

Preprint copy of: Spector, L., and J. Klein. 2006. Multidimensional Tags, Cooperative Populations, and Genetic Programming. In *Genetic Programming Theory and Practice IV*, edited by R.L. Riolo, T. Soule, and B. Worzel, pp. 97-112. New York: Springer-Verlag.

Multidimensional tags, cooperative populations, and genetic programming

Lee Spector¹ and Jon Klein²

¹ Cognitive Science, Hampshire College, Amherst, MA, 01002-3359 USA.
lspector@hampshire.edu

² Physical Resource Theory, Chalmers University of Technology & Göteborg University, Göteborg, Sweden. jk@artificial.com

Summary. We present new results on the evolution of tag-mediated cooperation, demonstrating that the use of multidimensional tags can enhance the emergence of high levels of cooperation. We discuss these results in the context of prior cases in which work on the evolution of cooperation has led to practical techniques for improving the problem-solving performance of genetic programming systems.

1 Cooperation and adaptive complexity

In this chapter we explore the conditions under which altruistic cooperation is produced by natural selection, and we do so as part of an effort to improve the problem-solving performance of genetic programming systems. But why would anyone think that the evolution of cooperation would provide clues for the improvement of problem solving systems? What does cooperation have to do with the kinds of adaptive complexity that we seek from our genetic programming systems?

One answer is that this connection—between the study of the evolution of cooperation and the improvement of problem-solving performance—has been fruitful in the past, so it might be worth exploring further (see Section 6 below). But a deeper answer might also draw on the observation that many biological systems appear to make use of cooperative interactions at several levels of organization, and that these interactions may be important for the evolution of adaptive complexity.

Adaptive complexity in nature is usually produced not by isolated individuals but rather by ecosystems that are structured by both genetic and economic relationships. Typically many of these relationships involve cooperation or other forms of mutualism. Indeed, some theorists have credited cooperation with an essential role in the evolutionary transitions that produced successive levels of adaptive complexity in the history of life. For example, in discussing the origin of multicellularity Michod writes:

We see the formation of cooperative interactions among lower-level individuals as a necessary step in evolutionary transitions; only cooperation transfers fitness from lower levels (costs to group members) to higher levels (benefits to the group). —[Michod, 2003, p. 292]

So there are reasons to think that cooperation can be important in adaptive evolutionary systems. It is therefore reasonable to expect useful insights, even for practical applications, to emerge from the study of the evolution of cooperation.

In the remainder of this chapter we discuss these issues in more concrete terms, focusing on a particular model of cooperation and on ways in which a new development within this model might be applied to genetic programming practice. In the next two sections we describe the model (tag-mediated cooperation) and then the new development (the use of multidimensional tags). We then provide quantitative results, a discussion of the meaning of those results, and several suggestions for incorporation of related mechanisms into genetic programming systems. These suggestions are speculative but we argue that they merit further exploration.

2 Tag-mediated cooperation

Altruism and cooperation are behavioral traits that have long drawn the attention of evolutionary theorists. From a naive reading of Darwinian theory one might expect that natural selection would produce only selfish agents, and that cooperation would be maladaptive and therefore rare. Early theorists recognized, however, that while natural selection may favor selfish *genes* it does not necessarily favor selfish *agents*. Beginning with Hamilton in the 1960s biologists have built and tested quantitative theories of the evolution of cooperative behavior among kin, grounded in the understanding that degree of kinship degree of genetic similarity are closely related [Hamilton, 1963, Hamilton, 1964]. More recent work has expanded the class of conditions under which cooperation can be expected to evolve; such conditions now include the presence of reciprocating partners [Axelrod and Hamilton, 1981, Trivers, 1972] and the presence of partners with known good reputations [Nowak and Sigmund, 1998].

Models of “tag-mediated” cooperation have the potential to explain the evolution of cooperation in an even wider range of cases. A tag is a simple marker, represented in most of the prior models as a floating point number, that is attached to each agent and is visible to other agents [Holland, 1995]. Tags can be used to model a variety of identity-based interactions in biological systems, some of which have physical implementations that are far simpler than those that underly judgments of kinship, reciprocation, or reputation.

In the models of tag-mediated cooperation first presented by Riolo et al. [Riolo et al., 2001] each agent has both a floating point tag and a floating

point “tolerance,” and one agent will make an altruistic donation to another if the tags of the two agents differ by no more than the donor’s tolerance. The simulation proceeds through rounds of donation attempts followed by reproductive tournaments, the winners of which become, after possible mutation, the agents for the next generation of donations and reproductive tournaments.

Roberts and Sherratt raised concerns about the original model and noted that cooperation failed to emerge if tolerances were allowed to drop below zero, thereby permitting truly selfish agents [Roberts and Sherratt, 2002]. While Roberts and Sherratt were correct in the context of the particular values that they chose for system parameters, subsequent work has shown that tag-mediated cooperation does indeed arise robustly under a variety of equally reasonable parameter settings [Riolo et al., 2002, Axelrod et al., 2004, Spector and Klein, 2006]. Some of our previous work, in particular, showed that cooperation readily emerges, even when tolerances are allowed to drop below zero, if mutation rates are low and/or if agent interactions are geographically limited [Spector and Klein, 2006]. These results inspired a practical technique for improving the performance of genetic programming systems, which we discuss briefly in Section 6. The new work presented in Sections 3, 4 and 5 extends our previous results and may have additional implications for genetic programming practice (also discussed in Section 6).

3 Multidimensional tags

Our observations of the previous models and their dynamics revealed that a group of cooperators can be destabilized in a number of ways. Downward tolerance drift is dangerous for a group of cooperators because a mutation resulting in low or even negative tolerance produces a selfish agent situated to exploit its generous neighbors with similar or identical tags. Upward tolerance drift is also dangerous because it may lead to overly generous agents that begin to donate to agents outside of the group. Finally, even agents with well-balanced tolerance values may be invaded by selfish agents that are able to mimic their tags. Because agents are not able to explicitly manipulate their own tags, this can occur only as a result of random mutation.

This last possibility is the focus of our current work. Invasions by selfish mimics can occur only when the invaders successfully “guess,” by means of mutation, the tags of nearby cooperators. One can think of each cooperator’s tag as a password and a successful invader as a hacker who has discovered that password, although the only hacker’s only means for discovery is random mutation. On the other hand, the hacker need not get the password exactly right; it will succeed if it falls within the cooperator’s tolerance of the cooperator’s tag.

This analogy leads naturally to a hypothesis that cooperation may be more easily maintained if tag values are made more difficult to mimic. To improve the security of a password-secured system we can increase the length

of the password; by analogy we can increase the “security” of tag-mediated cooperation by adding additional dimensions to the tags.

In this chapter we explore how the introduction of multidimensional real-valued tags can lead to improved cooperation in evolving populations. Earlier work on this theme has been conducted by Hales and Edmonds, who have experimented with multidimensional tags in the form of bit strings [Hales, D., 2005, Edmonds, 2006]. Some of these models allow only exact tag matching, while others use the Hamming distance between bit strings as a measure of tag difference. By contrast, we have retained the real-valued tag and tolerance scheme from the models of Riolo et al. [Riolo et al., 2001] but we have extended it to use tags consisting of sequences of real values.³

We begin with a model similar to those used in prior work deriving from Riolo et al. [Riolo et al., 2001] in which a population of 100 agents is run through a series of generations (30,000 in previous work; 400,000 here). At each generation, agents begin with a score of 0 and are given the opportunity to make donations to $P = 3$ randomly chosen agents. An agent makes a donation to a recipient if and only if the difference between their tag values is less than or equal to the potential donor’s tolerance value. If a donation event occurs then a cost C is deducted from the donor’s score (which is allowed to go negative) and a benefit B is added to the recipient’s score. In the original model, the values $C = 0.1$ and $B = 1.0$ were used; in this chapter and in our previous work we investigate a variety of cost to benefit ratios with $C = \{0.1, 0.5, 1.0\}$ and $B = 1.0$.

After all donation interactions have been completed, agents are selected for reproduction based on their scores. Each agent at position N is paired randomly with another agent; the agent with a higher score (or the agent at N in the event of a tie) is selected for reproduction at position N . The child inherits its tag and tolerance from the parent, but each is mutated with probability $m = \{0.001, 0.01, 0.1\}$. When a tag is mutated it is replaced by a new value, sampled uniformly from $[0, 1]$. Tolerance values are mutated by summation with random Gaussian noise with standard deviation 0.01. As in the work of Roberts and Sherratt [Roberts and Sherratt, 2002], the tolerance value is clamped to a lower bound of -10^{-6} .

As in our previous work on this model we introduce a simple form of spatial structure, or geography, that changes the way that agents interact. We model the population as a one dimensional ring in which agents consider only others in a local radius R when selecting recipients for donations or competitors for reproduction. We examine a variety of values for R from 1 up to 50. Note that the special case of $R = 50$ is equivalent to the non-spatial populations used in previous models.

³ Edmonds refers to systems with “multi-dimensional binary and continuous vectors” but we are not aware of published details [Edmonds, 2006].

In our multidimensional tag model the single floating point tag value is replaced by a vector of floating point values. The definition of the distance between two tags, t_1 and t_2 , is generalized for n dimensions as:

$$\sqrt{\sum_{i=1}^n (t_1[i] - t_2[i])^2} \quad (1)$$

As in previous models, tag mutation produces completely new tag values; each value in the tag vector is replaced by a new random value, uniformly sampled from $[0, 1]$.

Tolerance values remain as single floating point values and are interpreted as they were in previous models: an agent shares with another agent if and only if the difference in their tags is less than or equal to the donor’s tolerance. We note that while the maximum multidimensional tag distance, given by \sqrt{n} , exceeds the 1.0 limit found in the previous tag models, there is no upper bound on the tolerance value so that the model allows for cooperation between any pair of agents, even those with large tag differences. In practice, tolerance values tend to be far smaller than the maximum distance value.

4 Results

We performed 92 runs for each condition and collected average donation rates and tolerance values. The results are presented here as averages over all runs in each condition.

We first discuss the results in the conditions in which the cost (C) charged to a cooperating agent for a successful donation was 0.1. This is the cost structure that has received the greatest attention in the literature. As in the prior work we characterize the amount of cooperation in a condition by reporting the percentage of donation attempts that are successful across the entire population and across the entire simulation; we call this measure the “donation rate.”

Figure 1 shows the average donation rates under all $C = 0.1$ conditions. The vertical bars show the average donation rates for standard, one-dimensional tags for each combination of mutation rate and interaction radius. This data replicates that from our previous studies, although the data plotted here is from new, independent runs. In line with and generally above each vertical bar in Figure 1 are symbols plotting the average donation rates for runs with higher-dimensional tags.

The one-dimensional data demonstrates that cooperation readily emerges except in a few conditions with high mutation rates and large interaction radii (the very conditions that formed basis of the critique by Roberts and Sherratt [Roberts and Sherratt, 2002]). In all but the most recalcitrant of these conditions the levels of cooperation are further boosted by increasing the dimensions of the tags. Substantial improvements result from the first augmentation to




Fig. 1. Average donation rates as a function of mutation rate, interaction radius, and number of tag dimensions. For the data in this graph the cost (C) charged to a cooperating agent for each donation was 0.1, as in most of the prior research. Each plotted point represents the average of 92 independent runs.

two-dimensional tags, and significant improvements result from the next several augmentations as well. In most cases the payoff for additional dimensions eventually tapers off, presumably because the advantages that can be gained by making one’s tag hard to mimic have a natural limit; as the probability of randomly generating a particular tag approaches zero, little is to be gained by decreasing the probability further.

Figure 2 shows the average donation rates under all $C = 0.5$ conditions. In this higher cost regime cooperation is slightly harder to achieve, but significant cooperation nonetheless results from most of the parameter sets that we tested. Note that the levels of cooperation for higher numbers of dimensions are more spread out, and that in many cases a significant improvement can be achieved when increasing the number of dimensions from 5 to 6. This may reflect the fact that invasions by selfish agents, while rare in a simulation with five-dimensional tags, are highly disruptive when they do occur.

Figure 3 shows the average donation rates under all of the $C = 1.0$ conditions. In these “zero sum” conditions the cost to a donor is equivalent to the benefit gained by a recipient, and cooperation is significantly harder




Fig. 2. Average donation rates as a function of mutation rate, interaction radius, and number of tag dimensions. For the data in this graph the cost (C) charged to a cooperating agent for each donation was 0.5, which is five times the cost used in most of the prior research. Each plotted point represents the average of 92 independent runs.

to achieve. As was reported previously, significant levels of cooperation can nonetheless be achieved even with one-dimensional tags, with donation rates exceeding 12% in one of our tested configurations. As can be seen in Figure 3, however, one can achieve significantly higher levels of cooperation with higher-dimensional tags. Indeed, with a mutation rate of $m = 0.01$, an interaction radius of $R = 1$, and six-dimensional tags we observe a donation rate of 75.7%.

Figures 4, 5, and 6 show the average observed tolerances in the conditions with donation cost $C = 0.1$, $C = 0.5$, and $C = 1.0$ respectively. Note that cooperation is possible even with a tolerance of zero, although a zero-tolerance agent will cooperate only with others that have identical tags. The presence of larger tolerances indicates agents that will donate to a wider range of recipients. Note also that while negative tolerances are possible, all of the average tolerance values that we observed were non-negative; this is why no negative tolerances appear in our graphs.




Fig. 3. Average donation rates as a function of mutation rate, interaction radius, and number of tag dimensions. For the data in this graph the cost (C) charged to a cooperating agent for each donation was 1.0, the same as the benefit to the recipient. Each plotted point represents the average of 92 independent runs.

In nearly all cases the observed tolerance is higher with higher-dimensional tags. This can be explained as resulting from the fact that multidimensional tags increase the size of the “tag space” and thereby reduce the risk of outside invasion. By doing this they allow groups of cooperating agents to safely raise their tolerance values and thus protect themselves from the secondary threat of self-destruction due to a crash in tolerance values.

5 Discussion

In our previous analysis of the one-dimensional case we suggested that the mechanism behind tag-mediated cooperation can be thought of as a kind of “probabalistic kin selection” [Spector and Klein, 2006]. While the agents in the model have no access to explicit kinship information, and therefore cannot be certain that they are donating to kin, successful tag-mediated cooperation can nonetheless arise because kin are more likely to have similar tags.




Fig. 4. Average tolerances as a function of mutation rate, interaction radius, and number of tag dimensions. For the data in this graph the cost (C) charged to a cooperating agent for each donation was 0.1, as in most of the prior research. Each plotted point represents the average of 92 independent runs.

Our findings here bolster this notion by clarifying the evolutionary dynamic behind the perpetuation of cooperating subpopulations. The tag mutation scheme in our model, which was taken from the previous work in the literature [Riolo et al., 2001], does not generally produce new tag values that are close to those of their ancestors. Instead, tags mutate spontaneously to entirely new values, meaning that agents with a common ancestor have either identical tag values or, following a mutation event, tag values that are no more similar than those of completely unrelated agents. Tolerance values, on the other hand, are subject to incremental drift. Were it not for incremental tolerance drift a population of cooperating agents would probably settle on an infinitesimally small “optimal” value which would allow for cooperation with identically tagged agents while avoiding invasion from defectors. The incremental drift of tolerance values, however, creates a more complex evolutionary dynamic in which, on the group level, a subpopulation of successful cooperating agents must strive for *higher* tolerance values in order to avoid an accidental self-destructive tolerance drop below zero. At the same time, on




Fig. 5. Average tolerances as a function of mutation rate, interaction radius, and number of tag dimensions. For the data in this graph the cost (C) charged to a cooperating agent for each donation was 0.5, which is five times the cost used in most of the prior research. Each plotted point represents the average of 92 independent runs.

the individual level, tolerance values must be kept low to avoid exploitation by unrelated agents.

Invasion of a population of cooperators requires that a would-be freeloader's tag value mutates to be close to that of the cooperators, while maintaining a tolerance value close to zero. As discussed above, tag values can be thought of as passwords shared among kin, albeit passwords with “fuzzy” interpretations due to the effects of tolerance values. With multidimensional tags, the passwords become exponentially more difficult to guess. Note that the exponential nature of the “password guessing problem” means that multidimensional tags cannot be exploited simply by increasing the mutation rate: the complexity of the passwords grows too quickly for mutation to keep up. Note also that there is a tradeoff between the advantages of multidimensional tags and the need for more sophisticated mechanisms (cognitive, chemical or otherwise) to maintain and recognize them.

Although we feel that our findings support the notion of “probabilistic kin selection” in this model and similar models, our findings do not preclude the




Fig. 6. Average tolerances as a function of mutation rate, interaction radius, and number of tag dimensions. For the data in this graph the cost (C) charged to a cooperating agent for each donation was 1.0, the same as the benefit to the recipient. Each plotted point represents the average of 92 independent runs.

possibility that other mechanisms can also support tag mediated cooperation. For example, Hales has shown that tag-based models can lead to cooperation among groups of unrelated agents with diverse skills [Hales, 2002]. Our probabilistic kin selection interpretation is also compatible with other recent analyses of tag mediated cooperation (e.g. [Jansen and van Baalen, 2006]).

6 Cooperation and genetic programming

Researchers have previously drawn several connections between work on the evolution of cooperation and work on evolutionary computation. One example comes from our own prior work on tag-mediated cooperation, in which we highlighted the ways in which a particularly simple form of spatial structure can enhance the evolvability of cooperative behavior [Spector and Klein, 2006]. Spatial structure has also long been applied to evolutionary computation, often as a mechanism to preserve population diversity [Collins and Jefferson, 1991, Pettey, 1997, Fernandez et al., 2003, Folino et al., 2003]. But the models of

spatial structure previously employed in evolutionary computation have generally been more complex than the simple one-dimensional geographic scheme we used in our study of cooperation, and we were curious about whether the our scheme, which we called “trivial geography,” would provide benefits for genetic programming. In our contribution to last year’s meeting of the *Genetic Programming Theory and Practice* workshop we presented results showing that, indeed, trivial geography can enhance the problem-solving power of genetic programming systems, at least on the symbolic regression problems that we studied [Spector and Klein, 2005].

In the context of the new results presented above we are also curious about ways in which tags themselves might contribute to advances in genetic programming practice. Our comments about such possibilities here are speculative, but we believe there are several avenues worth pursuing.

One possibility is that tags, like trivial geography, might be used to better control the diversity of an evolving population. We have seen that tags can be used as surrogates for kin identifiers, and therefore as easily computed probabilistic indicators of genetic similarity or diversity. A great deal of recent work in genetic programming has focused on diversity metrics that may form the basis of mechanisms for diversity management [Burke et al., 2002]. A number of specific mechanisms for diversity management have also been developed, for example mechanisms based on mate selection [Fry and Tyrrell, 2003, Fry et al., 2005]. We speculate that multidimensional tags, which we have shown to facilitate the formation of clusters of cooperatively interacting agents, might also be used as the basis of mate selection schemes that combine in-breeding and out-breeding to balance exploration and exploitation.

Another way that these results might be exploited in genetic programming concerns work on the evolution of cooperative multiagent systems. Several researchers have previously used genetic programming to produce teams of cooperating agents, using a variety of mechanisms for team-member recognition and coordination [Luke and Spector, 1996, Soule, 2000, Brameier and Banzhaf, 2001]. The prospect raised by new theoretical work on tag-mediated cooperation is that the very simple mechanisms used for team-member recognition and coordination in these models might also serve as a solid foundation for cooperation in much more complex systems. We have previously shown that tag-mediated cooperation readily emerges in evolving multiagent swarms [Spector et al., 2005], and more recently we have seen tag-mediated cooperation emerge in evolving populations of blocks that can grow and divide. To the extent that any such multiagent systems can benefit from tag-based cooperation we might expect multidimensional tags to provide even greater advantages.

Finally, there is a sense in which genetic programming systems are themselves multiagent systems within which the “agents” — in this case the individual programs in an evolving population — might benefit from cooperation with one another. For example, in some of our prior work we allowed all of the

individuals in a population and across evolutionary time to access a shared indexed memory, or “culture” [Spector and Luke, 1996]. Several researchers have experimented with shared code in the form of automatically defined library functions [Racine et al., 1998, Ryan et al., 2004, Keijzer et al., 2005]. And several fitness sharing, niching, and crowding schemes involve the sharing or distribution of positions in fitness space [McKay, 2000, Gustafson, 2004]. Where there are common resources there are probably also opportunities for cooperation and coordination, for example to arbitrate the retention of commonly useful data or code fragments. Tags might provide a simple yet effective mechanism for achieving the requisite cooperation and coordination, and if they do so then again one would expect multidimensional tags to have even greater utility.

More generally the study of the evolution of cooperation, insofar as it challenges naive perspectives on Darwinian mechanisms, has helped to lay bare some of the fundamental dynamical properties of interdependent populations under natural selection. It would not be surprising, therefore, for such studies produce models with features that are applicable to evolutionary computation in general, and to genetic programming in particular. We note also that this research strategy, of borrowing specific pieces of theory from recent work in evolutionary biology and using them to enhance genetic programming systems, follows our general methodological suggestion from the first *Genetic Programming Theory and Practice* workshop [Spector, 2003].

7 Conclusions

We have presented the results of new experiments on computational models of tag-mediated cooperation. Our results demonstrate that the evolution of tag-mediated cooperation can be facilitated by the use of multidimensional tags. We analyzed these results as a form of “probabilistic kin recognition” and used the analogy of password-guessing to explain the observed patterns of cooperation and tolerance.

We discussed possible applications of these results to genetic programming. While our suggestions for application were speculative we noted that similar efforts have already borne fruit and that there are good reasons to be optimistic about future results. More generally, we argued that cooperative exchanges build networks of interaction that can support the evolution of adaptive complexity. For this reason we expect that the study of the evolution of cooperation will continue to produce important insights that can be applied to genetic programming.

Acknowledgments

This material is based upon work supported by the United States National Science Foundation under Grant No. 0308540 and Grant No. 0216344. Any

opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation. The authors also thank Jason Daida, Michael Korn, and the other participants in the Genetic Programming Theory and Practice IV Workshop for many helpful suggestions.

References

- [Axelrod and Hamilton, 1981] Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211:1390–1396.
- [Axelrod et al., 2004] Axelrod, R., Hammond, R. A., and Grafen, A. (2004). Altruism via kin-selection strategies that rely on arbitrary tags with which they coevolve. *Evolution*, 58:1833–1838.
- [Brameier and Banzhaf, 2001] Brameier, M. and Banzhaf, W. (2001). Evolving teams of predictors with linear genetic programming. *Genetic Programming and Evolvable Machines*, 2(4):381–407.
- [Burke et al., 2002] Burke, E., Gustafson, S., and Kendall, G. (2002). A survey and analysis of diversity measures in genetic programming. In Langdon, W. B., Cantú-Paz, E., Mathias, K., Roy, R., Davis, D., Poli, R., Balakrishnan, K., Honavar, V., Rudolph, G., Wegener, J., Bull, L., Potter, M. A., Schultz, A. C., Miller, J. F., Burke, E., and Jonoska, N., editors, *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 716–723, New York. Morgan Kaufmann Publishers.
- [Collins and Jefferson, 1991] Collins, R. J. and Jefferson, D. R. (1991). Selection in massively parallel genetic algorithms. In Belew, R. and Booker, L., editors, *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 249–256, San Mateo, CA. Morgan Kaufman.
- [Edmonds, 2006] Edmonds, B. (2006). The emergence of symbiotic groups resulting from skill-differentiation and tags. *Journal of Artificial Societies and Social Simulation*, 9(1).
- [Fernandez et al., 2003] Fernandez, F., Tomassini, M., and Vanneschi, L. (2003). An empirical study of multipopulation genetic programming. *Genetic Programming and Evolvable Machines*, 4(1):21–51.
- [Folino et al., 2003] Folino, G., Pizzuti, C., Spezzano, G., Vanneschi, L., and Tomassini, M. (2003). Diversity analysis in cellular and multipopulation genetic programming. In Sarker, R., Reynolds, R., Abbass, H., Tan, K. C., McKay, B., Essam, D., and Gedeon, T., editors, *Proceedings of the 2003 Congress on Evolutionary Computation CEC2003*, pages 305–311, Canberra. IEEE Press.
- [Fry et al., 2005] Fry, R., Smith, S. L., and Tyrrell, A. M. (2005). A self-adaptive mate selection model for genetic programming. In Corne, D., Michalewicz, Z., Dorigo, M., Eiben, G., Fogel, D., Fonseca, C., Greenwood, G., Chen, T. K., Raidl, G., Zalzal, A., Lucas, S., Paechter, B., Willies, J., Guervos, J. J. M., Eberbach, E., McKay, B., Channon, A., Tiwari, A., Volkert, L. G., Ashlock, D., and Schoenauer, M., editors, *Proceedings of the 2005 IEEE Congress on Evolutionary Computation*, volume 3, pages 2707–2714, Edinburgh, UK. IEEE Press.
- [Fry and Tyrrell, 2003] Fry, R. and Tyrrell, A. (2003). Enhancing the performance of GP using an ancestry-based mate selection scheme. In Cantú-Paz, E., Foster,

- J. A., Deb, K., Davis, D., Roy, R., O'Reilly, U.-M., Beyer, H.-G., Standish, R., Kendall, G., Wilson, S., Harman, M., Wegener, J., Dasgupta, D., Potter, M. A., Schultz, A. C., Dowsland, K., Jonoska, N., and Miller, J., editors, *Genetic and Evolutionary Computation – GECCO-2003*, volume 2724 of *LNCS*, pages 1804–1805, Chicago. Springer-Verlag.
- [Gustafson, 2004] Gustafson, S. (2004). *An Analysis of Diversity in Genetic Programming*. PhD thesis, School of Computer Science and Information Technology, University of Nottingham, Nottingham, England.
- [Hales, 2002] Hales, D. (2002). Smart agents don't need kin - evolving specialisation and cooperation with tags. Technical Report CPM Working Paper 02-89 (version 1), Centre for Policy Modelling.
- [Hales, D., 2005] Hales, D. (2005). Altruism For Free using Tags. In *Paris ECCS'05 Conference, Nov. 2005*.
- [Hamilton, 1963] Hamilton, W. D. (1963). The evolution of altruistic behavior. *American Naturalist*, 97:354–356.
- [Hamilton, 1964] Hamilton, W. D. (1964). The genetical evolution of social behaviour. i. *Journal of Theoretical Biology*, 7:1–16.
- [Holland, 1995] Holland, J. H. (1995). *Hidden Order: How Adaptation Builds Complexity*. Perseus Books.
- [Jansen and van Baalen, 2006] Jansen, V. A. A. and van Baalen, M. (2006). Altruism through beard chromodynamics. *Nature*, 440:663–666.
- [Keijzer et al., 2005] Keijzer, M., Ryan, C., Murphy, G., and Cattolico, M. (2005). Undirected training of run transferable libraries. In Keijzer, M., Tettamanzi, A., Collet, P., van Hemert, J. I., and Tomassini, M., editors, *Proceedings of the 8th European Conference on Genetic Programming*, volume 3447 of *Lecture Notes in Computer Science*, pages 361–370, Lausanne, Switzerland. Springer.
- [Luke and Spector, 1996] Luke, S. and Spector, L. (1996). Evolving teamwork and coordination with genetic programming. In Koza, J. R., Goldberg, D. E., Fogel, D. B., and Riolo, R. L., editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 150–156, Stanford University, CA, USA. MIT Press.
- [McKay, 2000] McKay, R. I. B. (2000). Fitness sharing in genetic programming. In Whitley, D., Goldberg, D., Cantu-Paz, E., Spector, L., Parmee, I., and Beyer, H.-G., editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2000)*, pages 435–442, Las Vegas, Nevada, USA. Morgan Kaufmann.
- [Michod, 2003] Michod, R. E. (2003). Cooperation and conflict mediation during the origin of multicellularity. In Hammerstein, P., editor, *Genetic and Cultural Evolution of Cooperation*, pages 291–307. The MIT Press, Cambridge, MA.
- [Nowak and Sigmund, 1998] Nowak, M. A. and Sigmund, K. (1998). Evolution of indirect reciprocity by image scoring. *Nature*, 393:573–577.
- [Pettey, 1997] Pettey, C. C. (1997). Diffusion (cellular) models. In Bäck, T., Fogel, D. B., and Michalewicz, Z., editors, *Handbook of Evolutionary Computation*, pages C6.4:1–6. Institute of Physics Publishing and Oxford University Press, Bristol, New York.
- [Racine et al., 1998] Racine, A., Schoenauer, M., and Dague, P. (1998). A dynamic lattice to evolve hierarchically shared subroutines: DL'GP. In Banzhaf, W., Poli, R., Schoenauer, M., and Fogarty, T. C., editors, *Proceedings of the First European Workshop on Genetic Programming*, volume 1391 of *LNCS*, pages 220–232, Paris. Springer-Verlag.

- [Riolo et al., 2001] Riolo, R. L., Cohen, M. D., and Axelrod, R. (2001). Evolution of cooperation without reciprocity. *Nature*, 414:441–443.
- [Riolo et al., 2002] Riolo, R. L., Cohen, M. D., and Axelrod, R. (2002). Riolo et al. reply. *Nature*, 418:500.
- [Roberts and Sherratt, 2002] Roberts, G. and Sherratt, T. N. (2002). Does similarity breed cooperation? *Nature*, 418:499–500.
- [Ryan et al., 2004] Ryan, C., Keijzer, M., and Cattolico, M. (2004). Favorable biasing of function sets using run transferable libraries. In O’Reilly, U.-M., Yu, T., Riolo, R. L., and Worzel, B., editors, *Genetic Programming Theory and Practice II*, chapter 7, pages 103–120. Springer, Ann Arbor.
- [Soule, 2000] Soule, T. (2000). Heterogeneity and specialization in evolving teams. In Whitley, D., Goldberg, D., Cantu-Paz, E., Spector, L., Parmee, I., and Beyer, H.-G., editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2000)*, pages 778–785, Las Vegas, Nevada, USA. Morgan Kaufmann.
- [Spector, 2003] Spector, L. (2003). An essay concerning human understanding of genetic programming. In Riolo, R. L. and Worzel, B., editors, *Genetic Programming Theory and Practice*, chapter 2, pages 11–24. Kluwer.
- [Spector and Klein, 2005] Spector, L. and Klein, J. (2005). Trivial geography in genetic programming. In Yu, T., Riolo, R. L., and Worzel, B., editors, *Genetic Programming Theory and Practice III*, volume 9 of *Genetic Programming*, chapter 8, pages 109–123. Springer, Ann Arbor.
- [Spector and Klein, 2006] Spector, L. and Klein, J. (2006). Genetic stability and territorial structure facilitate the evolution of tag-mediated altruism. *Artificial Life*, 12(4):1–8.
- [Spector et al., 2005] Spector, L., Klein, J., Perry, C., and Feinstein, M. (2005). Emergence of collective behavior in evolving populations of flying agents. *Genetic Programming and Evolvable Machines*, 6(1):111–125.
- [Spector and Luke, 1996] Spector, L. and Luke, S. (1996). Cultural transmission of information in genetic programming. In Koza, J. R., Goldberg, D. E., Fogel, D. B., and Riolo, R. L., editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 209–214, Stanford University, CA, USA. MIT Press.
- [Trivers, 1972] Trivers, R. (1972). The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46:35–57.