Université Libre de Bruxelles

*Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle*

IRIDIA

Incremental Particle Swarm-Guided Local Search for Continuous Optimization

Marco A. MONTES DE OCA, Ken VAN DEN ENDEN and
Thomas STÜTZLE

Incremental Particle Swarm-Guided Local Search for Continuous Optimization

Marco A. MONTES DE OCA[1]   mmontes@ulb.ac.be
Ken VAN DEN ENDEN[2]   kvdenden@vub.ac.be
Thomas STÜTZLE[1]   stuetzle@ulb.ac.be

[1]IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
[2]Vrije Universiteit Brussel, Brussels, Belgium

June 2008

Abstract

We present an algorithm that is inspired by theoretical and empirical results in social learning and swarm intelligence research. The algorithm is based on a framework that we call *incremental social learning*. In practical terms, the algorithm is a hybrid between a local search procedure and a particle swarm optimization algorithm with growing population size. The local search procedure provides rapid convergence to good solutions while the particle swarm algorithm enables a comprehensive exploration of the search space. We provide experimental evidence that shows that the algorithm can find good solutions very rapidly without compromising its global search capabilities.

1   Introduction

Many algorithms that have been successfully used for solving optimization problems can be thought of as being a collection of agents that interact with their environment and among themselves. In fact, this mental imagery has been used to design some of them. Some examples of this are genetic algorithms [18], ant colony optimization algorithms [12], and particle swarm optimization algorithms [20]. Following this same approach, we present here an algorithm that is inspired by results in social learning and swarm intelligence research. The algorithm is based on a multiagent learning framework that we call *incremental social learning* [23]. The multiagent scenario used for the design of the proposed algorithm is the following: A growing population of agents that can learn both individually and socially explores its environment in order to maximize the group's well-being. The strategy used by the agents is to use their social learning skills when they become part of the population and when learning individually is either too costly or deemed unproductive.

Technically speaking, the algorithm presented in this paper is a hybrid between the well-known Powell's direction set method [27] and a particle swarm optimization algorithm with growing population size [23]. Powell's method is

used as an agent's individual learning mechanism. If the local topography of the objective function (from an agent's perspective) allows it, this mechanism improves rapidly candidate solutions without any information exchange between agents. The particle swarm velocity- and position-update rules, which allow the global exploration of the search space, play the role of the agents' social learning mechanism. Together, these components form an incremental particle swarm-guided local search algorithm for solving continuous optimization problems (Section 3).

We compare the performance of the proposed algorithm with that of other four algorithms (Section 4). The proposed algorithm exhibits the fast convergence of Powell's method without compromising its global search capabilities (Section 5).

2   Incremental Social Learning

In previous work, we introduced a multiagent learning framework, called *incremental social learning*, that combines elements of social and individual learning in order to speed up learning and to allow scalability [23]. Social learning is a term that refers to the class of mechanisms that allow the transmission of knowledge between individuals without the use of genetic material [25, 10]. Individual learning, on the other hand, is a process that lets an individual acquire knowledge about its environment by interacting directly with it and without any social influence [1]. Applying social learning ideas to the multiagent learning problem is appealing because learning from others effectively gives agents a shortcut to adaptive information that otherwise would be expensive to acquire [21]. The framework is called incremental because it is based on a growing population of agents that adapts incrementally to the final multiagent environment. When a new agent is added to the population, it learns socially from more experienced agents and learns individually once it is part of it. This strategy is inspired by the fact that, in nature, newborn individuals are particularly favored by social learning because it allows them to learn many skills very rapidly from the adult individuals that surround them [15]. The algorithmic structure of the incremental social learning framework is outlined in Algorithm 1.

At the beginning, the environment and the primogenial population of agents are initialized. An agent addition schedule is used to control the rate at which agents are added to the environment. If no agents are to be added, the agents in the current population learn individually by directly interacting with the environment. If a new agent is scheduled to be part of the population, it learns socially from a subset of the agents in the population before it is added. The environment is then updated and the cycle is repeated until a stopping criterion is satisfied.

The learning process starts with a small number of agents as this reduces the interference caused by the co-existence of multiple learning agents. Agents are added according to a problem-dependent schedule in order to create time delays that allow agents that are already part of the population learn from the environment and to have something to teach to newcomers. When a new agent learns socially, it saves itself the effort required to learn individually what others already have. Incrementally growing the population size is also useful for allocating the minimum number of agents required to solve a particular problem.

Algorithm 1 Incremental social learning

/* Initialization */
$t \leftarrow 0$
Initialize environment $\mathbf{E}^t$
Initialize primogenial population of agents $\mathbf{X}^t$

/* Main loop */
while Stopping criteria not met **do**
  if Agent addition schedule or criterion is not met **then**
    $\mathbf{X}^{t+1} \leftarrow$ ilearn$(\mathbf{X}^t, \mathbf{E}^t)$ /* Individual learning */
  else
    Create new agent a_{new}
    slearn$(a_{new}, \mathbf{X}^t)$ /* Social learning */
    $\mathbf{X}^{t+1} \leftarrow \mathbf{X}^t \cup \{a_{new}\}$
  end if
  $\mathbf{E}^{t+1} \leftarrow$ update$(\mathbf{E}^t)$ /* Update environment */
  $t \leftarrow t + 1$
end while

---

3   Incremental Particle Swarm-Guided Local Search

The incremental social learning framework, although conceived for tackling multiagent learning problems, can also be used for designing population-based optimization algorithms. This is because, as we said in Section 1, the multiagent metaphor is usually useful for thinking about how this kind of algorithms work. In [23], a particle swarm optimization algorithm with growing population size (IPSO) was designed as an instantiation of the incremental social learning framework. In IPSO, individual learning is not implemented; instead, the algorithm comprises vertical and horizontal social learning mechanisms. Vertical social learning is a pattern of knowledge transmission between individuals of different generations while horizontal social learning happens between individuals of the same generation [6]. In IPSO, vertical social learning occurs when a new particle is added to the swarm and horizontal social learning is used instead of individual learning. IPSO exhibits a good solution quality vs. time trade-off since it finds solutions of at least the same quality than a particle swarm optimization algorithm with constant population size in fewer objective function evaluations.

Here we present an extension of IPSO that includes an individual learning mechanism. Individual learning is implemented as a local search procedure to simulate a particle's self-improvement process. The addition of an individual learning mechanism is justified not only from a practical point of view but also from a theoretical one. Indeed, since social learning is cheaper than individual learning, it is usually assumed to be always advantageous [21]. However, it has been argued that relying only on socially acquired knowledge is not always the best strategy [17]; instead, individuals should devote some of their time and energy to learn individually or to innovate [21].

As IPSO, the proposed algorithm is based on the particle swarm optimization algorithm in which particles (i.e., potential solutions to an optimization problem) move in the search space with a certain velocity. Each particle is attracted toward its own previous best position (with respect to an objective

function) and toward the best position found by the particles in its neighbor-
hood. Neighborhood relations are usually given in advance through a population
topology which can be defined by a graph $G = \{V, E\}$, where each vertex in V
corresponds to a particle in the swarm and each edge in E establishes a neigh-
bor relation between a pair of particles. The velocity and position updates of a
particle i over dimension j are as follows

$$v_{i,j}^{t+1} = \chi \cdot [v_{i,j}^t + \varphi_1 \cdot U_1 \cdot (p_{i,j}^t - x_{i,j}^t) + \varphi_2 \cdot U_2 \cdot (l_{i,j}^t - x_{i,j}^t)], \qquad (1)$$

and

$$x_{i,j}^{t+1} = x_{i,j}^t + v_{i,j}^{t+1}, \qquad (2)$$

where $v_{i,j}^t$ and $x_{i,j}^t$ are the particle's velocity and position at time step t respec-
tively, $p_{i,j}^t$ is the particle's best position so far, $l_{i,j}^t$ is the best position found by
the particle's neighbors, φ_1 and φ_2 are two parameters, U_1 and U_2 are two uni-
formly distributed random numbers in the range $[0, 1)$, and χ is a constriction
factor that is used in order to avoid an "explosion" of the particles' velocity.
Clerc and Kennedy [8] found the relation $\chi = 2k / \left| 2 - \varphi - \sqrt{\varphi^2 - 4\varphi} \right|$, where
$k \in [0, 1]$, and $\varphi = \varphi_1 + \varphi_2 > 4$, to compute it.

The local search procedure employed in the proposed algorithm is the well-
known Powell's direction set method [27] using Brent's technique [5] as the
auxiliary line minimization algorithm. The decision of using Powell's method
as the local search procedure is based on performance considerations. The
proposed algorithm calls repeatedly the local search procedure and thus an
efficient one was needed. In this sense, Powell's method quadratic convergence
can be very advantageous if the objective function is locally quadratic [28].

The last component of the proposed algorithm is the vertical social learning
mechanism which is implemented as a rule that moves the new particle's pre-
vious best position from its initial random location in the search space to one
that is closer to the previous best position of a particle that serves as a "model"
to imitate. The rule is applied in a component-wise fashion as follows

$$p_{new,j}' = p_{new,j} + U \cdot (p_{model,j} - p_{new,j}), \qquad (3)$$

where $p_{new,j}'$ is the new particle's updated previous best position, $p_{new,j}$ is the
new particle's original previous best position, $p_{model,j}$ is the model's previous
best position and U is a uniformly distributed random number in the range
$[0, 1)$.

The combination of the local search procedure and the particle swarm op-
timization algorithm with growing population size, as sketched in Algorithm
2, produces an incremental particle swarm-guided local search algorithm whose
operation is illustrated in Figure 1.

The proposed algorithm is designed for taking advantage of features that the
objective function may have. For example, if the objective function is convex and
separable, it should be optimized in a single local search run without the need
of using the exploration capabilities provided by the particle swarm algorithm.
Similarly, if the objective function has a few local optima, starting with only
one particle ensures the minimal use of function evaluations for the exploration
of the space. However, the objective function may also have features for which
local search and small initial populations are ineffective. For example, if an
objective function had large plateaus, the local search procedure would return

---

Algorithm 2 Incremental Particle Swarm-Guided Local Search

Input: Objective function f and maximum number of iterations t_{max}

  /* Initialization */
  Create initial particle p_1 and add it to the set of particles $\mathcal{P}$ which is initially empty
  Initialize position vector $\boldsymbol{p_1.x}$ to random values within the search range
  Initialize velocity vector $\boldsymbol{p_1.v}$ to zero
  Set $\boldsymbol{p_1.pb} = \boldsymbol{p_1.x}$

  /* Main Loop */
  Set $t = 0$ and $k = 1$
  repeat
    /* Individual learning */
    for $i = 1$ to k **do**
      Improve $\boldsymbol{p_i.pb}$ through a local search procedure
    end for

    /* Horizontal social learning */
    for $i = 1$ to k **do**
      Move $\boldsymbol{p_i.x}$ using Eqs. 1 and 2
      if $f(\boldsymbol{p_i.x})$ is better than $f(\boldsymbol{p_i.pb})$ **then**
        Set $\boldsymbol{p_i.pb} = \boldsymbol{p_i.x}$
      end if
    end for

    /* Population growth and vertical social learning */
    if Agent addition criterion is met **then**
      Create particle p_{k+1} and add it to the set of particles $\mathcal{P}$
      Initialize position vector $\boldsymbol{p_{k+1}.x}$ using Eq. 3
      Initialize velocity vector $\boldsymbol{p_{k+1}.v}$ to zero
      Set $\boldsymbol{p_{k+1}.pb} = \boldsymbol{p_{k+1}.x}$
      Set $k = k + 1$
    end if
    Set $t = t + 1$
    Set $\boldsymbol{sol} = \underset{p_i \in \mathcal{P}}{\mathrm{argmin}}\, f(\boldsymbol{p_i.pb})$
  until $f(\boldsymbol{sol})$ is good enough or $t = t_{max}$

---

with no solution improvement after having wasted several function evaluations. Under these circumstances, the algorithm would start making progress (thanks to the exploration capabilities of the particle swarm component) until a critical population size is reached.

4   Experimental Setup

The performance of the incremental particle swarm-guided local search algorithm (labeled IPSOLS) is compared to that of the following algorithms:
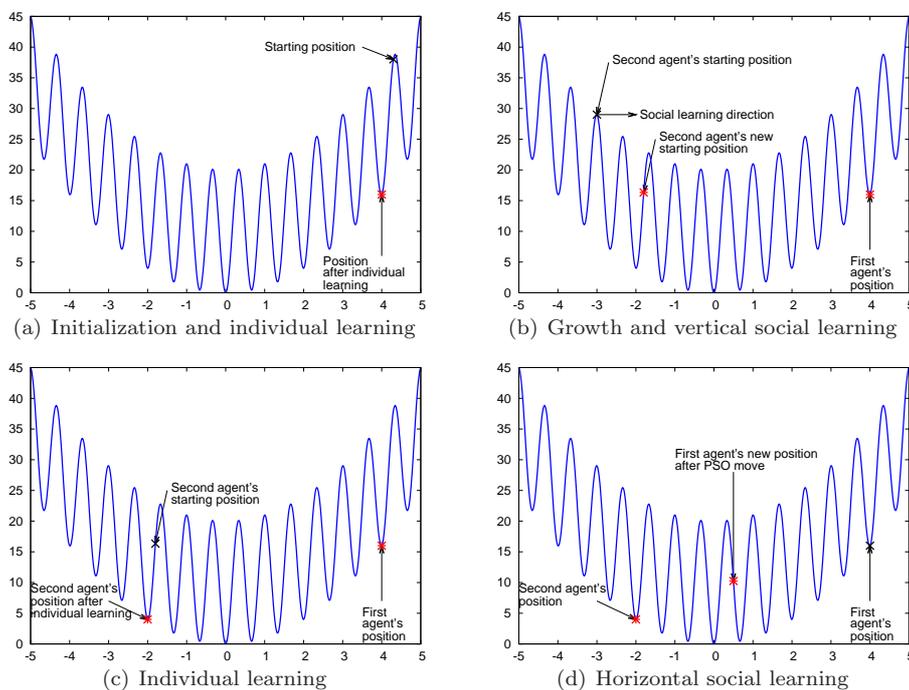
Figure 1: The incremental particle swarm-guided local search algorithm at work. In Figure (a), an initial particle is randomly placed in the search space. Immediately after, the first round of individual learning (local search) is run, moving the particle to, or close to, a local minimum. This is followed by a horizontal learning round (a normal PSO move) that does not have any effect in a swarm of one particle (the particles' velocity are initialized to zero). In Figure (b), the size of the population is increased. The vertical social learning rule (Eq. 3) is used to place the new particle in the search space. In Figure (c), the second iteration of the algorithm begins by running a round of individual learning. In Figure (d), through a round of horizontal learning, particles can move and discover more promising areas of the search space. The algorithm continues intertwining individual and social learning procedures until a stopping criterion is satisfied.

1. A traditional particle swarm optimization algorithm with constant population size (labeled PSO). Three population sizes of 10, 100 and 1000 particles are used.

2. An incremental particle swarm optimization algorithm as described in Section 3 (labeled IPSO). In [23], IPSO proved to work well with a fast agent addition schedule. Accordingly, we use a fast particle addition schedule in which a new particle is added every iteration (this schedule is also used in IPSOLS). The maximum size of the swarm is set to 1000 particles. Whenever a new particle is added, the particle that serves as model to imitate is the best of the swarm.

3. A hybrid particle swarm optimization algorithm with local search (la-

beled PSOLS). It is a constant population size particle swarm algorithm in which the particles' previous best positions undergo an improvement phase (via Powell's method) before the velocity update rule is applied. Three population sizes of 10, 100 and 1000 particles are used.

4. A random restart local search algorithm (labeled RLS). Every time the local search procedure (Powell's method) converges, it is restarted from a newly generated random solution. The best solution found so far is considered to be the output of the algorithm.

All particle swarm-based algorithms (PSO, IPSO, PSOLS and IPSOLS) are run with two population topologies: a fully connected topology, in which each particle is a neighbor to all others including itself, and the so-called ring topology, in which each particle has two neighbors apart from itself. In the incremental algorithms, the new particle is randomly placed within the topological structure. Other parameter settings are $\varphi_1 = \varphi_2 = 2.05$ and $\chi = 0.7298$.

Powell's method has also a number of parameters to be defined. The so-called tolerance, used to stop the procedure once a very small difference between two solutions is detected, is set to 0.01. In case such a difference is not found, the procedure is stopped after a maximum number of iterations is reached. In our experiments, this parameter is set to 10. Other settings were explored but no significant differences in the results were found. Finally, different step sizes (used for the line minimization procedure) were explored. The values tried for this parameter were 0.1, 1, 10, 20, and 33% of the length of the corresponding search range.

A set of six benchmark functions are used in our experiments. Their mathematical formulation and search ranges are listed in Table 1. In all cases, we used their 100-dimensional instantiations (i.e., $n = 100$).

Table 1: Benchmark optimization problems

Name	Definition	Range
Ackley	$-20e^{-0.2\sqrt{\frac{1}{n}\sum_{i=1}^{n} x_i^2}} - e^{\frac{1}{n}\sum_{i=1}^{n}\cos(2\pi x_i)} + 20 + e$	$[-32,32]^n$
Rastrigin	$10n + \sum_{i=1}^{n}\left(x_i^2 - 10\cos(2\pi x_i)\right)$	$[-5.12,5.12]^n$
Rosenbrock	$\sum_{i=1}^{n-1}\left[100(x_{i+1} - x_i^2)^2 + (x_i - 1)^2\right]$	$[-30,30]^n$
Expanded Schaffer	$ES(\boldsymbol{x}) = \sum_{i=1}^{n-1} S(x_i, x_{i+1}) + S(x_n, x_1),\text{where}$ $S(x,y) = 0.5 + \frac{\sin^2\left(\sqrt{x^2+y^2}\right)-0.5}{(1+0.001(x^2+y^2))^2}$	$[-100,100]^n$
Schwefel	$418.9829n + \sum_{i=1}^{n} -x_i\sin(\sqrt{|x_i|})$	$[-500,500]^n$
Sphere	$\sum_{i=1}^{n} x_i^2$	$[-100,100]^n$

Our data is based on 100 independent runs of up to 10^6 function evaluations each. In each run, all benchmark functions (except Schwefel's) were randomly shifted within the specified range. Particles or trial points outside the search range were forced to stay within by putting them on the boundary.

5   Results

The distribution of the solution quality obtained by the compared algorithms after 10^6 function evaluations is shown in Figure 2. These results correspond to the case in which all particle swarm-based algorithms use a ring topology and the local search step size is equal to 20% the length of the search range. In [24], the reader can find the complete set of results, including those from the statistical significance tests performed on the algorithms' output data[1]. The algorithms without a local search component are placed on the left side of each plot. This arrangement reveals that the algorithms with a local search component find solutions of equal or better quality than the algorithms without it. The use of a local search component also makes the algorithms more robust to changes in the population topology, that is, the algorithms with a local search component do not show a significant change in their performance if a different topology is used. Only PSO and IPSO are affected by a change in the population topology. For example, the PSO algorithm with 10 particles finds better solutions with the ring topology than with the fully connected topology while the opposite happens if larger populations are used. In any case, independently of the topology used, the algorithms with local search outperform those without it.

IPSOLS (at the center of all plots) performs at least as good as the other particle swarm-based algorithms with local search. In three problems (Expanded Schaffer, Schwefel, and Rastrigin), IPSOLS obtains better solutions than all other algorithms. With Ackley's function, IPSOLS, PSOLS and RLS obtain comparable results. With Rosenbrock's function, IPSOLS and PSOLS with a population of 10 particles obtain the best results. It is important to note, however, that even after 10^5 function evaluations, the only algorithm that consistently finds high-quality solutions to all the problems in our benchmark suite is IPSOLS.

Figure 3 shows the development over time of the median solution quality obtained by the compared algorithms. In order to make the plots more readable, we consider only the algorithms that obtained the best results after 10^6 function evaluations. These plots show more clearly the effects of the incremental and socially guided use of local search procedures. As expected, IPSOLS exhibits the same performance as RLS during the first 10^4 function evaluations. This happens because IPSOLS starts, just as RLS does, performing a local search from a random initial solution. Once the first local search procedure stops making progress, IPSOLS and RLS start differentiating. The beneficial effects of relying on socially acquired information are specially clear in the results obtained for Ackley, Rastrigin, Expanded Schaffer, and Schwefel's problems.

The results obtained while solving the Sphere problem show clearly that starting with one single particle is beneficial if the objective function is separable and convex. In this case, as can be seen in the plots, IPSOLS and RLS have exactly the same performance. This is so because a single run of the local search procedure is able to find the optimal solution with no extra algorithmic overhead. The difference between IPSOLS and RLS is that IPSOLS is able to adapt and cope with the difficulties posed by the features of other objective functions.

All the algorithms that contain Powell's method as an algorithmic compo-

---

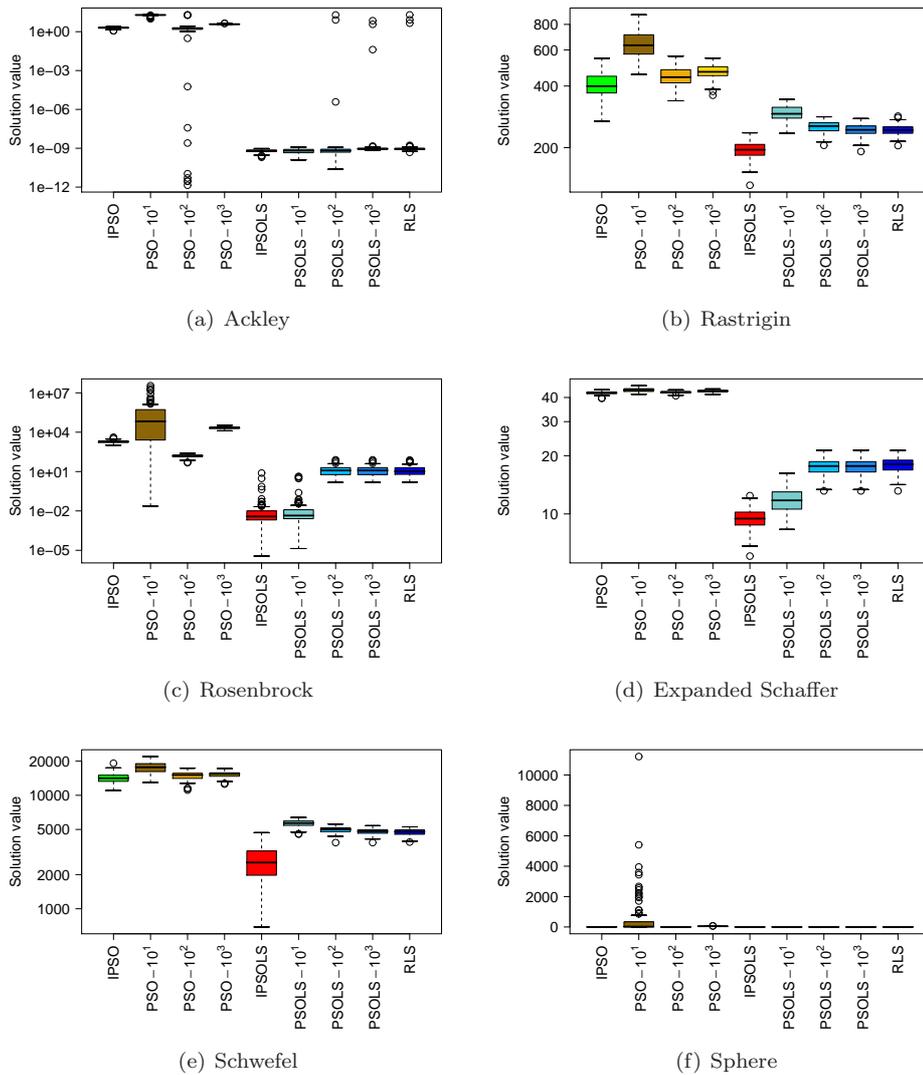[1]Unless otherwise noted, all our statements are supported by statistical evidence.

Figure 2: Boxplots showing the distribution of the solution quality obtained by the compared algorithms after 10^6 function evaluations. These results correspond to the case in which all particle swarm-based algorithms used a ring topology. The numbers next to the labels indicate the population size used.

(a) Ackley

(b) Rastrigin

(c) Rosenbrock
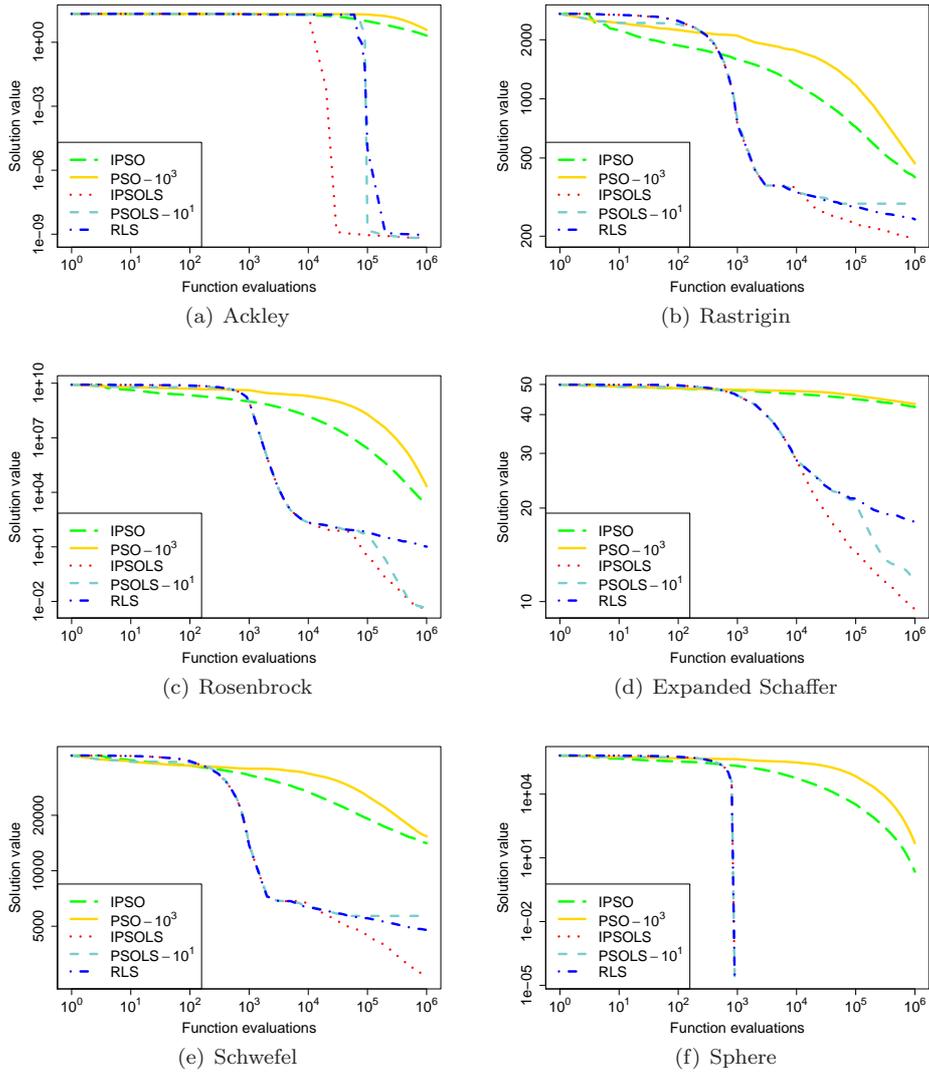
(d) Expanded Schaffer

(e) Schwefel

(f) Sphere

Figure 3: Development of the solution quality over time (median values). These results correspond to the case in which all particle swarm-based algorithms used a ring topology. The numbers next to the labels indicate the population size used.

nent are affected by the step size of the subsidiary line minimization algorithm (in our case, Brent's algorithm). Figure 4 shows the median solution quality development over time for different step sizes using RLS. As can be seen in the plots, the actual effect of the step size depends on the objective function. In our study, the greatest differences are seen on multimodal test problems (Ackley, Rastrigin, Expanded Schaffer, and Schwefel) while there is practically no difference on unimodal functions (Rosenbrock and Sphere). This result suggests that, on multimodal problems, large step sizes allow the local search procedure to "jump" over local optima, which effectively increases the chance of finding good quality local optima. The downside of this behavior is that if the step size parameter is set to the wrong value, a poor performance can be obtained.

6   Related Work and Discussion

The incremental particle swarm-guided local search algorithm (IPSOLS) can be seen as a population of agents with the capability of learning by themselves and from others. In order to obtain the greatest reward, individuals must use a strategy that allows them to intertwine their two modes of learning in a productive way. In IPSOLS, individual learning is preferred over social learning. Agents learn socially only when individual learning is either too costly or deemed unproductive. Practically speaking this means that local search is the mechanism that is always tried first. If local search alone cannot solve the problem at hand satisfactorily, either because it has converged to a local optimum (i.e., individual learning is deemed unproductive because no further improvement can be made using only local search) or because the maximum number of iterations of the local search procedure has been reached (i.e., individual learning is too costly because many functions evaluations are used), then social learning is used. Of course, it is possible to intertwine the two modes of learning in a different way. A study on the utility of different learning strategies seems a fruitful avenue for future research.

IPSOLS shares a number of features with other algorithms. In particular, a time-varying population size and a subsidiary local search algorithm. Research on population (re)sizing has been a topic of study within the field of evolutionary computation for many years. From that experience, it is now usually accepted that the population size in evolutionary algorithms should be proportional to the problem's difficulty [22]. The problem is that we usually know little about the real-world problem's difficulty. The approach taken in the design of the incremental particle swarm-guided local search algorithm (i.e., to start with small populations) makes sense because if the problem is not so difficult, a growing population will offer the best solution quality vs. time trade-off. If the problem is difficult, there will be a time penalty to pay for reaching very high quality solutions; however, acceptable solutions may be found early in a run.

Practically all resizing strategies consider the possibility of reducing the size of the population during an algorithm's run (see e.g. [2, 19, 4, 13, 14, 9]). An exception to this approach is the work of Auger and Hansen [3] in which the population size of a CMA-ES algorithm is doubled each time it is restarted. The feature that IPSOLS shares with Auger and Hansen's approach is that no population reduction mechanism is implemented; however, IPSOLS does not use restarts and has a slower population growth rate. Nevertheless, it may be that

(a) Ackley

(b) Rastrigin

(c) Rosenbrock
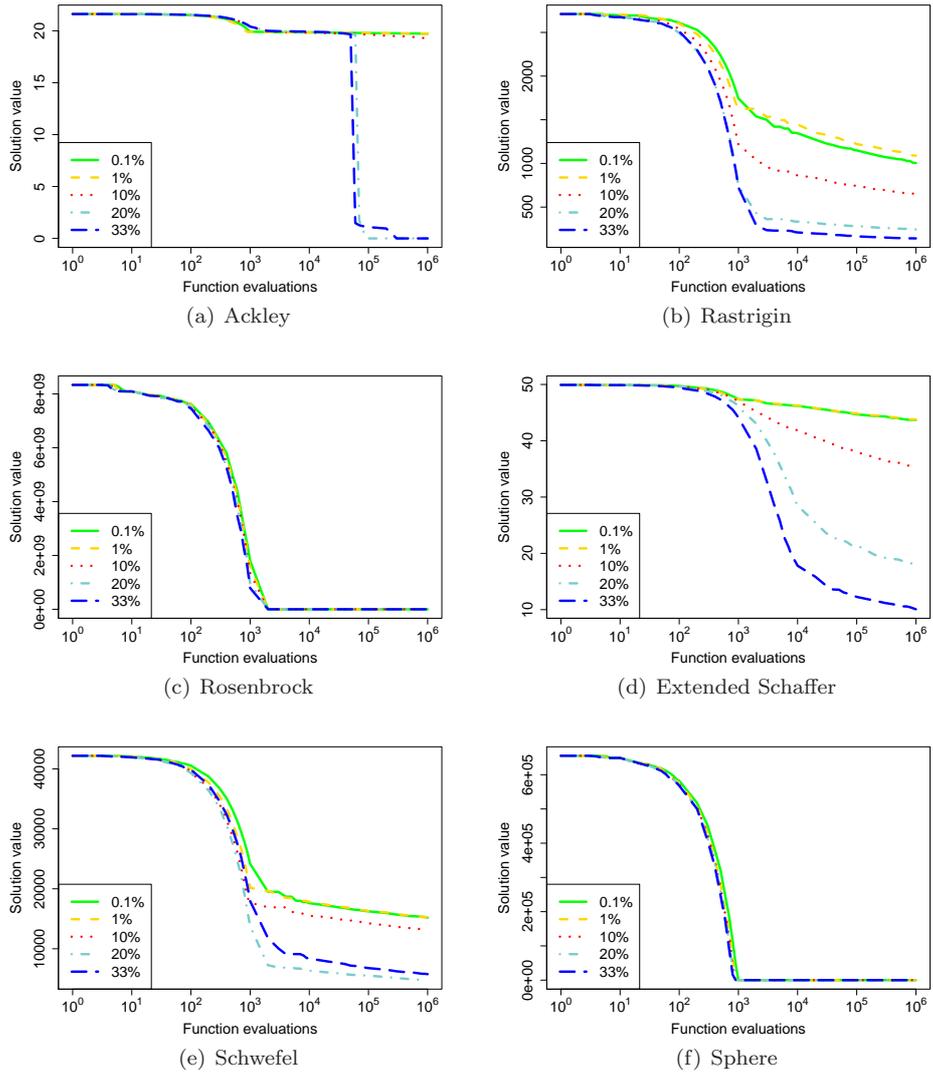
(d) Extended Schaffer

(e) Schwefel

(f) Sphere

Figure 4: Development of the solution quality over time (median values) for different local search step sizes using RLS.

the performance of IPSOLS is improved by decreasing the population size from time to time. This is because decreasing the population size may save many function evaluations. In terms of solution quality, however, it is not expected that this may impact negatively on the quality of the solutions found since the choice of removal would be done on converged particles (particles very close to each other with very small velocity and already in a local minimum).

The idea of combining local search techniques with particle swarm optimization algorithms comes from the observation that particles are attracted by their previous best positions. It is usually thought that the better the attractors are, the higher the chances to find even better solutions. The goal of most hybrid algorithms is thus to accelerate the placement of the particles' previous best positions in good spots. For example, Chen et al. [7] combined a particle swarm algorithm with a hill-climbing local search procedure; Gimmler et al [16] experiment with Nelder and Mead's simplex method as well as with Powell's method; Das et al. [11] also use Nelder and Mead's simplex method and propose the inclusion of an estimate of the local gradient into the particles' velocity update rule. Petalas et al. [26] report experiments with several local search-particle swarm combination schemes. All these previous approaches try to enhance the performance of the particle swarm algorithm; IPSOLS, on the contrary, guides a local search procedure in order to better explore the search space. By doing it incrementally, IPSOLS can exploit the features of the objective function.

7 Conclusions

A hybrid algorithm for solving continuous optimization problems, called incremental particle swarm-guided local search, has been presented. The algorithm has two main components: (i) a particle swarm optimization algorithm with growing population size and (ii) a subsidiary local search procedure that operates on the particles' previous best positions. We presented the experimental evidence that allows us to conclude that the algorithm exhibits a very good exploitation behavior that does not compromise its global search capabilities.

The design of the proposed algorithm is inspired by results in social learning and swarm intelligence research. Using a multiagent systems metaphor, the algorithm can be thought of as a growing population of agents that can learn both socially and individually. The strategy used by the agents to intertwine these two modes of learning is the following: when a naive agent becomes part of the main population, it learns from more experienced agents. Once it is part of the main population, an agent learns socially only when learning individually is considered to be too costly or when doing it is deemed unproductive.

Acknowledgments

References

[1] K. Aoki, Y. Wakano, and M. W. Feldman. The emergence of social learning in a temporally changing environment: A theoretical model. *Current Anthropology*, 46(2):334–340, 2005.

[2] J. Arabas, Z. Michalewicz, and J. J. Mulawka. GAVaPS – A genetic algorithm with varying population size. In *Proceedings of the IEEE Conference on Evolutionary Computation*, pages 73–78, Piscataway, NJ, USA, 1994. IEEE Press.

[3] A. Auger and N. Hansen. A restart CMA evolution strategy with increasing population size. In *Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2005)*, pages 1769–1776, Piscataway, NJ, USA, 2005. IEEE Press.

[4] T. Bäck, A. E. Eiben, and N. A. L. van der Vaart. An empirical study on GAs "without parameters". In *Parallel Problem Solving from Nature, PPSN VI*, pages 315–324, Berlin, Germany, 2000. Springer-Verlag.

[5] R. P. Brent. *Algorithms for Minimization Without Derivatives*. Prentice-Hall, Englewood Cliffs, NJ, USA, 1973.

[6] L. L. Cavalli-Sforza and M. W. Feldman. *Cultural Transmission and Evolution. A Quantitative Approach*. Princeton University Press, Princeton, NJ, USA, 1981.

[7] J. Chen, Z. Qin, Y. Liu, and J. Lu. Particle swarm optimization with local search. In *Proceedings of the International Conference on Neural Networks and Brain (ICNN&B 2005)*, pages 481–484, Piscataway, NJ, USA, 2005. IEEE Press.

[8] M. Clerc and J. Kennedy. The particle swarm–explosion, stability, and convergence in a multidimensional complex space. *IEEE Transactions on Evolutionary Computation*, 6(1):58–73, 2002.

[9] A. L. V. Coelho and D. G. de Oliveira. Dynamically tuning the population size in particle swarm optimization. In *Proceedings of the ACM Symposium on Applied Computing (SAC'08)*, pages 1782–1787, New York, NY, USA, 2008. ACM Press.

[10] D. Curran and C. O'Riordan. Increasing population diversity through cultural learning. *Adaptive Behavior*, 14(4):315–338, 2006.

[11] S. Das, P. Koduru, M. Gui, M. Cochran, A. Wareing, S. M. Welch, and B. R. Babin. Adding local search to particle swarm optimization. In *Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2005)*, pages 428–433, Piscataway, NJ, USA, 2006. IEEE Press.

[12] M. Dorigo and T. Stützle. *Ant Colony Optimization*. Bradford Books. MIT Press, Cambridge, MA, USA, 2004.

[13] A. E. Eiben, E. Marchiori, and V. A. Valkó. Evolutionary algorithms with on-the-fly population size adjustment. In *Parallel Problem Solving from Nature, PPSN VIII*, pages 41–50, Berlin, Germany, 2004. Springer-Verlag.

[14] C. Fernandes and A. Rosa. Self-regulated population size in evolutionary algorithms. In *Parallel Problem Solving from Nature, PPSN IX*, pages 920–929, Berlin, Germany, 2006. Springer-Verlag.

[15] B. G. Galef Jr. and K. N. Laland. Social learning in animals: Empirical studies and theoretical models. *BioScience*, 55(6):489–499, 2005.

[16] J. Gimmler, T. Stützle, and T. E. Exner. Hybrid particle swarm optimization: An examination of the influence of iterative improvement algorithms on performance. In M. Dorigo et al., editors, *LNCS 4150. Ant Colony Optimization and Swarm Intelligence. 5th International Workshop, ANTS 2006*, pages 436–443, Berlin, Germany, 2006. Springer-Verlag.

[17] L.-A. Giraldeau, T. J. Valone, and J. J. Templeton. Potential disadvantages of using socially acquired information. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357:1559–1566, 2002.

[18] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison-Wesley, Reading, MA, USA, 1989.

[19] G. R. Harik and F. G. Lobo. A parameter-less genetic algorithm. In W. Banzhaf et al., editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 1999)*, pages 258–265, San Francisco, CA, USA, 1999. Morgan Kaufmann.

[20] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of IEEE International Conference on Neural Networks*, pages 1942–1948, Piscataway, NJ, USA, 1995. IEEE Press.

[21] K. N. Laland. Social learning strategies. *Learning & Behavior*, 32(1):4–14, 2004.

[22] F. G. Lobo and C. F. Lima. *Parameter Setting in Evolutionary Algorithms*, volume 54/2007 of *Studies in Computational Intelligence*, chapter Adaptive Population Sizing Schemes in Genetic Algorithms, pages 185–204. Springer, Berlin, Germany, 2007.

[23] M. A. Montes de Oca and T. Stützle. Towards incremental social learning in optimization and multiagent systems. In *Workshop on Evolutionary Computation and Multiagent Systems Simulation of the Genetic and Evolutionary Computation Conference (GECCO 2008)*, New York, NY, USA, 2008. ACM Press. In press.

[24] M. A. Montes de Oca, K. Van den Enden, and T. Stützle. Incremental particle swarm-guided local search for continuous optimization: Complete results, 2008. Supplementary information page at `http://iridia.ulb.ac.be/supp/IridiaSupp2008-013/`.

[25] C. L. Nehaniv and K. Dautenhahn, editors. *Imitation and Social Learning in Robots, Humans and Animals: Behavioral, Social and Communicative Dimensions.* Cambridge University Press, Cambridge, United Kingdom, 2007.

[26] Y. G. Petalas, K. E. Parsopoulos, and M. N. Vrahatis. Memetic particle swarm optimization. *Annals of Operations Research*, 156(1):99–127, 2007.

[27] M. J. D. Powell. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *The Computer Journal*, 7(2):155–162, 1964.

[28] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C. The Art of Scientific Computing.* Cambridge University Press, second edition, 1992.