002] Marcus Schaefer and Christopher Umans. Completeness in the polynomial-time hierarchy: A compendium. *SIGACT news*, 33(3):32–49, 2002.
- [Siegel, 2013] Aaron N Siegel. *Combinatorial game theory*. American Mathematical Society, 2013.
- [Umans, 1999] Christopher Umans. Hardness of approximating  $\Sigma_2^P$  minimization problems. In Paul Beame, editor, *Proceedings of the Fortieth Annual Symposium on Foundations of Computer Science (FOCS 1999)*, pages 465–474. IEEE Computer Society, 1999.
- [von Stengel and Koller, 1997] Bernhard von Stengel and Daphne Koller. Team-maxmin equilibria. *Games and Economic Behavior*, 21(1):309–321, 1997.
- [von Stengel, 1996] Bernhard von Stengel. Efficient computation of behavior strategies. *Games and Economic Behavior*, 14(2):220–246, 1996.
- [Zhang *et al.*, 2023] Brian Hu Zhang, Gabriele Farina, and Tuomas Sandholm. Team belief DAG: generalizing the sequence form to team games for fast computation of correlated team max-min equilibria via regret minimization. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the Fortieth International Conference on Machine Learning (ICML 2023)*, volume 202 of *Proceedings of Machine Learning Research*, pages 814–839. PMLR, 2023.