UC Irvine

UC Irvine Previously Published Works

Title

Fortify particle swarm optimizer (PSO) with principal components analysis: A case study in improving bound-handling for optimizing high-dimensional and complex problems

Permalink

<https://escholarship.org/uc/item/5sn5c0mz>

ISBN

9781424478347

Authors

Chu, W
Gao, X
Sorooshian, S

Publication Date

2011-08-29

DOI

10.1109/CEC.2011.5949812

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Fortify Particle Swarm Optimizer (PSO) with Principal Components Analysis

A case study in improving bound-handling for optimizing high-dimensional and complex problems

Wei Chu, Xiaogang Gao, and Soroosh Sorooshian

Department of Civil and Environmental Engineering
University of California
Irvine, CA, USA
wchu2@uci.edu

Abstract—It is reported that the absorbing bound-handling approach may paralyze PSO when it is applied to high-dimensional and complex problems. In this study, we introduce principal components analysis (PCA) into PSO in order to remedy the problem caused by the absorbing bound-handling approach. The experiments on 100-D composition functions demonstrate the effectiveness of PCA. Furthermore, the strong influence of bound-handling on PSO is also evidently revealed by the results. The fact that none of the studied bound-handling methods excels on all of the benchmark functions highlights the necessity of developing more sophisticated and robust bound-handling approaches that can facilitate the application of PSO on high-dimensional problems.

Keywords—constrained optimization; bound-handling approach; partial swarm optimization; principle components analysis; high-dimensional problems

I. INTRODUCTION

Particle Swarm Optimizer (PSO) [1], one of the most successful and popular evolutionary algorithms, has been intensively studied, modified and applied to many practical problems[2-4]. PSO simulates the dynamics of communication in a swarm of insects or a school of fish and guides a population of particles searching over the feasible space of a objective function for the global optimum/optima. Many research efforts on improving PSO have been drawn to the study of several aspects, including dynamical inertia weights, swarm structure/dynamics, learning scheme, etc. However, for constrained high-dimensional problems, how to handle particles that fly outside the boundary of the search space is not a trivial issue, since boundary violation increases enormously as dimensionality increases [5].

In fact, our recent study [6] reveals that proper bound-handling approach is a prerequisite for PSO to function properly on constrained high-dimensional problems. Three most popular bound-handling approaches-reflecting, absorbing, and random, were studied. Experimental results show that only the reflecting approach enables PSO to function properly on all of the test functions, whereas random and absorbing approaches may paralyze PSO on some benchmark functions. The mechanisms that are responsible for the failure of the two

bound-handling approaches are investigated and summarized in [6]. In the failure of the random approach, the population evolves extremely slow and fitness values of particles tend to oscillate, due to the fact that the bound violation is prevailing in high-dimensional bounded search space and the random sample procedure depresses the PSO process; In the failure of the absorbing approach, all of the particles tend to adhere to the bounds and, eventually, converge to a single point on the boundary on certain dimensions which are referred to as “lost dimensions”.

In this study, we introduce Principal Components Analysis (PCA) as a tool to identify the occurrence of lost dimensions and to remedy the adverse effects. Experimental results show that, PCA enables PSO to function properly with absorbing bound-handling approach and even yield better results on some test functions compared with the reflecting bound-handling approach.

Furthermore, our study illustrates the importance of bound handling when applying PSO to high-dimensional and complex problems. Further investigation is required to fully understand the effect of bound handling and to design effective and robust bound-handling approaches.

II. PRINCIPAL COMPONENTS ANALYSIS

A. The Concept of PCA

Principal Components Analysis is a multivariate analysis tool that transforms a given dataset to a new orthogonal independent coordinate system so that the first coordinate (called the first principal component, PC) has the largest variance of projections from the dataset; the second coordinate has the second largest variance and so on. In certain cases, some lower-rank PCs will have negligible variances, which means that the dataset cannot span the dimensions represented by these lower-rank PCs and these dimensions are the lost dimensions.

For instance, if we have a population of particles:

$$\{x_i \in R^m\}_{i=1}^{NS}$$

where NS is the size of the population, and m is the dimensionality of the problem.

We can have the population transformed to the PC coordinate system $\{y_i \in R^m\}_{i=1}^{NS}$ by

$$y_i = Ax_i \quad (1)$$

$$\lambda_k P_k = BP_k \quad (2)$$

where matrix A has PCs ($P_k, k=1, \dots, m$) as its columns, B is the covariance matrix of the particle population in the original coordinate system, and eigenvalue λ_k is the data variance along P_k .

If we have $\frac{\lambda_L}{\sum_{j=1}^m \lambda_j} \rightarrow 0$, the vector component $y_{ip} \rightarrow$ constant for every $p \in \{L \leq p \leq m\}$. This indicates that the dataset is located within a subspace spanned by $(L-1)$ independent orthogonal vectors (PCs) denoted as R^{L-1} .

B. The Algorithm of Implementing PCA in PSO

Adopted from [7], a module for checking the lost dimensions and remedy its adverse effects is described below:

We use a matrix, C , to represent the current positions and the best history positions of all particles in the swarm, such that coordinates of each position is a column in C . Therefore, C has the size of $m \times 2NS$:

$$C = [c_{ij}] \text{ with } i=1, \dots, m, \text{ and } j=1, \dots, 2NS \quad (3)$$

At the end of every loop of evolution, the following steps are applied to C :

Step 1. Check the dimensionality of the space spanned C

(1) Transform the original coordinated system to a normalized coordinated system by centering and normalizing each row of and get C' . This normalization can reduce the effect of differences in the units of different parameters in real problems. The following operations of this module are all discussed as in this normalized space.

(2) Calculate the covariance matrix of C' and denote it as R . Obtain eigenvectors and eigenvalues of R . Each eigenvector is a principal component (PC), and its corresponding eigenvalue measures the variance of C' along the direction defined by that PC.

(3) By examining eigenvalues, we can determine if there is any dimension lost and, if yes, how many are lost. Theoretically, the population should fully span the m -dimensional parameter space, which means that C' should have comparable variance along all of the directions defined by every PC. If the variance along the direction of one PC is

too small, it means that the population does not span well over that direction, and that dimension is lost. On a lost dimension, we can use the centroid of C' to represent all of the particles since they have very small variance on this dimension. In a m -dimensional space, for an isotropic C' , the expected variance along each PC is $1/m$ of the total variance. Therefore, in this study, if a PC has variance less than 10% of the expected variance, we treat it as a lost dimension.

Step 2. Search along lost dimensions. For each lost dimension detected in Step 1, do the following random search along the PC that represents it:

(1) Sample a new position from the positive side of along the PC.

$$\bar{x}' = \bar{c}' + ar\bar{l} \quad (4)$$

where a is a random number generated from normal distribution with mean = 2 and variance = 1, $\bar{l}$ is the unit vector representing the PC, and r is the radius of C' , defined by

$$r = \max(r_i) \text{ with } i=1, \dots, m, \quad (5)$$

and

$$r_i = \max(|c'_{ij} - c'_{ik}|) \text{ with } j, k=1, \dots, 2NS \quad (6)$$

Then, transform $\bar{x}'$ back to the original coordinates and evaluate the function at it. If the function value is smaller than that of the worst particle in the swarm, move the worst particle to this new position, and the search on this PC is over. Otherwise, discard this new position and continue to (2).

(2) Sample a new position from the negative side of along the PC.

$$\bar{x}' = \bar{c}' - ar\bar{l} \quad (7)$$

Again, transform back to the original coordinates and evaluate the function at it. If the function value is smaller than that of the worst particle in the swarm, move the worst particle to this new position, and the search on this PC is over. Otherwise, discard this new position. The search on this PC terminates.

In summary, Step 2 is designed to quickly explore over lost dimensions to see if there is evident slope along them. If there is, the random sampling is likely to capture it, and the new position mingled into the swarm will enable the swarm to search along this lost dimension. The flow chart of the module is presented in Figure 1.

III. EXPERIMENTS AND RESULTS

The module described above is integrated in a standard PSO, as described in [6]. This modified PSO with the absorbing bound-handling approach is benchmarked by the same suite of 100-D composition functions, CF1-6 [8], and the

results are compared with those by the standard PSO with the reflecting and absorbing bound-handling approaches.

A. Benchmark Functions

As shown in [6], in order to explore and examine the possible problems of bound handling in PSO when applied to complicated real-world applications, composition benchmark functions with irregular fitness landscapes are preferable whereas the widely used standard benchmark functions that have symmetric fitness landscapes and global optima at the center of the search space can sometimes obscure problems.

Figure 1. Flowchart of the module that implements PCA to identify lost dimensions and remedy their adverse effects.

Each of the composition functions are composed of ten standard functions, chosen from Sphere, Ackley, Griewank, Rastrigin, and Weierstrass functions. The ten standard functions are weighted and summed to form the composition function.

In general, their fitness landscapes have no symmetry or any other obvious patterns, which poses tremendous difficulties when the dimensionality is high. Figure 2 shows the fitness landscapes of the functions in two-dimensional space. CF1 is

constructed with 10 sphere functions, with a evident global attractive region. CF2 and CF3 are constructed with multimodal Griewank and Rastrigin functions and, therefore, have increased complexity. The attractive regions of global minima are much harder to find than the one of CF1. CF4, CF5 and CF6 are all hybrid functions, constructed with more than one standard function. The global area of CF4 are occupied by the narrow optimum basin of Ackley's function and the global minimum is hidden in a bad fitness area. In CF4 and CF5, the global areas are dominated by Rastrigin function with many local minima concealing the global minimum. Compared with CF5, CF6 has a narrower global attractive region and flatter local optimum areas. More details about the construction of composition functions can be found in [8].

Figure 2. Response surfaces of benchmark functions in 2-D. (a) through (f) correspond to CF1-CF6.

B. Experimental Setting

The algorithmic coefficients of PSO are set at widely used values. The inertia weight is held at a constant $c_0 = 0.5$; c_1 and c_2 are both set at 2, and $v_{\max}$ equals half of the search range. A population of 1,000 particles is randomly generated with a uniform distribution in the search space. For all functions, the lower and upper bounds of every dimension are -5 and 5, respectively. The global minimum and the major local minima are all randomly set within the range of $[-4.5, 4.5]$ in every dimension.

IV. RESULTS AND DISCUSSIONS

A. Experimental Results

For each of the three PSO algorithms (the modified PSO with absorbing bound-handling approach, the standard PSO with absorbing bound-handling approach, and the standard

Figure 3. The best fitness value versus the number of function evaluations on 100-D CF1 to CF6 functions: standard PSO with reflecting bound-handling (left column), modified PSO with absorbing bound-handling (middle), and standard PSO with absorbing (right) bound-handling approaches.

PSO with reflecting approach), fifty independent runs are conducted on each of the test functions. The mean and standard deviation of the final best function values are listed in Table I. In [6], we have already revealed that the absorbing approach paralyzes the standard PSO on some of the test functions, and

the results shown in Table I indicate that the standard PSO with absorbing approach yields the worst results on all of the benchmark functions. Therefore, we only compare the modified PSO with absorbing approach and the standard PSO with the reflecting approach. The Wilcoxon rank sum test was

conducted to test the significance of differences between the modified PSO runs and the standard PSO runs. It turned out that the modified PSO outperforms the standard PSO on functions CF3 and CF4 at the 5% significance level, whereas the standard PSO retrieve better results than does the modified PSO on functions CF1 and CF5 at the same confidence level. For CF2, the two algorithms yield results that are not significantly different. As for CF6, both algorithms are trapped by the local minimum of the 10th standard function which is at the center of the search space.

TABLE I. EXPERIMENTAL RESLUTS

	<i>PSO Reflecting</i>	<i>Modified PSO Absorbing</i>	<i>PSO Absorbing</i>	<i>P-value</i>
CF1	32.00 ± 47.12	66.73 ± 30.43	156.7 ± 54.80	1.158e-04
CF2	210.1 ± 112.9	236.9 ± 117.2	312.5 ± 117.7	Not significant
CF3	434.9 ± 362.5	249.8 ± 127.7	717.3 ± 319.5	0.0476
CF4	900.0 ± 0.000	248.8 ± 121.2	917.5 ± 40.40	1.211e-12
CF5	74.20 ± 82.26	136.5 ± 50.45	185.2 ± 95.05	1.800e-3
CF6	900.0 ± 0.000	900.0 ± 0.000	901.9 ± 7.780	Not significant

a. The bolded numbers are the best results at the 5% significance level by the Wilcoxon rank sum test .

The fitness curves of all fifty runs are plotted in Figure 3, demonstrating the efficiency of each PSO. The purpose of plotting all curves is to graphically demonstrate the variations across the ensemble of fifty independent runs. On CF3 and CF4, the modified PSO achieve better final fitness values compared with the other two PSOs. Especially on CF4, only the modified PSO successfully escapes the trap posed by the 10th standard function which is at the center of the search space. In addition to the better final value, modified PSO also exhibits higher efficiency with greater speeds of decreasing objective function values. For the standard PSO with reflecting approach, the fitness curves converge to one or several lines on some functions, which is caused by the fact that the swarm was trapped to one of the local minima and converged there. In contrast, the modified PSO with absorbing approach has much more diverse fitness curves, which means that the swarm was not frequently trapped to local minima. CF6 is so difficult that both PSOs were consistently trapped to the minimum of the 10th standard function, which is a sphere function.

B. Discussions

The experimental results substantiate that the PCA can help PSO overcome the effects of lost dimensions when using the absorbing bound-handling approach on high-dimensional and complex problems. With the absorbing approach, the modified PSO show comparable skill to that of the standard PSO with reflecting bound-handling approach.

Conducting PCA does cause additional computation cost. However, when the dimension is lower than 1000, the computation of the PCA is very fast (less than 4 seconds on a computer with Pentium 4 CPU 2.53GHz and 1.00 GB of RAM), which is much faster compared with the run time of many high-dimensional practical problems. In addition, the PCA is only conducted at the end of each evolution loop. Therefore, in most cases of high-dimensional optimization , conducting of PCA is only a trivial part (less than 0.1%) of the total computation cost, which is the case in the experiments of this study.

The experimental results highlight the importance of proper treatment of bound violation. For the same function, different bound-handling approach may yield significantly different results. Unfortunately, none of the tested bound-handling approach consistently excels on all the test functions. Therefore, we have to admit that more deep and intensive study on the effect of bound-handling is required for applying PSO to optimization of high-dimensional and complex problems.

The issue of lost dimension is not unique to PSO, and it may also occur in other evolutionary algorithms [9]. PCA is a very effective tool of identify and remedy this problem, and it deserves attentions from developers and users of evolutionary optimization algorithms.

ACKNOWLEDGMENT

The authors are grateful to Dr. P. N. Suganthan for kindly providing Matlab codes of the composition benchmark functions and the PSO algorithm used in this study.

REFERENCES

- [1] J. Kennedy and R.C. Eberhart, "Swarm Intelligence", Morgan Kaufmann Academic Press, San Matteo, CA, 2001.
- [2] R.C. Eberhar and S. Yuhui, "Tracking and optimizing dynamic systems with particle swarms, in: Proceedings of the Congress on Evolutionary Computation", 2001, pp. 94-100.
- [3] A. Salman, I. Ahmad, and S. Al-Madani, "Particle swarm optimization for task assignment problem", Microprocessors and Microsystems, vol. 26, pp. 363-371, 2002.
- [4] J. Robinson and Y. Rahmat-Samii, Particle swarm optimization in electromagnetics, IEEE Transactions on Antennas and Propagation vol. 52, pp. 397-407, 2004.
- [5] S. Helwig, R. Wanka, Particle swarm optimization in high-dimensional bounded search spaces, in: Proceedings of IEEE Swarm Intelligence Symposium, pp. 198-205, 2007.
- [6] W. Chu, X. Gao, and S. Sorooshian, "Bound handling strategy is critical to the success of particle swarm optimization on high-dimensional complex problems," Information Sciences, 2010, doi:10.1016/j.ins.2010.11.030.
- [7] W. Chu, X. Gao, and S. Sorooshian, "A new evolutionary search strategy for global optimization of high-dimensional problems," Information Sciences, under review.
- [8] J.J. Liang, P.N. Suganthan, K. Deb, "Novel composition test functions for numerical global optimization", in: Proceedings of IEEE Swarm Intelligence Symposium, 2005, pp. 68-75.
- [9] W. Chu, X. Gao and S. Sorooshian, "A solution to the crucial problem of population degeneration in high-dimensional evolutionary optimization", IEEE Systems Journal, under review.