

Effective Tag Mechanisms for Evolving Coordination

Matthew Matlock
matt-matlock@utulsa.edu

Sandip Sen
sandip@utulsa.edu

Mathematical & Computer Sciences Department
University of Tulsa
Tulsa, Oklahoma, USA

ABSTRACT

Tags or observable features shared by a group of similar agents are effectively used in real and artificial societies to signal intentions and can be used to infer unobservable properties and choose appropriate behaviors. Use of tags to select partners has been shown to produce stable cooperation in agent populations playing the Prisoner's Dilemma game. Existing tag mechanisms, however, can promote cooperation only if that requires identical actions from all group members. We propose a more general tag-based interaction scheme that facilitates and supports significantly richer coordination between agents. Our work is motivated by previous research that showed the ineffectiveness of current tag schemes for solving games requiring divergent actions. The mechanisms proposed here not only solves those problems but are effective for other general-sum games. We argue that these general-purpose tag mechanisms allow new application possibilities of multiagent learning algorithms as they allow an agent to reuse its learned knowledge about one agent when interacting with other agents sharing the same observable features.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Algorithms, Performance, Economics

Keywords

tags, Prisoner's Dilemma, coordination, evolution, learning, games

1. INTRODUCTION

Both humans and artificial agents have bounded cognitive capabilities and other resource limitations. To effectively respond to real environments, therefore, such agents need to

leverage their experience and extrapolate conclusions. Such inductive generalizations are not guaranteed to work or even produce preferable outcomes, e.g., social stereotypes used in human societies are often discriminatory. However, they form a core component of human reasoning which allows us to manage the scale and complexity of life's challenges [13, 28]. We are particularly interested in the kind of generalization which allows agents to reuse their knowledge gleaned from interaction with one agent in interactions with a similar agent.

Most of the research on learning in agents and multiagent systems involves learning to interact effectively with another agent [21, 33]. In single agent learning, one agent is privy to the past behavior of another agent and learns a model of the latter to interact effectively with that agent [6, 34]. In multiagent learning, two or more agents repeatedly interact with the goal of converging to some equilibrium policies with desirable properties [19]. An interesting approach related to this idea is that of learning by imitation wherein an agent mimics the behavior of a capable agent to improve its problem-solving performance [22]. Evolutionary and adaptive strategies used in this and other similar research seek to expand and study this imitation [3, 5].

The focus of current work in this field, however, is complementary in the sense that rather than simply adopting another agent's successful strategy we want to investigate how a successful strategy against one agent can be used against other opponents. Another viewpoint on this is found in the following question: if a strategy is found to be effective in interactions with one opponent, against which other agents could it also be effectively used? In that sense the current work is akin to the work on recommender systems [9, 16, 23, 29] where the goal is to identify groups of users who share common interests. Thus, if a new user can be accurately identified as belonging to an existing user group, recommendations that were useful for the existing group can be effectively reused for the new user. Though recommender systems have been an active area of research in recent years with high potential applications¹, the possible use of user-stereotypes have been recognized early in user modeling, adaptive interfaces, and information retrieval communities [4, 7, 20, 25, 24, 30, 35]

Unfortunately, very little work in multiagent systems has

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'07 May 14–18 2007, Honolulu, Hawai'i, USA.
Copyright 2007 IFAAMAS .

¹Web-based commercial recommender systems are profitable and even announce challenge problems with significant payoffs, e.g., Netflix recently announced a 1 million dollar prize for any system that performs significantly better than their recommendation engine.

been done in an effort to classify agents into groups with similar behavior. Such clustering in agent space can significantly reduce the problems faced by agents when learning to interact with other agents in an environment. The only relevant work in this area can be traced to John Holland’s proposal of “tags” as a primitive means of communication that can aid in the evolution of a group [12, 26]. Tags are observable features of individuals that do not influence their behavior but are shared due to common heritage. Both tags and behaviors are external manifestations of inherited genetic material. As a result, even though tags do not directly encode or constrain behavior, they correlate closely with behavioral similarity between individuals with shared ancestry. As agents sharing tags generally behave in a similar fashion, a strategy effective against one agent is likely to be effective against another agent with the same or similar tag! This can allow us to reuse the policy we learned against one agent against all agents belonging to the same tag cluster.

Whereas it is possible to study knowledge reuse in arbitrary multiagent environments, we will use stylized interactions represented as stage games [18]. There has been considerable recent research in multiagent systems involving one-shot or repeated interactions of stage games [8, 15]. In particular, the Prisoner’s Dilemma (PD) game (see Figure 1) has received widespread attention in both game theory and multiagent learning research. In this game, the only Nash Equilibrium is the strategy profile (D,D) which is also the only non-Pareto-optimal outcome! The (D,D) strategy profile is dominated by the (C,C) strategy profile. Unfortunately, in a single-shot PD game, rational play will produce the Nash Equilibrium strategy profile. In repeated or iterated play, however, learning approaches can produce higher payoff by choosing the (C,C) strategy profile. Numerous researchers in game theory and multiagent systems have attempted various mechanisms to induce cooperation in iterated PDs [2, 31, 32].

	C	D
C	R, R	S, T
D	T, S	P, P

Figure 1: Utilities to players in a two-player Prisoner’s Dilemma game. Constraints on the utility values are $T > R > P > S$ and $2R > T + S > 2P$.

We are particularly interested in recent work using tags in a population of interacting agents [10, 11]. Tags have also been used by other researchers to promote cooperation in variations of PDs [26, 27]. Whereas these papers provided a reasonable high-level explanation of how the use of tags promotes cooperation, a detailed analysis that clearly explains the fundamental subtleties of the interactions in the population was missing. As a result, design of tag systems was based on trial and error and did not explain why certain parameter choices for such systems succeeded in inducing cooperation whereas others did not. More recently, McDonald and Sen [17] claimed that current tag mechanisms can promote cooperation only when cooperation can be achieved by imitation of behavior. Though McDonald and Sen were able to provide a systematic explanation of how tag-based cooperation evolved in the iterated PD, they failed to evolve cooperation in simple congestion games, e.g., the “anti-coordination” (AC) game (Figure 2), where

agents received high-payoff only if they chose different actions. The claim was that imitation does not produce high payoff in this game, and hence cooperation (or coordination) could not be sustained.

	0	1
0	L, L	H, H
1	H, H	L, L

Figure 2: Utilities to players in Anti-Coordination. Constraints on the utility values are $H > L$.

We believe that though McDonald and Sen raise an interesting point, a richer framework and interpretation of tags will address the limitations they observe. The current tag mechanisms are constrained to be self-matching types. Thus an agent only interacts with other agents with identical or similar tags. To extrapolate such constraints to human societies, this would mean that we limit our interactions to people with whom we share some external features, e.g., fashion choices. Though some people do limit themselves to small groups or clans who dress and behave similarly, most of us interact with a much larger and more diverse group of people. In some real-world situations, however, we do reuse knowledge learned while interacting with one individual to interact with another individual who belongs to the same group or cluster, e.g., knowledge of certain social norms, religious beliefs, or even habits of one person can be used to predict the preferences and actions of another person. We then routinely choose to interact with members of groups other than our own. In some cases, we might even be better off partnering with members of a different “group”, e.g., when complementary resources, capabilities, knowledge, etc. are required.

In multiagent domains in general, cooperation requires a richer, more diverse collection of behaviors than imitation. It was unclear from the current state of knowledge whether self-matching tags can support cooperation in a broad spectrum of multiagent problems. From the discussion above, however, it seems that it would be a natural next step to generalize tag matching to allow for interaction between individuals with different tags, i.e., belonging to different social groups. While there has been isolated work in evolutionary cooperation in PD by dividing the population into clans who employ different strategies for playing with agents within and outside their clans [14], there has been no prior work of inter-tag-group matching for evolving coordination. We propose a novel inter-tag-group matching mechanism that allows agents to coordinate effectively with members of other tag groups. We introduce the concept of matching tag patterns and other auxiliary mechanisms to facilitate the evolution of coordination and evaluate these schemes on the anti-coordination and prisoners dilemma games.

2. TAG-BASED EVOLUTION OF COOPERATION

The use of tags to bias interactions in agent populations playing the Prisoner’s Dilemma game has been suggested by [1, 12]. Riolo [26] modeled agents with a stochastic strategy, based on Tit-For-Tat, combined with a real-valued tag and a real-valued bias, both on the interval $[0, 1]$. Agents then attempt to pair up, where the difference between the

agents’ tags is less than each agents’ bias. If no suitable pairing can be found within a small number of trials, the agent simply chooses a partner at random. When each agent has an identical, fixed bias, Riolo’s model results in increased performance for the society. However, when each agent’s bias is allowed to evolve, behavior varies drastically according to initial conditions and the results are less clear.

Hales and Edmonds [10] used a different model where the population consists of a collection of agents represented as a binary string of $l + 1$ bits. The first bit represents a pure strategy (always cooperate or always defect), while the remaining l bits are the tag. In each population generation, every agent plays a PD game against one other agent with an identical tag. If an agent has a unique tag in the population, it plays against a randomly selected opponent. The next generation is formed via fitness proportionate reproduction where the fitness of an agent is the payoff received in this round of play. Mutation is then applied to each bit. This process is described in Algorithm 1. Hales was able to develop sustained cooperation but with only large tag sizes. We recount Hales’ [10] explanation of how tags help promote cooperation in the PD game. A homogeneous group of cooperators will prosper and grow. When such a group is invaded, via mutation of the strategy or tag, by a defector, the defector will prosper, resulting in imitators in the next generations. Over time, the group will fill with defectors, resulting in declining performance and eventual extinction. Thus defectors, even if formed by chance, will not live long, and hence a majority of the population will be cooperators. There was no clear explanation of why large tags were required to sustain cooperation.

Algorithm 1 Hales and Edmond’s model of population evolution with tags.

```

for some number of generations do
  for each agent  $a$  in the population do
    Select a game partner agent  $b$  with the same tag (if possible)
    Agent  $a$  and  $b$  invoke their strategies and  $a$  gets corresponding payoff.
  end for
  Reproduce agents in proportion to their average payoff (with some low level of mutation)
end for

```

McDonald and Sen [17] observe that an individual with a unique tag, i.e., a singleton, can prosper if the randomly chosen individual it interacts with is cooperative. If the singleton is cooperative, it will perform well, probably leading to more agents copying its tag bits and strategy, causing the group to expand as a group of cooperators. If the singleton is a defector, it will obtain a high reward, but when others copy its strategy and tag, all agents in the resulting group will be defectors. They will then perform poorly in the following generations, and the group will die out. They combined this observation with that of Hales’ explanation of population evolution presented above to provide a more complete and detailed picture of the evolution of cooperation in PD with tags: At any point in time there are many groups in the population. Homogeneous groups of cooperators expand through preferential selection. Ultimately such an expanding group is invaded, mostly through strategy mutation, by a defector. As the defector prospers, more

copies of it are made, and the corresponding deteriorating performance leads to group extinction. To sustain cooperation, therefore, at any point in time, there has to be a sufficient number of homogeneous cooperative groups in the population such that mutation cannot simultaneously infect a majority of them with defectors. Mutation is the primary mechanism to create this necessary diversity of tag groups². If tag mutation rate is low, large tags are needed to guarantee sufficient diversity of groups. To support this analysis McDonald and Sen were able to sustain cooperation with much smaller tag lengths when they increased the tag mutation rate.

Another, perhaps more important, but controversial claim by McDonald and Sen involved the limitation of self-matching tags. They argued that the basic premise of tags was that agents with similar tags had similar strategies and hence when agents are playing other agents with the same tag they are playing the same strategies. In games like PD, such behavior imitation leads to preferable Pareto-optimal outcomes. If identical strategies do not produce good outcomes, e.g., in the “Anti-Coordination” game then tag based evolution will fail to produce cooperation. Their experimental results with the AC game confirmed this conjecture.

We, however, believe that the limitation identified by McDonald and Sen holds only for self-matching tags. More general tag frameworks, such as the ones that incorporate the following proposals, transcend these limitations.

We now propose three new tag mechanisms with corresponding motivations:

Tag matching patterns (one and two-sided): To allow an agent to interact with members belonging to other tag groups, we include a tag-matching string, \mathcal{M} defined over the alphabet $\{0, 1, *\}$ where we assume that the tags are binary strings of length L , and the $*$ is a don’t care symbol. In *one-sided matching*, an agent i can interact with an agent j if \mathcal{M}_i matches T_j , the tag of agent j . \mathcal{M}_i matches T_j iff $\forall k = 1, \dots, L, \mathcal{M}_i(k) = \times \vee \mathcal{M}_i(k) = T_j(k)$. If i matches j , only i , and not j receives the payoff from i ’s interaction with j . In *two-sided matching*, an agent i can interact with an agent j if \mathcal{M}_i matches T_j and \mathcal{M}_j matches T_i . While one-sided matching allows an agent to select partners, two-sided matching allows more stringent pairings that require mutual consent.

Payoff sharing: In some real-life situations side-payments are allowed, e.g., payoff distributions in coalitions. Payoff-sharing is a similar mechanism by which an agent shares a fraction, α of its payoff with its opponent. Such sharing can encourage self-interested agents to converge to social-welfare maximizing outcomes rather than outcomes obtained with safe or greedy strategies.

Paired reproduction: This is a more subtle operator motivated by McDonald and Sen’s analysis of the operation of self-matching tags in PD games. Sustaining

²Mutation will often form singletons in the population. Note that the current framework allows singletons to play with a randomly selected member from the population. This is a rather key “boundary condition”, if singleton agents were not allowed to play, i.e., their fitness was to remain at zero, they would not be reproduced and cooperation could not be sustained in the PD.

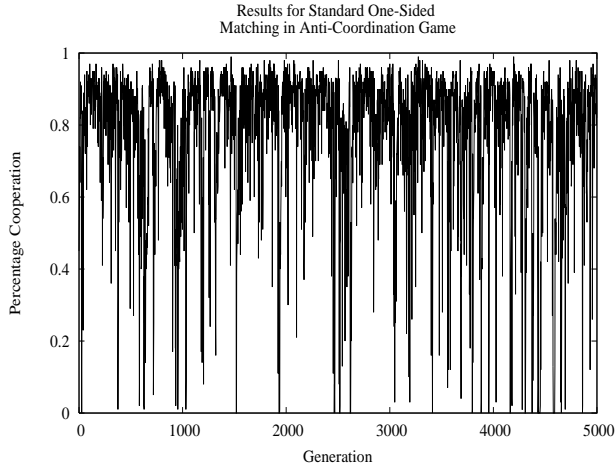


Figure 3: One-sided matching in the AC game.

coordination requires multiple “groups of cooperators” in the population so that if a few of these are disrupted other cooperators can fill the niche. While tag mutation is sufficient to create such diversity when self-matching tags are used, special reproduction operators which make copies of matching pairs of individuals with mutations at corresponding places on the tag of one and the matching string of the other, to preserve the match after mutation, is necessary. Paired reproduction, then is an infrequent special reproduction operator used when tag matching patterns are used, and is applied with probability p_{pr} . It copies over a randomly selected matching pair of individuals and alters the tag and matching tag bits in the pair with the probability of mutation such that they continue to match thereafter.

3. EXPERIMENTAL RESULTS

Now we present experimental results to evaluate our new tag mechanisms for promoting coordination in the AC and PD games. Unless otherwise noted, the experimental results were obtained using a population of 100 agents with tag lengths of 8 bits. The AC and PD games represent situations where homogeneous and heterogeneous optimal strategies are required for effective coordination. Imitation is effective in PD but detrimental in AC.

McDonald and Sen [17] noted that self-matching tags could only generate 50% coordination rate for the AC game which is similar to random action selection by the agents. Significantly higher coordination rates were obtained on this problem when one-sided matching was used (see Figure 3). Though the mean coordination rate improved significantly, the standard deviation is significant which indicates that it is difficult for the system to maintain a stable equilibrium. This suggests that the one-sided matching method needs further adjustment.

The problem with a simple one-sided matching scheme is that it does not necessarily support the creation of pairwise optimal interactions. For example, in this scheme, if agent A 's matching string \mathcal{M}_A matches agent B 's tag T_B , A can play with B without \mathcal{M}_B matching T_A . Therefore, even if A receives a high payoff playing with B , B may receive a

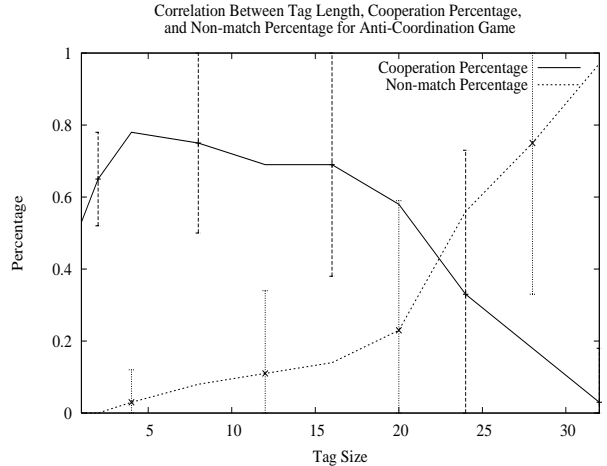


Figure 4: Effect of tag length in the AC game.

poor payoff playing someone else. Thus B may be replaced and this in turn may jeopardize A 's survival. This instability can be demonstrated by varying the size of the tag space as shown in Figure 4. As tag length increases, coordinated actions decrease as it becomes more difficult to find matches in the tag space, and even more so to find optimal matches. Simply choosing a smaller tag space has its own pitfalls though. Tag groups must remain homogeneous, otherwise when an agent A matches a group of agents G in the population, and one or more of G 's member's strategies mutate, it will cause A to have a non-zero probability of non-optimal interaction. This defector in group G can ultimately lead to the death of the group that A belongs to, which may lead to instability in group G and its ultimate demise. Thus, if too few groups exist due to a small tag space, the system will be highly unstable due to a lack of homogeneous groups. The choice of an appropriate tag space is clearly important. However, simply optimizing the size of tag spaces will not produce a solution to all the problems posed by tag matching schemes.

We next tried the one-sided matching scheme on the PD game (see Figure 5). Much to our chagrin, there is an extreme lack of cooperation in the population in this case. In fact, over many runs of the one-sided matching system, convergence to the Nash Equilibrium of the PD was the consistent result. To have cooperation in PD we must have two cooperative agents with matching strings that match each other's tags. However, in random initializations, all of these criteria being satisfied simultaneously is both rare and unlikely to last due to mutations which cause heterogeneity to be lost in any cooperating group. Additionally, as observed earlier, many such cooperating groups are needed in order to ensure the survival of cooperating agents. It is much more likely in a one-sided matching system that a cooperator A matches a cooperator B who matches a defector C , leading to the ultimate death of both cooperators. Thus, the system converges to a non-cooperative equilibrium.

One solution to this problem lies in ensuring that both interacting pairs remain alive in the population. One way of doing this is by sharing a percentage of the payoff between the two agents playing. With payoff sharing, as in the previous example, an agent A can now share some of

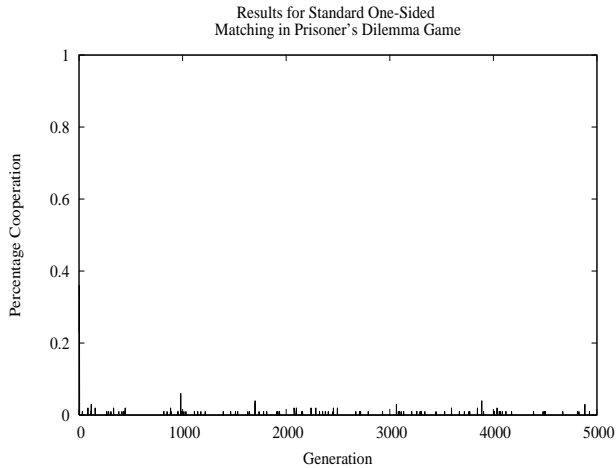


Figure 5: One-sided matching in the PD game.

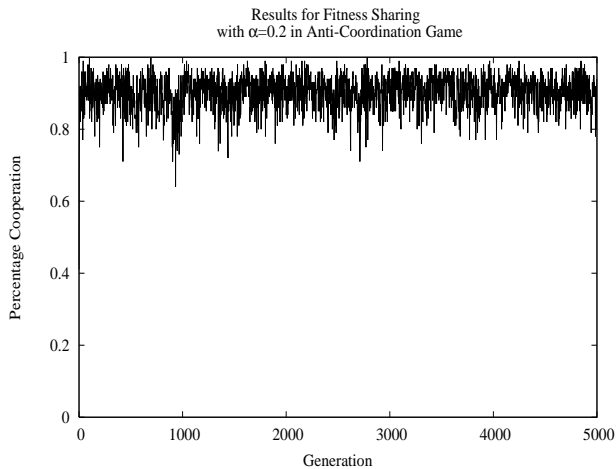


Figure 6: Payoff sharing in the AC game.

its high payoff with B , thus keeping it alive even though it is exploited by C . Similarly, a defector will share some of its high payoff with an interacting cooperator to keep it alive. In Figure 6 we observe that the payoff-sharing scheme effectively promotes coordination in populations of agents playing the AC game for $\alpha = 0.2$.

For runs of the PD game, payoff sharing produced interesting results in creating populations of cooperators. Figures 7 (a) and (b) show runs of the PD game with varying amounts of payoff sharing. As the sharing percentage α increases from 0.2 to 0.4, the population increases in cooperative interactions and stability for the same reasons that the Anti-Coordination game gains stability. That is, in the case that agent A , a cooperator, matches agent B , also a cooperator, but B does not match A , B will get a payoff from its interactions with A , regardless of whether or not it's other interactions are optimal. The higher fitness values here provide B a higher chance for survival during the evolutionary selection process, and thus makes the system more stable. The only drawback is that, though payoff sharing works for both AC and PD games, it may be used only when side payments are allowed.

When side payments are not allowed, we need other mechanisms. Next we evaluate two-sided matchings on the AC game (see Figure 8). This system produces a high number of coordinated interactions but cannot match the stability provided by payoff sharing. Groups that undergo mutation and the subsequent loss of homogeneity have no fall-back payoff to help them survive until homogeneity is restored. Thus it is more difficult for the system to return to equilibrium and there is a short, but noticeable down period for coordinated interactions before the system can grow new groups to replace the lost groups.

Double matching completely fails to solve the fundamental problems in the PD game (see Figure 9). Cooperating groups are formed in this system, seen as spikes in the figure. However, once homogeneity is lost in any of these groups, they die out almost immediately.

The payoff sharing and matching approaches form parts of a solution to the problem presented in forming large numbers of interacting groups which can survive indefinitely. Using the two methods simultaneously leads to a system which satisfies the properties necessary for the survival of cooperating groups. Payoff sharing provides a suboptimal group with the chance of survival until it can restore its own equilibrium, while two-sided matching ensures that interactions will occur in a pairwise format. Figures 10 (a) and (b) demonstrate the effectiveness of this system for both Anti-Coordination and Prisoner's Dilemma. The system is not quite as stable as simple payoff sharing for Anti-Coordination, since the double matching makes it easier for random mutations to leave an agent with no group to interact with, either due to a matching string mutation or tag mutation which in turn causes double matching to be unable to locate a suitable partner. However, the system does achieve some improvements in the PD game. Though the difference is marginal between Figure 7 (b) and Figure 10 (b), the combined system performs consistently better on average than payoff sharing alone.

Finally, we introduce the paired reproduction operator that is specifically tailored to the needs of tag based systems with matching strings. From our experiments, we have seen that the necessary criteria of successful interaction is that tag groups form pairwise interactions with each other, that these interactions are optimal, and that there are enough groups in the tag space to tolerate the introduction of defectors and the subsequent loss of groups. If we add paired reproduction to the standard set of genetic operators in tag systems, we obtain an elegant solution to the problem. The performance of this mechanism is just in the preliminary stages of testing, however the results for both Prisoner's Dilemma and Anti-Coordination games are encouraging (see Figures 11 (a) and (b)). Though the corresponding coordination levels are lower than that with payoff sharing, this scheme can be used when side payments, and hence payoff sharing, are disallowed. It should be noted that in our examples we give this new operator a relatively low frequency. This is due to the problems of matching discussed earlier in the paper. Since the operator is designed to create new interacting groups in the tag space so as to address the loss of groups, it adds many more groups to the space than would exist under the tag systems we have demonstrated. Thus, if we use this operator too often we will end up with many small groups which have a higher likelihood of being replaced. This would lead to their partner group's death and

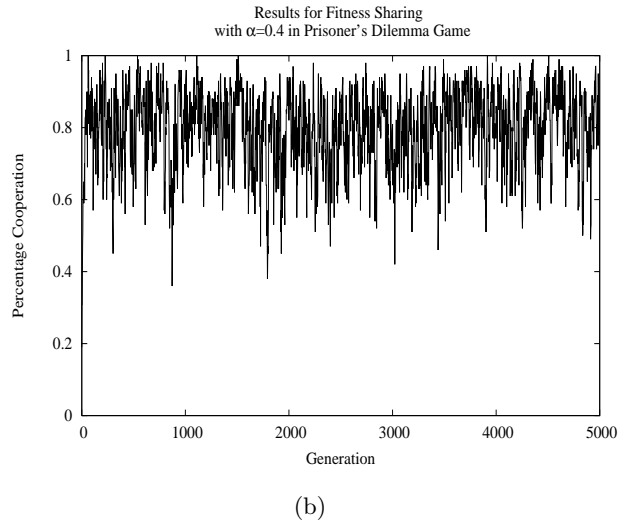



Figure 7: Payoff sharing in the PD game: (a) $\alpha = 0.2$, (b) $\alpha = 0.4$.

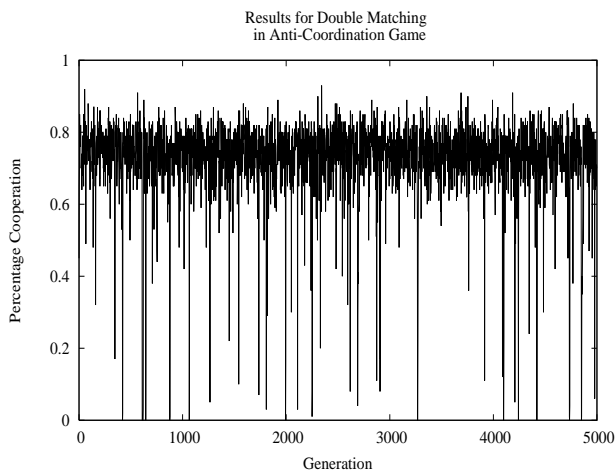


Figure 8: Two-sided matching in the AC game.

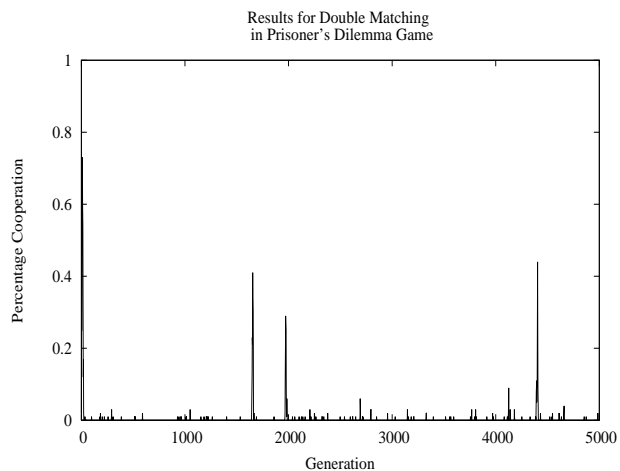


Figure 9: Two-sided matching in the PD game.

a less stable equilibrium. Thus, we chose a low frequency of reproduction to counter-balance this effect.

4. CONCLUSIONS

We argue for the development of mechanisms for reusing effective policies for interacting with one agent against a similar agent. Identifying “similar” agents can be difficult in general. When agents copy both internal strategies and external features of other, successful, agents in the population, we have an evolutionary process that groups agents with similar genetic makeup and external non-coding features. Such external features or tags can help identify individual agent types and allow us to reuse learned coordination policies with one agent with other agents of the same type.

Previous research on using self-matching tags to evolve cooperation had succeeded only on problems where imitation (identical moves) is necessary for cooperation. We introduce

and validate more general tag mechanisms that allow intra-group matchings and is found effective on both problems that require identical and problems that require distinct actions to produce coordinated behavior. In particular, fitness sharing seems to be particularly promising when side payments are allowed, and two-sided matching with paired reproduction works reasonably well in other cases. We plan to run further experiments with combinations of these mechanisms. Additional experimentation with larger range of parameter values and combination of parameter ranges (tag length, payoff sharing fraction, etc.) will also be performed.

While the new tag mechanisms are reasonably effective, there is scope for improvement in coordination performance. We plan to study the inefficiencies of the current method to identify improvements. We plan to test these mechanisms on a larger class of coordination problems including the Chicken Game, the El Farol bar problem, etc. We also intend to evaluate tag-based coordination in domains with delayed,




Figure 10: Payoff sharing combined with two-sided matching: (a) AC game, $\alpha = 0.2$, (b) PD game, $\alpha = 0.4$.

rather than immediate, feedback.

We believe that we have merely scratched the surface of the possibilities for using tags for evolving coordinated actions between agent clusters. Tags can be generalized to arbitrary patterns and richer representations to support significantly more complex interactions between agents. We are far from representing the complex and nuanced visual, verbal, and other physical cues that serve as “tags” for communication between human and animal groups, but careful study, abstraction, modeling, and implementation of such external feature-based communication and their relation to successful coordination, signaling, and negotiation can provide important insights into developing effective coordination strategies for groups of artificial agents.

Acknowledgment: US National Science Foundation award IIS- 0209208 and Tulsa Undergraduate Research Challenge (TURC) program partially supported this work.

5. REFERENCES

- [1] P. D. Allison. The cultural evolution of beneficent norms. *Social Forces*, 71:279–301, 1992.
- [2] R. Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.
- [3] R. Axelrod. *The complexity of cooperation: Agent-based models of conflict and cooperation*. Princeton University Press, Princeton, NJ, 1997.
- [4] A. Ballim and Y. Wilks. Beliefs, stereotypes and dynamic agent modeling. *User Modeling and User-Adapted Interaction*, 1(1):33–65, 1991.
- [5] J. Bendor and P. Swistak. The evolutionary stability of cooperation. *American Political Science Review*, 91(2):290–307, June 1997.
- [6] D. Carmel and S. Markovitch. Learning models of intelligent agents. In *Thirteenth National Conference on Artificial Intelligence*, pages 62–67, Menlo Park, CA, 1996. AAAI Press/MIT Press.
- [7] A. J. Courtney. Control-display stereotypes for multicultural user systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(4):681–687, July-Aug. 1992.
- [8] M. Genesereth, M. Ginsberg, and J. Rosenschein. Cooperation without communications. In *Proceedings of the National Conference on Artificial Intelligence*, pages 51–57, Philadelphia, Pennsylvania, 1986.
- [9] S. H. Ha. Digital content recommender on the internet. *IEEE Intelligent Systems*, 21(2):70–77, 2006.
- [10] D. Hales and B. Edmonds. Evolving social rationality for mas using “tags”. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 497–503, Melbourne, Australia, July 2003. ACM Press.
- [11] D. Hales and B. Edmonds. Applying a socially-inspired technique (tags) to improve cooperation in p2p networks. *IEEE Transactions in Systems, Man and Cybernetics - Part A: Systems and Humans*, 35(3):385–395, 2005.
- [12] J. H. Holland, K. Holyoak, R. Nisbett, and P. Thagard. *Induction: Processes of Inferences, Learning, and Discovery*. MIT Press, Cambridge, MA, 1986.
- [13] C. M. Judd and B. Park. Definition and assessment of accuracy in social stereotypes. *Psychological Review*, 100(1):109–128, Jan. 1993.
- [14] B. Julstrom. The development of cooperation in genetic algorithms for the iterated prisoner’s dilemma. In *Proceedings of the 33rd Annual Midwest Instruction and Computing Symposium (MICS 2000)*, April 2000.
- [15] P. t. H. K. Tuyls and B. Vanschoenwinkel. An evolutionary dynamical analysis of multi-agent learning in iterated games. *Autonomous Agents and Multi-Agent Systems*, 12(1):115–153, 2006.
- [16] J. A. Konstan. Introduction to recommender systems: Algorithms and evaluation. *ACM Transactions on Information Systems*, 22(1):1–4, Jan. 2004.
- [17] A. McDonald and S. Sen. The success and failure of




Figure 11: Pair reproduction: (a) AC game, (b) PD game.

- tag-mediated evolution of cooperation. In K. Tuyls, P. J. Hoen, K. Verbeeck, and S. Sen, editors, *LAMAS*, volume 3898 of *Lecture Notes in Computer Science*, pages 155–164. Springer, 2005.
- [18] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.
- [19] E. Nudelman, J. Wortman, K. Leyton-Brown, and Y. Shoham. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *Proceedings of the Third International Joint Conference on Autonomous Agents and MultiAgent Systems*, pages 880–887, 2004.
- [20] G. Paliouras, V. Karkaletsis, C. Papatheodorou, and C. Spyropoulos. Exploiting learning techniques for the acquisition of user stereotypes and communities. In *Proceedings of the International Conference on User Modeling*, pages 169–178. Springer-Verlag, 1999.
- [21] L. Panait and S. Luke. Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems*, 11(3):387–434, 2005.
- [22] B. Price and C. Boutilier. Accelerating reinforcement learning through implicit imitation. *Applied Artificial Intelligence Journal*, 19:569–629, 2003.
- [23] P. Resnick and H. R. Varian. Recommender systems. *Communications of the ACM*, 40(3):56–58, March 1997.
- [24] E. Rich. Stereotypes and user modeling. In A. Kobsa and W. Wahlster, editors, *User Models in Dialog Systems*, pages 35–51. Springer-Verlag, Berlin, 1989.
- [25] E. A. Rich. User modeling via stereotypes. *Cognitive Science*, 3:329–354, 1979.
- [26] R. Riolo. The effects and evolution of tag-mediated selection of partners in populations playing the Iterated Prisoner’s Dilemma. In *Proceedings of the Seventh International Conference on Genetic Algorithms*, pages 378–385. Morgan Kaufmann Publishers, Inc., 1997.
- [27] R. Riolo, M. D. Cohen, and R. Axelrod. Cooperation without reciprocity. *Nature*, 414:441–443, 2001.
- [28] M. Root. How we divide the world. In D. A. Howard, editor, *PSA’1998: Proceedings of the 1998 Biennial Meetings of the Philosophy of Science Association, Part II: Symposium Papers*, pages S628–S639. Philosophy of Science Association, Newark, Delaware, 1998.
- [29] J. B. Schafer, J. Konstan, and J. Riedl. Recommender systems in e-commerce. In *Proceedings of the First ACM Conference on Electronic Commerce*, pages 158–166, New York: NY, 1999. ACM Press.
- [30] B. Shapira, P. Shoval, and U. Hanani. Stereotypes in information filtering systems. *Inf. Process. Manage*, 33(3):273–287, 1997.
- [31] K. Sigmund and M. A. Nowak. The alternating prisoner’s dilemma. *Journal of Theoretical Biology*, 38:262–275, 1994.
- [32] J. L. Stimpson, M. A. Goodrich, and L. C. Walters. Satisficing and learning cooperation in the prisoner’s dilemma. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, pages 535–540, 2001.
- [33] P. Stone and M. Veloso. Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, 8(3), July 2000.
- [34] D. Zeng and K. Sycara. Bayesian learning in negotiation. *International Journal of Human Computer Studies*, 48, 1998.
- [35] X. Zhang and H. Han. An empirical testing of user stereotypes of information retrieval systems. *Inf. Process. Manage*, 41(3):651–664, 2005.