Learning to Identify Winning Coalitions in the PAC Model

Ariel D. Procaccia Jeffrey S. Rosenschein
School of Engineering and Computer Science
The Hebrew University of Jerusalem
Jerusalem, Israel
{arielpro, jeff}@cs.huji.ac.il

ABSTRACT

We consider PAC learning of *simple* cooperative games, in which the coalitions are partitioned into "winning" and "losing" coalitions. We analyze the complexity of learning a suitable concept class via its Vapnik-Chervonenkis (VC) dimension, and provide an algorithm that learns this class. Furthermore, we study constrained simple games; we demonstrate that the VC dimension can be dramatically reduced when one allows only a single minimum winning coalition (even more so when this coalition has cardinality 1), whereas other interesting constraints do not significantly lower the dimension.

Categories and Subject Descriptors

F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity;
I.2.6 [**Artificial Intelligence**]: Learning—*Concept Learning*;
I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems*

General Terms

Algorithms, Theory

Keywords

PAC Learning, Coalition Formation

1. INTRODUCTION

A significant portion of recent research in multiagent systems has focused on learning. Nevertheless, very few investigations have been devoted to learning in coalition formation, an area of game theory that is exceedingly relevant to multiagent systems. Classical models of coalition formation assume that the values of all coalitions are known *a priori*, but this assumption is unreasonable in many (indeed, almost all) multiagent settings. Surprisingly, even fewer researchers have attempted to apply PAC[1] (Probably Approximately Correct) learning theory to multiagent settings, although this model has been extensively studied by researchers in computational learning theory (readers who are not familiar with this model are urged to consult [4]).

In this paper, we endeavor to remedy these shortcomings by studying PAC learning of *simple* cooperative games, in which each coalition is either "winning" or "losing". A basic structure is induced by the assumption that if a coalition is winning, any coalition that contains it is also winning. Simple games are a suitable model for numerous n-person conflict situations.

The Vapnik-Chervonenkis (VC) dimension is a combinatorial measure of the "richness" of a concept class[2], and is proportional to the difficulty of learning the class. We define several concept classes, which characterize settings of learning simple games with different restrictions, and identify the complexity of learning these classes by calculating their VC dimensions.

2. THE MODEL

A *cooperative n-person game in characteristic form with side payments* is a pair $(N; v)$, where $N = \{1, 2, \dots, n\}$ is a set of players, and v is the *characteristic function*, which assigns a real number $v(C)$ to each *coalition* $C \subseteq N$. $v(C)$ is the *value* of C.

Simple games are games where each coalition has a value of either 1 or 0. A coalition C is said to be *winning* if $v(C) = 1$, and *losing* if $v(C) = 0$. 2^N, the powerset of N, is partitioned into $\mathscr{W}$, the set of winning coalitions, and $\mathscr{L}$, the set of losing coalitions. A standard assumption is that this partition satisfies:

$$[C_1 \in \mathscr{W} \wedge C_1 \subseteq C_2] \Rightarrow C_2 \in \mathscr{W}. \qquad (1)$$

We wish to find a way to identify winning coalitions. In other words, given a coalition, we would like to label it by "winning" or "losing", or equivalently, by 1 or 0, respectively. Essentially, we would like to learn a function from coalitions to $\{0, 1\}$.

The assumption given in Equation (1) allows us to easily represent such functions. Indeed, we say that $C \subseteq N$ is a minimum winning coalition if and only if

$$C \in \mathscr{W} \wedge \forall i \in C,\ (C \setminus \{i\}) \in \mathscr{L}.$$

[1]The PAC model is also known as the *formal* model.
[2]A concept class is a binary function whose domain is the set of possible samples.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...$5.00.

A simple game can be concisely represented by a set of minimum winning coalitions. Furthermore, a set of minimum winning coalitions may be regarded as a function from the set of coalitions to $\{0,1\}$: a coalition is winning (labeled by 1) if and only if it is a superset (in the weak sense) of one of the minimum winning coalitions.

Remark 1. Surprisingly, learning to identify minimum winning coalitions (with the above formalization in mind) is equivalent to learning monotone[3] DNF formulas. Indeed, one can associate players with variables, and coalitions with terms. A hypothesis consisting of a set of minimum winning coalitions is essentially a disjunction of terms. When learning minimum winning coalitions, the sample space is the space of all coalitions; these can be identified with assignments to variables, where a coalition C induces an assignment of 1 to the variables associated with the players in C, and 0 to all other players. This equivalence allows us to leverage some of the research on monotone DNF in order to prove certain results.

For the sake of completeness, we give a swift description of the VC dimension; more details can be found in [4].

Definition 1. Let $\mathcal{C} : X \to \{0,1\}$ be a concept class, $S = \{x_1, x_2, \ldots, x_m\} \subseteq X$, and let

$$\Pi_{\mathcal{C}}(S) = \{\langle h(x_1), h(x_2), \ldots, h(x_m)\rangle : h \in \mathcal{C}\}.$$

If $|\Pi_{\mathcal{C}}(S)| = 2^m$, then S is considered *shattered* by $\mathcal{C}$.

In other words, if S is shattered by $\mathcal{C}$, then $\mathcal{C}$ realizes all possible dichotomies on S.

Definition 2. The *Vapnik-Chervonenkis (VC) dimension* of a concept class $\mathcal{C}$, denoted VC-dim($\mathcal{C}$), is the size of the largest set S that is shattered by $\mathcal{C}$. If $\mathcal{C}$ shatters arbitrarily large sets, then VC-dim$(\mathcal{C}) = \infty$.

It is known that the VC dimension of a concept class almost completely characterizes the complexity of learning it in the PAC model. Therefore, we focus on calculating the VC dimension of several concept classes of interest.

3. LEARNING UNCONSTRAINED GAMES

In order to calculate the VC dimension of the concept class of minimum winning coalitions, when no restrictions are imposed on the game, we first need to address a natural combinatorial problem.

Given the numbers $\{1, 2, \ldots, n\}$, we would like to find a maximal antichain of subsets, i.e., a family of subsets such that for any two subsets, neither one is contained in the other. Finding an antichain of size $\binom{n}{\lfloor n/2 \rfloor}$ is easy: we simply choose all the subsets of size $\lfloor n/2 \rfloor$ (see Table 1). But can one do better? The theorem gives a negative answer.

Theorem 1 (Sperner's theorem). *Let $\mathcal{F}$ be a family of subsets of $\{1, 2, \ldots, n\}$, such that for all $A, B \in \mathcal{F}$: $A \nsubseteq B$. Then $|\mathcal{F}| \leq \binom{n}{\lfloor n/2 \rfloor}$.*

Now we can straightforwardly prove the main theorem.

[3]DNF formulas in which all the literals are positive.

			$\emptyset$		
	$\{1\}$	$\{2\}$	$\{3\}$	$\{4\}$	
{1,2}	**{1,3}**	**{1,4}**	**{2,3}**	**{2,4}**	**{3,4}**
	$\{1,2,3\}$	$\{1,2,4\}$	$\{1,3,4\}$	$\{2,3,4\}$	
			$\{1,2,3,4\}$		

Table 1: Subsets of {1,2,3,4}, sorted by size. A maximal antichain is constructed by choosing all subsets of size 2; Sperner's Theorem states that one cannot construct a larger antichain.

Theorem 2. *Let $\mathcal{C}^*$ be a concept class, in which each concept is a set of minimum winning coalitions. Then:*

$$\textit{VC-dim}(\mathcal{C}^*) = \binom{n}{\lfloor n/2 \rfloor}.$$

Proof. We first show that the VC-dimension is at least $\binom{n}{\lfloor n/2 \rfloor}$, by producing a set of size $\binom{n}{\lfloor n/2 \rfloor}$ which is shattered by $\mathcal{C}^*$. Consider the set S of all coalitions of size $\lfloor n/2 \rfloor$. Clearly, any dichotomy on S can be realized by the concept that contains as minimum winning coalitions exactly the coalitions in S that are labeled by 1.

On the other hand, the VC-dimension is at most $\binom{n}{\lfloor n/2 \rfloor}$. Indeed, consider a set S of more than $\binom{n}{\lfloor n/2 \rfloor}$ coalitions. By Theorem 1, there are two coalitions $C_1, C_2 \in S$ such that $C_1 \subseteq C_2$. The dichotomy in which C_1 is labeled by 1 and C_2 is labeled by 0 cannot be realized. □

Contract, given as Algorithm 1, is a learning algorithm for $\mathcal{C}^*$: it returns a set of minimum winning coalitions which is consistent with the given samples (assuming one exists).

Algorithm 1

procedure Contract($\{(C_i, y_i)\}_{i=1}^m$) ▷ $y_i \in \{0,1\}$ is the label of coalition C_i
  $\mathscr{W}^m \leftarrow \emptyset$ ▷ Set of minimum winning coalitions
  for $i = 1$ to m **do** ▷ All sample coalitions
    if $y_i = 1 \ \wedge \ \forall R \in \mathscr{W}^m, \ R \nsubseteq C_i$ **then**
      $\mathscr{W}^m \leftarrow \mathscr{W}^m \setminus \{R \subseteq N : C_i \subseteq R\}$
      $\mathscr{W}^m \leftarrow \mathscr{W}^m \cup \{C_i\}$
    end if
  end for
  return $\mathscr{W}^m$
end procedure

4. LEARNING CONSTRAINED GAMES

Veto games are cooperative games where any coalition with nonzero value contains a distinguished player, called the *veto* player. In some veto games, the inclusion of the veto player in a coalition is also a sufficient condition for the coalition to be winning, not just a necessary one; in this case, we say the veto player is a *dictator*. Some real-world n-person conflict situations are simple games with a dictator.

That being the case, in order to identify winning coalitions in a simple game with a dictator, it is enough to pinpoint this distinguished player. The set of minimum coalitions contains exactly one coalition of cardinality 1: the dictator.

The following proposition follows from [1].

PROPOSITION 3. *Let $\mathcal{C}_1$ be a concept class, in which each concept is a single minimum winning coalition of cardinality 1. Then:*

$$\text{VC-dim}(\mathcal{C}_1) = \lfloor logn \rfloor.$$

It is worthwhile to generalize our requirement of a single dictator player; we now assume there is a coalition W such that for all coalitions C, C is winning if and only if $W \subseteq C$. Since W is the only minimum winning coalition, the goal of the learning process in such games is to recognize W.

By the equivalence between minimum winning coalitions and monotone DNF, it is possible to derive the following proposition from [2].

PROPOSITION 4. *Let $\mathcal{C}_1^*$ be a concept class, in which each concept is a single minimum winning coalition. Then:*

$$\text{VC-dim}(\mathcal{C}_1^*) = n.$$

REMARK 2. Another possible generalization is having at most k minimum winning coalitions; denote this concept class by $\mathcal{C}_k^*$. An upper bound for this case is given by Schmitt [3]. In particular, if $k = O(1)$, then VC-dim$(\mathcal{C}_k^*) = O(n)$.

Proper simple games are another type of common constrained games. In such games it holds that

$$\forall S \subseteq N,\ S \in \mathscr{W} \Rightarrow N \setminus S \in \mathscr{L}. \tag{2}$$

However, this constraint does little to reduce the VC dimension, compared with unconstrained simple games.

PROPOSITION 5. *Let $\mathcal{C}_p^*$ be a concept class, in which each concept is a set of minimum winning coalitions in a proper simple game. Then:*

$$\text{VC-dim}(\mathcal{C}_p^*) \geq \binom{n-1}{\lfloor \frac{n-1}{2} \rfloor}.$$

PROOF. We must exhibit a set S of size $\binom{n-1}{\lfloor \frac{n-1}{2} \rfloor}$ which is shattered by $\mathcal{C}_p^*$. Let

$$S = \left\{ C \subseteq N :\ 1 \in C \wedge |C| = \left\lfloor \frac{n-1}{2} \right\rfloor + 1 \right\}.$$

It holds that the cardinality of S is as desired. Moreover, for any dichotomy on S, one may choose exactly the subsets labeled by 1 as the minimum winning coalitions: since the intersection of all subsets in S is not empty, the constraint (2) is not violated. □

Another popular constraint is the elimination of *dummy* players.

$$\forall i \in N \exists C \subseteq N \text{ s.t. } i \in C \wedge C \in \mathscr{W} \wedge (C \setminus \{i\}) \in \mathscr{L}. \tag{3}$$

Informally, for every player there is a winning coalition which cannot do without it. The elimination of dummy players also does not substantially reduce the VC dimension.

PROPOSITION 6. *Let $\mathcal{C}_d^*$ be a concept class, in which each concept is a set of minimum winning coalitions in a simple game with no dummy players. Then:*

$$\text{VC-dim}(\mathcal{C}_d^*) \geq \binom{n-1}{\lfloor \frac{n-1}{2} \rfloor}.$$

PROOF. We must exhibit a set S of size $\binom{n-1}{\lfloor \frac{n-1}{2} \rfloor}$ which is shattered by $\mathcal{C}_d^*$. Let

$$S = \left\{ C \subseteq N :\ 1 \notin C \wedge |C| = \left\lfloor \frac{n-1}{2} \right\rfloor \right\}.$$

It holds that the cardinality of S is as desired. Given a dichotomy on S, let S^+ be the set of coalitions in S that are labeled by 1, and let B be the set of players that are not members of any of the coalitions in S^+. The set of minimum winning coalitions is $S^+ \cup \{B\}$; the purpose of including B is to ensure that constraint (3) is not violated. It remains to show that this is a legitimate set of minimum coalitions in a proper simple game, which realizes the dichotomy. Since $1 \in B$, B is not a subset of any of the coalitions in S, and in particular is not a subset of any of the coalitions in $S \setminus S^+$ — so it is not the case that the addition of B to the set of minimum winning coalitions mistakenly labels a coalition in $S \setminus S^+$ by 1. Clearly, neither is B a superset of any of the coalitions in S^+; thus it is not the case that there are two winning coalitions such that one is a superset of the other. Since the coalitions in S^+ are included in the set of minimum winning coalitions, the dichotomy is realized. Moreover, constraint (3) is satisfied: every $i \in N$ is contained in one of the coalitions in S^+, or in B, and these are all *minimum* winning coalitions. □

5. CONCLUSIONS

Simple games have natural interpretations in multiagent systems. Such games can be concisely represented by the set of minimum winning coalitions; this set may be regarded as a function from the set of coalitions to $\{0, 1\}$.

We have shown that the VC-dimension of the concept class which contains sets of minimum winning coalitions is $\binom{n}{\lfloor n/2 \rfloor}$; this implies that in general, it is very difficult to learn to identify winning coalitions in the PAC model.

We have also discussed constrained simple games, and have proven that restricting the set of minimum winning coalitions to a single coalition, or a dictator player, greatly reduces the VC-dimension. These constrained games are thus efficiently learnable. Nevertheless, other popular constrained games, such as proper simple games and games with no dummy players, are almost as hard to learn as unconstrained games.

6. ACKNOWLEDGMENT

This work was partially supported by grant #039-7582 from the Israel Science Foundation.

7. REFERENCES

[1] N. Littlestone. Learning quickly when irrelevant attributes abound. *Machine Learning*, 2(4):285–318, 1988.

[2] T. Natschlager and M. Schmitt. Exact VC-dimension of boolean monomials. *Information Processing Letters*, 59:19–20, 1996.

[3] M. Schmitt. On the capabilities of higher-order neurons: A radial basis function approach. *Neural Computation*, 17(3):715–729, 2005.

[4] M. Vidyasagar. *A Theory of Learning and Generalization*. Spinger-Verlag, 1997.