Phys. Med. Biol. 50 (2005) 3491–3514 doi:10.1088/0031-9155/50/15/002

A particle swarm optimization algorithm for beam angle selection in intensity-modulated radiotherapy planning

Yongjie Li[1], Dezhong Yao[1], Jonathan Yao[2] and Wufan Chen[1]

[1] School of Life Science and Technology, University of Electronic Science and Technology of China, Chengdu 610054, People's Republic of China
[2] Topslane Inc., Pleasant Hill, CA 94523, USA

E-mail: dyao@uestc.edu.cn

Received 7 February 2005, in final form 13 June 2005
Published 13 July 2005
Online at stacks.iop.org/PMB/50/3491

Abstract
Automatic beam angle selection is an important but challenging problem for intensity-modulated radiation therapy (IMRT) planning. Though many efforts have been made, it is still not very satisfactory in clinical IMRT practice because of overextensive computation of the inverse problem. In this paper, a new technique named BASPSO (Beam Angle Selection with a Particle Swarm Optimization algorithm) is presented to improve the efficiency of the beam angle optimization problem. Originally developed as a tool for simulating social behaviour, the particle swarm optimization (PSO) algorithm is a relatively new population-based evolutionary optimization technique first introduced by Kennedy and Eberhart in 1995. In the proposed BASPSO, the beam angles are optimized using PSO by treating each beam configuration as a particle (individual), and the beam intensity maps for each beam configuration are optimized using the conjugate gradient (CG) algorithm. These two optimization processes are implemented iteratively. The performance of each individual is evaluated by a fitness value calculated with a physical objective function. A population of these individuals is evolved by cooperation and competition among the individuals themselves through generations. The optimization results of a simulated case with known optimal beam angles and two clinical cases (a prostate case and a head-and-neck case) show that PSO is valid and efficient and can speed up the beam angle optimization process. Furthermore, the performance comparisons based on the preliminary results indicate that, as a whole, the PSO-based algorithm seems to outperform, or at least compete with, the GA-based algorithm in computation time and robustness. In conclusion, the reported work suggested that the introduced PSO algorithm could act as a new promising solution to the beam angle optimization problem and potentially other optimization problems in IMRT, though further studies need to be investigated.

1. Introduction

Intensity-modulated radiotherapy (IMRT) is a powerful technology to potentially improve the therapeutic ratio by using modulated beams to produce highly three-dimensionally conformal dose distributions to the target, while sparing organs-at-risk (OARs) and normal tissues as much as possible. The conventional IMRT planning starts with the selection of suitable beam angles, followed by an optimization of beam intensity maps using inverse optimization methods under the guidance of an objective function and the predefined dose prescription (Webb 2000, Chvetsova *et al* 2005). Currently, the selection of beam angles is generally based upon the experience of the human planner. Several trial-and-error attempts are normally needed in order to find a group of acceptable beam angles. Although this leads to good treatment plans, they are not necessarily optimal (Rowbottom *et al* 1998). This is partly because intensity modulation can often compensate for suboptimally placed beams and optimum beam configurations may sometimes be counterintuitive (Stein *et al* 1997, Wang *et al* 2004).

The requirement of automatic selection of optimal beam angles is particularly highlighted in intensity-modulated radiation therapy (IMRT), in which highly three-dimensionally conformal dose distributions are expected. It has been proved that the selection of suitable beam angles is most valuable for plans with a small number of beams ($\leqslant$5) (Bortfeld and Schlegel 1993, Soderstrom and Brahme 1995, Stein *et al* 1997, Rowbottom *et al* 1998), and is also clinically meaningful for plans with a large number of beams ($\geqslant$9) in some complicated cases, where the tumour volume surrounds a critical organ, or is surrounded by multiple critical organs (Pugachev *et al* 2000, 2001, Gaede *et al* 2004). Furthermore, it has been demonstrated that the plans with fewer but optimized beams could be as good as or better than plans using a larger number of unoptimized beams (Stein *et al* 1997, Hou *et al* 2003, Wang *et al* 2004).

To date, extensive efforts have been devoted to facilitating the technique of automatic beam angle selection for IMRT planning (Soderstrom and Brahme 1992, 1995, Bortfeld and Schlegel 1993, Gokhale *et al* 1994, Stein *et al* 1997, Haas *et al* 1998, Rowbottom *et al* 1998, 2001, Pugachev *et al* 2000, 2001, Pugachev and Xing 2001, Das *et al* 2003, Meedt *et al* 2003, Djajaputra *et al* 2003, Hou *et al* 2003, Gaede *et al* 2004, Li *et al* 2004, Meyer *et al* 2004, Souza *et al* 2004, Schreibmann and Xing 2004, Schreibmann *et al* 2004, Wang *et al* 2004). Though fruitful results have been achieved, the improvements are still not satisfactory, especially on the optimization efficiency because of the handicap of extensive computation in IMRT optimization. Tradeoffs always have to be made between the optimization accuracy and the computation time. For example, to make the optimization problem tractable, previous investigators employed compromising strategies for accelerating IMRT dose calculations (Stein *et al* 1997, Pugachev *et al* 2000, John Cho *et al* 2004) and for identifying the most preferred beam directions (Pugachev and Xing 2002). The influence of these approximations on the optimization results still needs future investigation (Stein *et al* 1997, Wang *et al* 2004).

Among those published optimization algorithms adopted in radiotherapy planning, there exist a bunch of natural algorithms, which borrow the intelligence from the organisms in nature and inherit a global search mechanism. Some typical examples are the genetic algorithm (GA) (Hou *et al* 2003, Schreibmann *et al* 2004, Li *et al* 2004) and simulated annealing (SA) (Bortfeld and Schlegel 1993, Stein *et al* 1997, Pugachev *et al* 2000, 2001, Rowbottom *et al* 2001, Pugachev and Xing 2002, Meyer *et al* 2004). The results reported were promising and encouraging for further research in this field. However, SA is intrinsically a sequential algorithm because it describes a Markovian process, and it is consequently difficult to take advantage of potential efficiency improvements that could result from parallel computation.

In addition, as recognized by many other researchers (Hou *et al* 2003, Meedt *et al* 2003, Wang *et al* 2004), SA algorithms explore a large number of candidate solutions, which may inevitably prolong the optimization process. It has been commonly accepted that the SA algorithm is good at local convergence, but its ability to converge to the global minima is relatively poor. In contrast, GA is intrinsically a population-based parallel algorithm and usually discovers the promising regions of search space very quickly, but it often needs too many computations to reach local minima (Goldberg 1989, Renders and Flasse 1996, Gen and Cheng 2000).

Particle swarm optimization (PSO) is a recent addition to evolutionary algorithms (EA) first introduced by Kennedy and Eberhart in 1995. The foundation of PSO is based on the hypothesis that social sharing of information among conspecifics offers an evolutionary advantage. Partially inspired by animal social behaviours such as flocking of birds, PSO originally intends to graphically mimic the graceful way in which they find their food sources and save themselves from predators (Eberhart and Kennedy 1995). PSO is a population-based stochastic optimization paradigm, in which each agent, named *particle*, of the population, named *swarm*, is thought of as a collision-proof bird and used to represent a potential solution. PSO is similar in some ways to other existing EAs, such as GA, but is defined in a social context as opposed to a biological context. Compared to GA and SA, PSO is generally characterized as simple in concept, easy to implement and computationally efficient (Eberhart and Shi 2001). In addition, PSO appears to be robust to control parameters (Kennedy and Spears 1998, Trelea 2003).

As an emerging technology, PSO has received a lot of attention in recent years, in some conferences, such as Congress on Evolutionary Computation (CEC) and Genetic and Evolutionary Computation Conference (GECCO), devoted solely to this topic (Hu *et al* 2004). To date, PSO has been successfully applied in many fields, such as artificial neural network training (Eberhart and Shi 1998b, Messerschmidt and Engelbrecht 2004), the optimal power flow (OPF) problem (Abido 2002), the task assignment problem (Salman *et al* 2002), the unit commitment (UC) problem (Ting *et al* 2003), quantitative structure–activity relationship (QSAR) model construction (Cedeno and Agrafiotis 2003), multiple sequence alignment (MSA) (Rasmussen and Krink 2004), multi-modal biomedical image registration (Wachowiak *et al* 2004), multi-objective optimization (Coello Coello *et al* 2004), electromagnetic optimizations (Robinson and Rahmat-Samii 2003, Boeringer and Werner 2004), blind source separation (BBS) (Gao and Xie 2004), protein motif discovery (Chang *et al* 2004), etc. Most of these applications demonstrated that PSO could give competitive or even better results in a faster and cheaper way, compared to other heuristic methods such as GA. Generally, all the application areas that the other evolutionary computation techniques are good at are the good application areas for PSO (Eberhart and Shi 2004).

To the best of our knowledge, the increasingly popularized PSO algorithm has not been investigated in the area of radiotherapy planning so far. In this paper, a new approach named BASPSO (Beam Angle Selection using the Particle Swarm Optimization algorithm) is proposed in order to improve the efficiency of the beam angle optimization for IMRT planning.

The remainder of this paper is organized as follows. The principle of the PSO, the strategies of parameter selection and the implementation of the problem-dependent PSO are presented in section 2 in detail. In section 3, a simulated test case with known optimal beam angles and two clinical cases are employed to verify the validity and efficiency of BASPSO by comparing the computation time with a GA-based algorithm and by comparing the dose distributions of the optimized plan with those of the manual plan. Finally, some discussion and conclusions are presented in section 4.

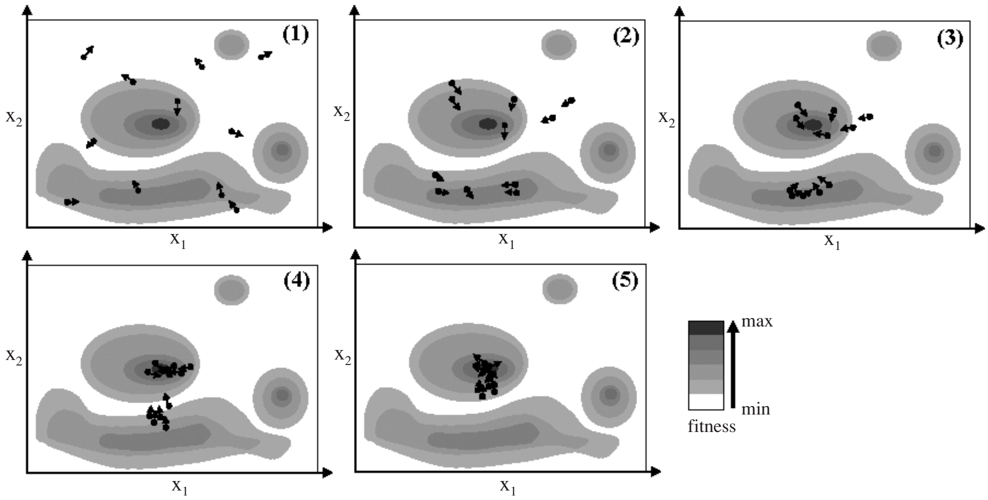

Figure 1. An illustration of the convergence of PSO in a two-dimensional search space.

2. Materials and methods

In order to simplify the optimization and decrease the computation burden, the beam angle selection and the beam intensity map optimization are normally separated into two iterative steps (Stein *et al* 1997, Pugachev *et al* 2000, Hou *et al* 2003, Wang *et al* 2004). A similar strategy is adopted in this paper. In BASPSO, a specified number of angle candidates contained in particles are selected from a discrete angle candidate pool using PSO evolution strategies, and then the corresponding intensity maps of these angles are optimized using a conjugate gradient (CG) method under the guidance of an objective function. The principles and the implementation are described in detail in this section.

2.1. Particle swarm optimization

PSO is similar to EA techniques in that a population of potential solutions to the problem under consideration is used to probe the search space. However, in PSO, each agent (i.e. particle) representing a potential solution moves in the search space and adaptively updates its velocity and position according to its own flying experience and the flying experience of its neighbours, aiming at a better position for itself. Moreover, each individual has a memory, remembering the best position of the search space it has ever visited. Thus, its movement is an aggregated acceleration towards its best previously visited position and towards the best individual of a topological neighbourhood. For a better understanding of the global optimization paradigm of PSO, the convergence of PSO in a two-dimensional search space is graphically illustrated in figure 1 (Krink 2002).

Two variants of the PSO algorithm were developed (Parsopoulos and Vrahatis 2002): one with a global neighbourhood, and one with a local neighbourhood. According to the global variant, each particle moves towards its best previous position and towards the best particle in the whole swarm. On the other hand, according to the local variant, each particle moves towards its best previous position and towards the best particle in its restricted neighbourhood. In this study, the global variant is adopted, though the local variant with a carefully designed local neighbourhood may provide better properties for specified applications (Kennedy and Mendes 2002).

PSO starts with the random initialization of a population of individuals (particles) in the search space and works on the social behaviour of the particles in the swarm. Each particle is treated as a mass-less and volume-less point in a D-dimensional space. The ith particle is represented as $\vec{x}_i = (x_{i1}, x_{i2}, \ldots, x_{iD})$. The best previous position of the ith particle that gives the best fitness value is represented as $\vec{p}_i = (p_{i1}, p_{i2}, \ldots, p_{iD})$. The best particle among all the particles in the population is represented by $\vec{p}_g = (p_{g1}, p_{g2}, \ldots, p_{gD})$. Velocity, the rate of position change for the ith particle, is represented as $\vec{v}_i = (v_{i1}, v_{i2}, \ldots, v_{iD})$. At every iteration, the velocity and the position of each particle are updated by using the two best values according to the following equations:

$$v_{id}^{k+1} = \omega \cdot v_{id}^{k} + c_1 \cdot \mathrm{rand}_1() \cdot \left(p_{id}^{k} - x_{id}^{k}\right) + c_2 \cdot \mathrm{rand}_2() \cdot \left(p_{gd}^{k} - x_{id}^{k}\right) \tag{1}$$

$$x_{id}^{k+1} = x_{id}^{k} + v_{id}^{k+1}. \tag{2}$$

Here k is the iteration number, $d = 1, 2, \ldots, D$, $i = 1, 2, \ldots, N$, and N is the size of the population (swarm). c_1 and c_2 are two positive values called acceleration constants, $\mathrm{rand}_1()$ and $\mathrm{rand}_2()$ are two independent random numbers that uniformly distribute between 0 and 1 and are used to stochastically vary the relative pull of $\vec{p}_i$ and $\vec{p}_g$ (Clerc and Kennedy 2002). The introduction of such random elements into the optimization is intended to simulate the slightly unpredictable component of natural swarm behaviour. ω is the inertial weight introduced by Shi and Eberhart in order to improve the performance of the particle swarm optimizer (Shi and Eberhart 1998b).

The first term on the right-hand side of equation (1) refers to the inertial effects of the movement, so that an individual is not allowed to change the flight direction too fast. The second and third terms refer to the memory of the individual and the group as a whole, respectively. Summarily, equation (1) is used to calculate the particle's new velocity according to its previous velocity and the distances of its current position from both its own best historical position and the group's best historical position. Such an adjustment of the particle's movement through the space causes it to search around the two best positions. This means that if a particle discovers a promising new solution, all the other particles will move closer to it, exploring the region more thoroughly in the process (van den Bergh and Engelbrecht 2004).

The manipulation of a particle's velocity according to equation (1) is regarded as the central element of the entire optimization, thorough understanding of which is the key to understand the optimization as a whole (Robinson and Rahmat-Samii 2004). The second term in equation (1) introduces a stochastic tendency to return towards the individual's previous best position. Psychologically, it is a tendency to remember and return to regions in the psychological space that have demonstrated promise, i.e., to combinations of beliefs that have seemed good or reinforcing (Kennedy 1997). The third term in equation (1) is derived from the successes of others, which is considered as a *social influence* term. Effectively, the variable $\vec{p}_g$ keeps moving, as neighbours find better and better points in the search space. As a particle swarm searches over time, individuals are drawn towards one another's successes, with the usual result being the clustering of individuals in optimal regions of the space (figure 1). Clearly, a particle further from the global best or the local best is more strongly pulled from its location, and therefore moves more rapidly than a closer particle. The particles continue to gain speed in the direction of the locations of greatest fitness until they pass over them. At that point they begin to be pulled back in the opposite direction. It is this overflying of the local and global extrema that is believed to be one secret to the success of PSO (Kennedy and Eberhart 1997, Kennedy 1998, Robinson and Rahmat-Samii 2004).

During the evolution of the swarm, the performance of each particle is evaluated by a predefined fitness function, which is problem dependent. Assuming the fitness function

```
Initialize population
Do
    For k = 1 to Maximum Iteration Number K
        For i = 1 to Population Size N
            Update the personal best position using Eq. (3)
            Update the global best position using Eq. (4)
            For d = 1 to Dimension D
                Update the velocity v_id using Eq. (1)
                Update the position x_id using Eq. (2)
            Next d
        Next i
    Next k
Until termination criterion is met
```

Figure 2. The pseudocode of the particle swarm algorithm.

$\mathrm{Fitness}(\vec{x}_i)$, described in section 2.4, to be maximized, the personal best position of each particle $\vec{p}_i$ and the global best position $\vec{p}_g$ are updated after each iteration (generation) using the following two equations, respectively:

$$\vec{p}_i^{k+1} = \begin{cases} \vec{x}_i^{k+1}: & \mathrm{Fitness}\big(\vec{x}_i^{k+1}\big) > \mathrm{Fitness}\big(\vec{p}_i^{k+1}\big) \\ \vec{p}_i^{k}: & \mathrm{Fitness}\big(\vec{x}_i^{k+1}\big) \leqslant \mathrm{Fitness}\big(\vec{p}_i^{k+1}\big) \end{cases} \tag{3}$$

$$\vec{p}_g^{k+1} = \arg\max_{p_i} \mathrm{Fitness}\big(\vec{p}_i^{k+1}\big). \tag{4}$$

The pseudocode of the PSO algorithm described above is outlined in figure 2, from which we can observe that PSO can be implemented in a few lines of a computer code, and it requires only primitive mathematical operators.

The real strength of the particle swarm derives from the interactions among particles as they search the space collaboratively (Clerc and Kennedy 2002). Numerous examples coming from nature have demonstrated that social sharing of information among the individuals of a population may provide an evolutionary advantage. This was the core idea behind the development of PSO (Parsopoulos and Vrahatis 2002). In a social science context, PSO enhances individuals by both sharing information between each other and their individually acquired knowledge.

Originally, PSO works based on the social adaptation of knowledge, and all individuals are considered to be of the same generation. In contrast, GA works based on the evolution from generation to generation, so the changes of individuals in a single generation are not considered. There are several key issues that should be pointed out between PSO and GA (Rahmat-Samii 2003). (1) One advantage of the PSO over the GA is its algorithmic simplicity. The GA typically requires three major operators: selection, crossover, and mutation. In PSO, however, there is one simple operator: velocity calculation. (2) Both GA and PSO have several numerical parameters that need to be carefully selected. However, the robustness to control parameters makes their selection even easier for PSO (Trelea 2003). Clerc and Kennedy (2002) even suggested that no problem-specific parameters might need to be specified. (3) Another difference between GA and PSO is the ability to control convergence. Crossover and mutation rates can subtly affect the convergence of GA. De Jong *et al* (1997) had noted that the crossover might not perform especially well on functions featuring high modality. However, Kennedy and Spears (1998) had demonstrated that PSO was hardly affected at all by increasing the modality of problems. Furthermore, it has been shown that the decrease of

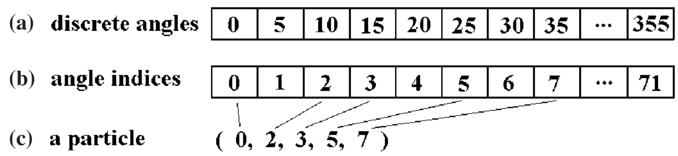

Figure 3. An illustration of particle representation.

inertial weight dramatically increases the swarm's convergence. Stagnation may occur in GA. In PSO, however, this effect can be controlled or prevented easily, for example, by decreasing the inertial weight during the evolution process (Eberhart and Shi 1998a, Clerc and Kennedy 2002). (4) The conceptual bases of these two algorithms rest upon two completely different philosophies. PSO is based upon social swarm behaviour, and GA is based upon genetic encoding and natural selection. These conceptual differences could be potentially utilized to develop some hybrid methods with a better performance, which will be further discussed in the last section of the paper.

2.2. *PSO for beam angle optimization*

In this section, we describe the formulations of the PSO algorithm for the beam angle optimization in IMRT planning. The beams in this study are restricted to the coplanar ones, and the beam angles are specified according to the International Electrotechnical Commission (IEC) convention. The search space covering the total 360° gantry rotational angles are discretized into equally spaced directions with a given angle step, such as 5° or 10°. These discrete angles are called trial angles and they comprise the search space of the optimization problem, an illustration of which is shown in figure 3(a).

One of the key issues to successfully apply PSO to a specified engineering problem is the representation scheme, i.e., finding a suitable mapping between the problem solution and the PSO particle. Those discrete angles in this paper are positive integers and comprise a non-continuous integer search space. In order to solve this optimization problem using the PSO algorithm, these discrete angles are indexed by a group of continuous integers in increasing order, which is demonstrated in figure 3(b). These indices constitute a continuous integer search space with a range of [0, 71] when the whole 2π gantry rotational angle is divided by angle steps of 5°. Then the particles are represented using these indices. The dimension of the search space D, i.e., the number of elements in a particle, is equal to the planner-defined beam number in a treatment plan, and the values of the elements (i.e., the coordinate components of the position) are the angle indices. Figure 3(c) illustrates a representation of a particle $\vec{x}_i = (x_{i1}, x_{i2}, x_{i3}, x_{i4}, x_{i5})$ as (0, 2, 3, 5, 7) for a five-beam plan, in which five indices are contained: 0, 2, 3, 5 and 7, corresponding to the beams with angles 0°, 10°, 15°, 25° and 35°.

Since the original PSO works on a real-valued search space, especially, the particle positions (i.e., the angle indices in this paper) calculated using equation (2) are real numbers, a conversion is needed between the real-valued positions and the positive-integer-valued indices. To facilitate the process, this conversion is simply realized by dropping the sign and truncating the real-valued position to the closest integer. This conversion is only applied to x_{id}^{k+1} in equation (2) after it is calculated, and the rest of the items in equations (1) and (2) are the same as the original PSO. Such facilitation has been proved feasible and does not seem to significantly affect the performance of the PSO algorithm (Salman *et al* 2002, Laskari *et al* 2002, Parsopoulos and Vrahatis 2002).

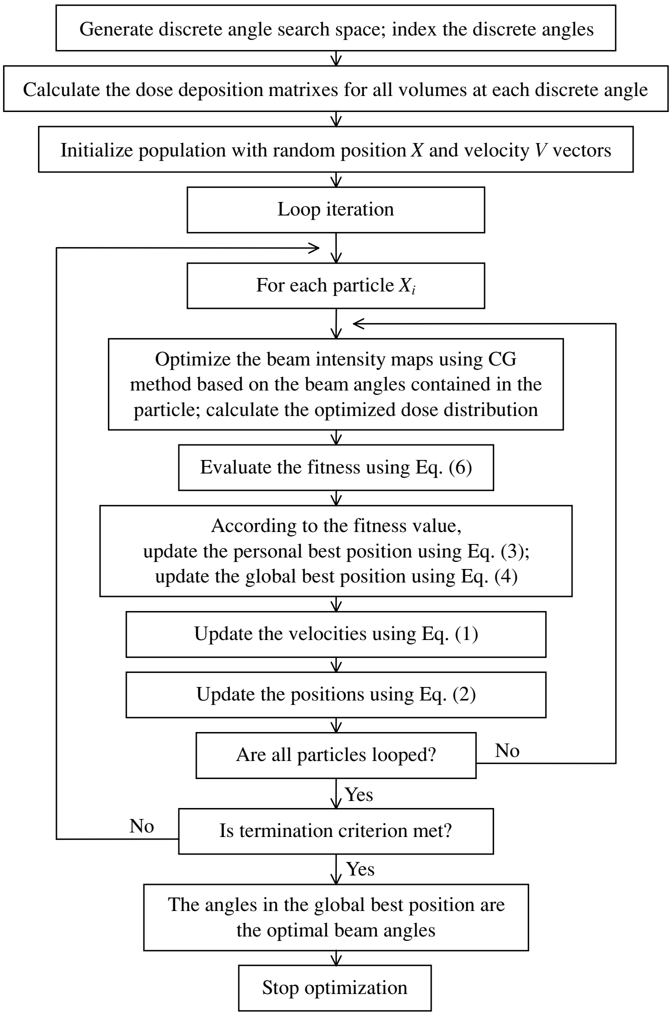

Figure 4. The flow diagram of the presented BASPSO algorithm.

The individuals in the swarm are initialized by randomly setting their positions and velocities using the angle indices. During the iteration, if the value of a new position component calculated using equation (2) is greater than the up limitation of the search space, it will be reset. For example, if the value of a new position component is 75, which is greater than 71, the up limitation of [0, 71] when the whole 360° gantry rotational angle is discretized in increments of 5°, the value 75 will be replaced by 3, the remainder of 75 divided by 72. It should be noted that, at any time of the optimization process, no two components in an individual are allowed to have the same values, which means that beams with the same angles are not allowed in a treatment plan. For an individual, if a new component value calculated using equation (2) already exists in it, a random small integer will be added to this value till no collision exists.

Figure 4 illustrates the flow diagram of the presented BASPSO algorithm. After the initialization of the swarm, the optimization will be implemented iteratively. At each iteration (generation), the updating of the velocity and position is applied to all particles. For each particle with a new position, an optimization of beam intensity maps is carried out using the conjugate gradient (CG) method with the beam angles contained in the particle. The corresponding optimized dose distributions are calculated for the evaluation of the fitness of the particle.

The optimization iteration will be terminated if the predefined maximum iteration number is reached, or the best particle position of the whole swarm cannot be further improved after a sufficiently large number of successive iterations. The angles contained in the best individual $\vec{p}_g = (p_{g1}, p_{g2}, \ldots, p_{gD})$ will be regarded as the optimal set of beam angles.

2.3. Parameter selection in PSO

The inertial weight ω is employed to control the impact of the previous history of velocities on the current one. The role of ω is considered critical for the convergence behaviour of PSO (Parsopoulos and Vrahatis 2002). The parameter ω has characteristics that are reminiscent of the temperature parameter in simulated annealing and is used to regulate the tradeoff between the global exploration and local exploitation abilities of the swarm. A large inertia weight facilitates global exploration; while a small one tends to facilitate local exploitation, i.e., fine-tuning the current search area. The empirical investigations conducted by Shi and Eberhart (1998a, 1998b) indicated that the PSO with the inertia weight in the range [0.9, 1.2] on average would have a better performance. Furthermore, by linearly decreasing the inertia weight from a relatively large value to a small value through the course of a PSO run, PSO tends to have more global search ability at the beginning of the run while having more local search ability near the end of the run.

In this study, the inertial weight ω is selected to be a monotonically decreasing function of the iteration k. Given a user-specified maximum weight $\omega_{\max}$ and a minimum weight $\omega_{\min}$, the inertial weight ω is updated as follows:

$$\omega^{k+1} = \omega^k + \frac{(\omega_{\min} - \omega_{\max})}{K}. \tag{5}$$

Here K is the predefined maximum number of iterations. It has been demonstrated that 0.9 for $\omega_{\max}$ and 0.4 for $\omega_{\min}$ can greatly improve the performance of PSO (Shi and Eberhart 1998a, 1998b).

c_1 and c_2 are also called *cognitive* and *social* parameters, respectively, and they are used to bias the particle's search towards its own best previous position and towards the best experience of the swarm. In the psychological metaphor, the cognitive term represents the tendency of organisms to repeat past behaviours that have proven successful or have been reinforced by their environment, whereas the social term represents the tendency to emulate the successes of others, which is fundamental to human sociality (Cedeno and Agrafiotis 2003). These two parameters are not critical for the convergence of PSO. However, proper fine-tuning may result in faster convergence and alleviation of local minima. An extended study of the acceleration parameter in the first version of PSO is given by Kennedy (1998), in which $c_1 = c_2 = 2$ were proposed as default values. Recent work reports that it might be even better to choose a larger parameter c_1 than parameter c_2, but with $c_1 + c_2 \leqslant 4$ (Carlisle and Dozier 2001).

Since there was no actual mechanism for controlling the velocity of a particle, it is common in particle swarm optimization to define a maximum velocity $V_{\max}$ for each modulus of the velocity vector. Usually, the upper limit of the each modulus of the position vector is used as the maximum velocity. This limitation enhances the local exploration of the problem space and it realistically simulates the incremental changes of human learning. It was found that PSO works well if $V_{\max}$ is set to the value of the dynamic range of each variable (Eberhart and Shi 2001). As for the beam angle optimization studied in this paper, for example, the value of the dynamic range of the angle index is 71 when the whole 360° gantry rotational angle is discretized in increments of 5°. Then V_{max} will be set to 71.

Work done by Eberhart and Shi (2000) indicated that the effect of the population size on the performance of the PSO method is of minimum significance. The typical range for the number of particles N is 20 to 40. Actually, for most of the problems $N = 10$ is large enough to get good results. For some difficult or special problems, one can try 100 or 200 particles as well (Shi and Eberhart 1998a, 1998b, Eberhart and Shi 2000, Clerc and Kennedy 2002, Hu *et al* 2004).

2.4. The fitness function

The quality of each individual (i.e., a plan) is evaluated by a fitness value, and the purpose of the optimization is to find the individual with the maximum fitness. The fitness value of an individual is calculated by

$$\mathrm{Fitness}(\vec{x}_i) = F_{\mathrm{max}} - F_{\mathrm{obj}}(\vec{x}_i), \qquad \vec{x}_i = (x_{i1}, x_{i2}, \ldots, x_{iD}). \tag{6}$$

Here $\vec{x}_i$ is an individual containing angle candidates, and here the swarm size D is the beam number in a treatment plan. $F_{\mathrm{obj}}(\vec{x}_i)$ is the value of the objective function of individual $\vec{x}_i$, and F_{max} is an approximate estimation of the maximum value of the objective function (Li *et al* 2004). The beams with angles in $\vec{x}_i$ are divided into pencil beamlets (also called rays), and all the rays are expressed as a vector $\vec{z}$. The purpose of the optimization is to find a treatment plan which minimizes the difference of the dose distributions between the prescribed and the calculated dose under the guidance of an objective function. The basic form of the objective function used in this paper can be written as

$$F_{\mathrm{obj}}(\vec{z}) = \sum_{j=1}^{NT} \delta \cdot w_j \cdot (d_j(\vec{z}) - p_j)^2 \tag{7}$$

$$d_j(\vec{z}) = \sum_{m=1}^{N_{\mathrm{ray}}} a_{jm} \cdot z_m. \tag{8}$$

Here NT is the number of sample points in the volume, $\delta = 1$ when the point dose in the volume breaks the constraints, else $\delta = 0$, w_j is the weight of the jth point, d_j is the calculated dose of the jth point in the volume, p_j is the prescribed dose of the jth point in the volume, N_{ray} is the total number of rays, a_{jm} is the dose deposited on the jth point from a unit weight of the mth ray, and z_m is the intensity of the mth ray.

2.5. Beam intensity map optimization using CG

In order to evaluate the performance of each individual according to their corresponding optimized dose distributions, an optimization of the intensity map is applied to each new individual using the beam angles contained in it. The optimization is efficiently implemented using a CG method (Spirou and Chui 1998, Zhang *et al* 2004, Li *et al* 2004), which can be written as

$$\vec{z}^{(k+1)}(\lambda) = \vec{z}^{(k)} + \lambda \cdot \vec{h}^{(k+1)} \tag{9}$$

$$\vec{h}^{(k+1)} = -\nabla F_{\mathrm{obj}}(\vec{z}^{(k+1)}) + \frac{[\nabla F_{\mathrm{obj}}(\vec{z}^{(k+1)}) - \nabla F_{\mathrm{obj}}(\vec{z}^{(k)})] \cdot \nabla F_{\mathrm{obj}}(\vec{z}^{(k+1)})}{\nabla F_{\mathrm{obj}}(\vec{z}^{(k)}) \cdot \nabla F_{\mathrm{obj}}(\vec{z}^{(k)})} \cdot \vec{h}_k. \tag{10}$$

Here k means the kth iteration, $\vec{z}^{(k)}$ is the vector of beam intensity maps of the kth iteration, $F_{\mathrm{obj}}(\vec{z}^{(k)})$ is the objective function value calculated using equations (7) and (8), $\nabla F_{\mathrm{obj}}(\vec{z}^{(k)})$ is the gradient of the objective function at point $\vec{z}^{(k)}$, and λ is the step size of the kth iteration, which can be easily calculated using a one-dimensional linear optimization, such as the golden section method.

Equation (10) is initialized by setting $\vec{h}^{(1)} = -\nabla F_{\text{obj}}(\vec{z}^{(1)})$, and the optimization of intensity maps is terminated when $(F_{\text{obj}}(\vec{z}^{(k)}) - F_{\text{obj}}(\vec{z}^{(k-1)}))/F_{\text{obj}}(\vec{z}^{(k)})$ is less than a sufficiently small value (0.001 is experientially used in this study) after the kth iteration, or the maximum iteration number is reached. The final objective function value reflecting the dose distribution difference between the optimized and the prescribed is saved to calculate the individual fitness using equation (6).

Mathematically, rays (i.e., beamlets) with negative weights will appear during the optimization of the intensity maps, which is not physically realizable. To deal with this problem, the non-negativity of ray weights is treated as a hard constraint that may not be violated under any circumstances in our study (Li *et al* 2004).

As for the problem of being potentially trapped in local minima by using a local search method such as CG, several investigators have studied it in detail (Deasy 1997, Borgers and Quinto 1999, Wu and Mohan 2002, Alber *et al* 2002, Llacer *et al* 2003, Jeraj *et al* 2003, Zhang *et al* 2004). One of the main conclusions is that dose–volume constraints make the feasible space possibly non-convex and therefore possibly lead to multiple minima. Fortunately, it has been demonstrated that even when multiple minima do exist for some specially designed cases, those minima are very close to each other in cost and the resulting treatment plans are practically identical, with no clear evidence having been found that multiple minima present impediments to finding a suitable IMRT solution (Wu and Mohan 2002, Llacer *et al* 2003, Zhang *et al* 2004). Even though these limited studies cannot be considered conclusive, the results do increase our confidence in the validity of gradient methods in IMRT optimization.

2.6. Dose calculation

The dose calculation in BASPSO is based on a pencil-beam-based three-dimensional full scatter convolution (FSC) algorithm, in which both the scatter effects and tissue heterogeneity correction are involved with high precision. Instead of simply omitting the scatter doses for the computation efficiency during the optimization process, the BASPSO technique takes into account the majority of those scatter effects. To speed up the calculation under the condition of not markedly sacrificing the accuracy, except the commonly used tricks such as pre-calculation of the dose deposition for all volumes at each discrete angle, some special measures are adopted, such as the indexing of the dose deposition matrixes. The details can be found in our previously published paper (Li *et al* 2004).

3. Results

In order to verify the presented BASPSO technique, first a simulated phantom case with obvious optimal beam angles is carefully designed to study if BASPSO can find the optimal beam angles within a reasonable computation time. Then a clinical prostate tumour case is used to demonstrate the advantage of the presented PSO-based optimization technique again. The results of the PSO-based optimization are compared with those of a GA-based optimization algorithm developed in our previous study (Li *et al* 2004). All the tests are run on a personal computer (Intel Pentium 4, 3.0 GHz, 512 MB RAM) with a Microsoft Windows XP Professional operating system.

The following optimization parameters are selected for all test cases in this section. The search space is sampled with a discrete angle step of 5°, which results in 72 discrete angles in

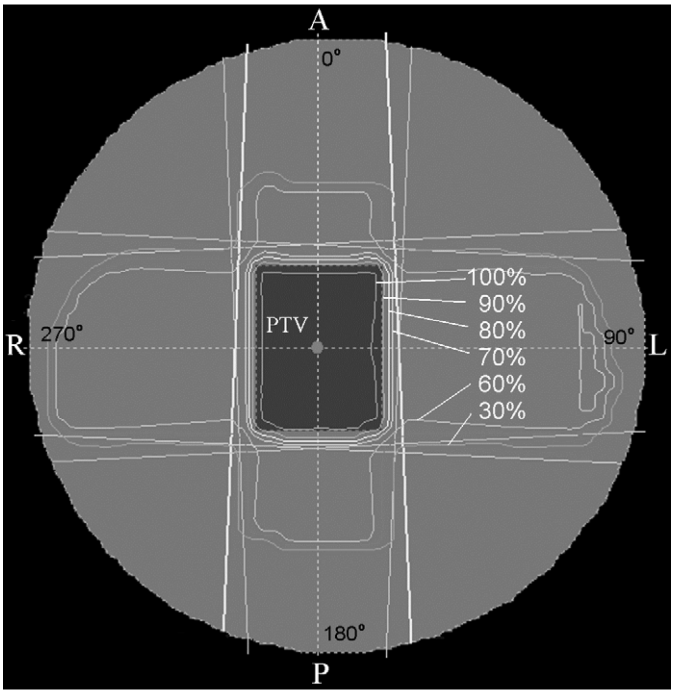

Figure 5. The simulated phantom case including a rectangle-shaped PTV with known optimal beam angles and the dose distribution of the plan with optimized angles. For a four-beam plan, the obvious optimal coplanar beam angles are 0°, 90°, 180° and 270° (the white dotted straight lines). The optimized beam angles are the same as the known optimal ones.

the candidate pool among the whole 360° coplanar gantry angle. Based on the principles of the PSO parameter selection described in section 2.3, we use a swarm of 30 particles, acceleration constants $c_1 = 2.5$, $c_2 = 1.5$, and the maximal velocity $V_{\max} = 71$. These parameter values are empirically determined based on adequate test runs and conform to most settings in our existing applications, though a better optimization performance could be potentially achieved by fine-tuning of these parameters. For a relatively fair comparison, a population size of 30 is used for the GA-based optimization, the same as the swarm size of PSO. 0.95 and 0.01 are used for the crossover and mutation probability of GA, respectively. Also, both of the PSO and GA algorithms start with a random initialization and adopt the same termination criteria, i.e., the optimization is terminated if the maximum iteration (generation) number of 200 is reached, or the global best particle (individual) is not updated in 10 successive generations.

3.1. A simulated test case

Figure 5 shows a simulated case with a rectangle-shaped planning tumour volume (PTV) in a cylinder phantom. It is obvious that 0°, 90°, 180° and 270° are the known optimal beam angles if four coplanar beams are used for the treatment plan (Soderstrom and Brahme 1992, Li *et al* 2004), shown as the white dotted straight lines in figure 5. The size of the PTV is 3.0 cm × 4.0 cm × 2.0 cm, and the size of the voxel volume is set to 0.3 cm × 0.3 cm × 0.5 cm for the PTV. The size of the pencil beam is 0.5 cm × 0.5 cm at the isocentre plane.

For an elaborate analysis, the optimization task runs 10 times. It should be noted that, because the known optimal angles of this simulated case are equispaced, here for all of

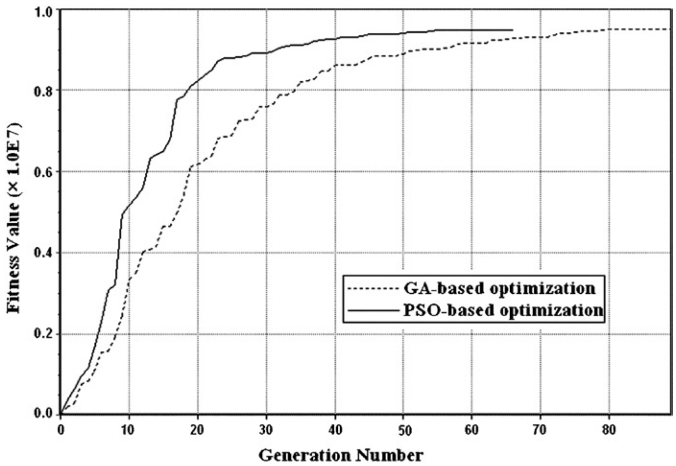

Figure 6. The fitness value versus generation number curves of the PSO- and GA-based optimizations for the simulated case. Note that the fitness values are those of the current best individual (plan).

Table 1. Performance comparisons of the PSO- and GA-based optimizations for the simulated case.

Optimization algorithm	Run times	Minimum computation time	Maximum computation time	Mean computation time	Success times	Success rate[a]
GA based	10	3 min 50 s	5 min 20 s	4 min 35 s	10	1.0
PSO based	10	3 min 15 s	4 min 30 s	3 min 40 s	10	1.0

[a] The fraction of the number of optimization runs in which the optimal result was found.

the 10 runs, particles with equispaced angles are not allowed to appear during the random initialization, in order to fully test the efficiency of the proposed BASPSO algorithm. The same policies are also applied to the GA-based optimization.

As expected, the same optimal angles are successfully found for all the runs of PSO- and GA-based optimizations, which are 0°, 90°, 180° and 270°, the same as the known optimal ones. The corresponding dose distributions are shown in figure 5. Table 1 lists the minimum, maximum and the mean computation time of these runs. From the table we can observe that all the runs are finished within a computation time of 6 min, a very short time because of the relatively small volume size of the PTV and simple anatomy structures of the simulated phantom. Meaningfully, though both of the two algorithms find the optimal beam angles successfully, the PSO-based optimizations need relatively less mean computation time than the GA-based algorithm. The improvement in computation efficiency is not noticeable, which may be attributed to the fact that the computation time listed here includes not only the optimization time, which is optimization engine dependent, but also the computation time spent on the pre-calculation of the dose deposition matrixes and the final refined dose calculation when the optimization task is terminated, which is independent of the optimization engine. Regardless of this independent computation time, the improvement of the optimization resulting from the PSO is sure to be noticeable.

Figure 6 shows the fitness value versus generation number curves of one GA-based run and one PSO-based run for this simulated case. Note that these two runs are with the least computation time among those 10 runs for GA and PSO algorithms, respectively, and

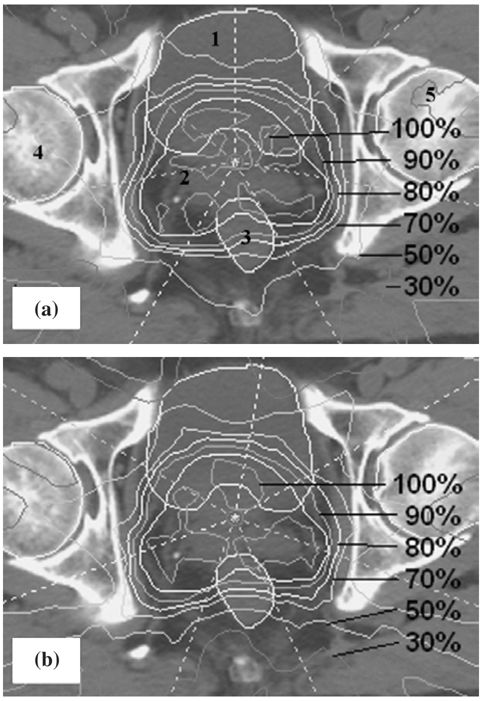

Figure 7. The anatomic structure of the clinical prostate case and the dose distributions of the manual plan (a) and the optimal plan (b). There are in total five volumes involved: 1—bladder, 2—rectum, 3—PTV (prostate), 4—right femur head, 5—left femur head.

the fitness values are those of the current best individual (plan), i.e., the previous global best particle in PSO, and the current best chromosome in GA. This figure shows that the PSO-based optimization converges faster than the GA-based optimization.

3.2. A prostate case

Figure 7 shows a relatively complicated clinical prostate (PTV) case, in which the PTV is surrounded by four OARs, namely, bladder, rectum, the left femur head and the right femur head. The sizes and relative positions of the volumes change substantially from slice to slice, and on most of the slices the contours of the rectum and bladder overlap with the PTV. The CT slice intervals are of 0.5 cm, the size of the voxel volume is set to 0.5 cm $\times$ 0.5 cm $\times$ 0.5 cm for the PTV, bladder, left and right femur heads, and 0.3 cm $\times$ 0.3 cm $\times$ 0.5 cm for the rectum, because of the relatively small volume size. The dose prescription to the PTV is set to 7600 cGy, which is normalized to 100%. Table 2 lists the defined constraints for all the volumes.

Optimization with seven coplanar beams is studied using the PSO- and GA-based beam angle selection algorithms, and the improvement in dose distribution resulting from beam angle optimization is studied by comparing with that of a manual plan with a set of seven coplanar beams with angles of $0°$, $50°$, $100°$, $150°$, $210°$, $260°$ and $310°$, which has become an informal standard in some institutions and oncology centres for the beam set-up for clinical prostate cases. For a fair comparison between the two algorithms, the optimization runs 20 times for both of the algorithms. More runs are performed compared to the simulated case in section 3.1, mainly because this prostate case is more complicated.

Table 2. The constraints for the volumes of the prostate case shown in figure 7.

Volume	Constraints
PTV (prostate)	Prescription: 7600 cGy
Bladder	The volume receiving doses higher than 6000 cGy: <20%
Rectum	The volume receiving doses higher than 6000 cGy: <25%
Left femur head	The volume receiving doses higher than 4500 cGy: <10%
Right femur head	The volume receiving doses higher than 4500 cGy: <10%

Table 3. Performance comparisons of the PSO- and GA-based optimizations for the prostate case.

Optimization algorithm	Run times	Minimum computation time	Maximum computation time	Mean computation time	Success times	Success rate[a]
GA based	20	28 min 43 s	35 min 20 s	32 min 36 s	19	0.95
PSO based	20	23 min 20 s	31 min 38 s	26 min 25 s	20	1.0

[a] The fraction of the number of optimization runs in which the optimal result was found.

Independently, all the 20 runs of the PSO-based algorithm and 19 runs of the GA-based algorithm find the same optimal beam angles: 10°, 60°, 110°, 155°, 200°, 250° and 300°. The remaining one run of the GA-based algorithm gives suboptimal beam angles: 10°, 60°, 110°, 150°, 200°, 250° and 300°. The only difference between the suboptimal and the optimal angles is a 5° angle for the fourth beam. The finally obtained fitness value is $0.905\,3712 \times 10^8$ for this suboptimal result, and $0.905\,3507 \times 10^8$ for the optimal one. This difference is so small that almost no difference could be found on the dose distributions. Rigorously, the reasons for GA producing a suboptimal result need further investigation in order to enhance the robustness of the GA-based algorithm. An imaginable reason for the failure to achieve the optimal result is due to a premature convergence to a local one.

The IMRT dose distributions of the manual plan and the optimal plan found by the optimization algorithm are also shown in figure 7. From these dose distributions we can see that, with a higher conformity, the PTV in the optimal plan is more tightly surrounded by high doses, and the OARs are spared high doses for larger volumes, especially for the rectum, compared to the dose distributions of the manual plan. The dose–volume histogram (DVH) comparisons shown in figure 8 also demonstrate the better results of beam angle optimization. There is no big difference in beam directions between the two plans, whereas the optimization algorithm gives much better dose distributions; this improvement is mainly because of the complicated anatomy structures and the great influence of the intensity-modulated beams on the dose distribution.

Statistics on the optimizations of the two algorithms are listed in table 3. From the table we can find that most of the runs are finished within 25 to 35 min, which is relatively long but clinically acceptable. The longer optimization time may be mainly explained by the relatively complicated anatomy structures and large volume sizes. Remarkably, the table shows that the PSO-based algorithm is more efficient and robust than the GA-based one. Compared to the GA-based algorithm, the PSO-based algorithm reduces the mean computation time from 32 min 36 s to 26 min 25 s, about 19% of the computation time is saved.

Figure 9 shows the fitness value versus generation number curves of one GA-based run and one PSO-based run for this prostate case. Similar to the simulated case in section 3.1, the selected two runs are the fastest runs among the 20 runs in the PSO- and GA-based optimizations. Again, this figure demonstrates the faster convergence of the PSO algorithm

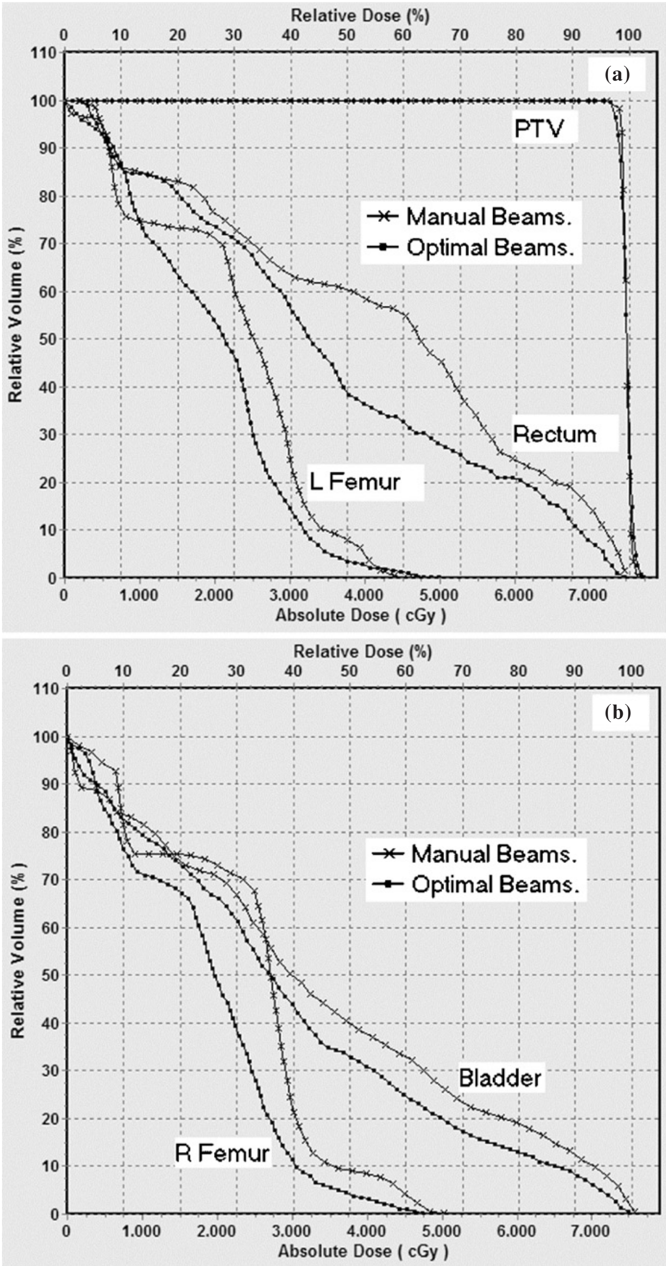

Figure 8. DVH comparisons of the manual plan and the optimal plan for the prostate case.

compared with the GA. The difference in the convergence could be partially explained by the smoothness of the curves. The curve of PSO is smoother, especially the part before the generation number of 60, which means that the global best particle is continuously updated in each generation. In contrast, there are many steps on the curve of GA just after the generation number of 10, which indicates that the best individual in GA is not further improved during those generations.

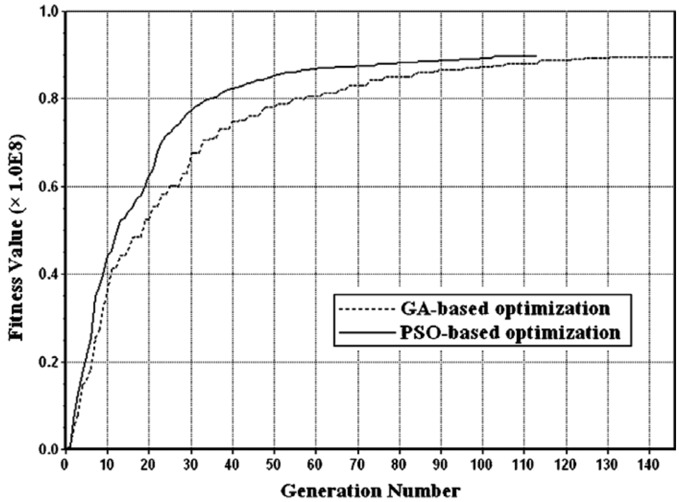

Figure 9. The fitness value versus generation number curves of the PSO- and GA-based optimizations for the prostate case. Note that the fitness values are those of the current best individual (plan).

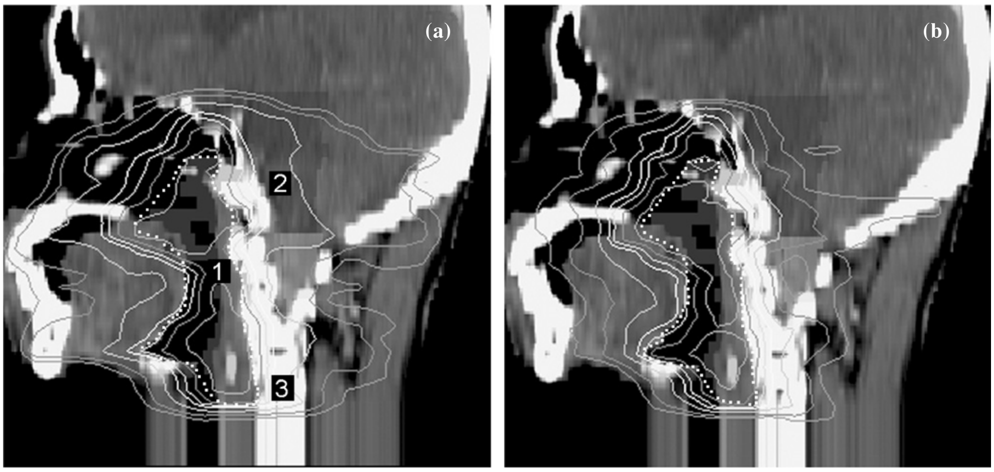

Figure 10. The anatomic structure of the head-and-neck case and the dose distributions of the manual plan (a) and the optimal plan (b) at the sagittal plane. Besides the visible volumes at this sagittal plane, i.e., 1—PTV, 2—brain stem, and 3—spinal cord, both eyes are also involved in the fields. The isodose lines are 100%, 90%, 80%, 70%, 50%, 30% and 10%.

3.3. A head-and-neck case

A head-and-neck tumour case shown in figure 10 is employed to demonstrate the performance of the proposed algorithm again. Besides the PTV, there are four critical structures involved in the fields: the left and right eyes, the brain stem and the spinal cord. From the sagittal plane shown in figure 10 we can see that the sizes and relative positions of the PTV vary substantially from slice to slice. The CT slice intervals are 0.5 cm, the size of the voxel volume is set to 0.5 cm $\times$ 0.5 cm $\times$ 0.5 cm for the PTV, and 0.2 cm $\times$ 0.2 cm $\times$ 0.5 cm for the left and right

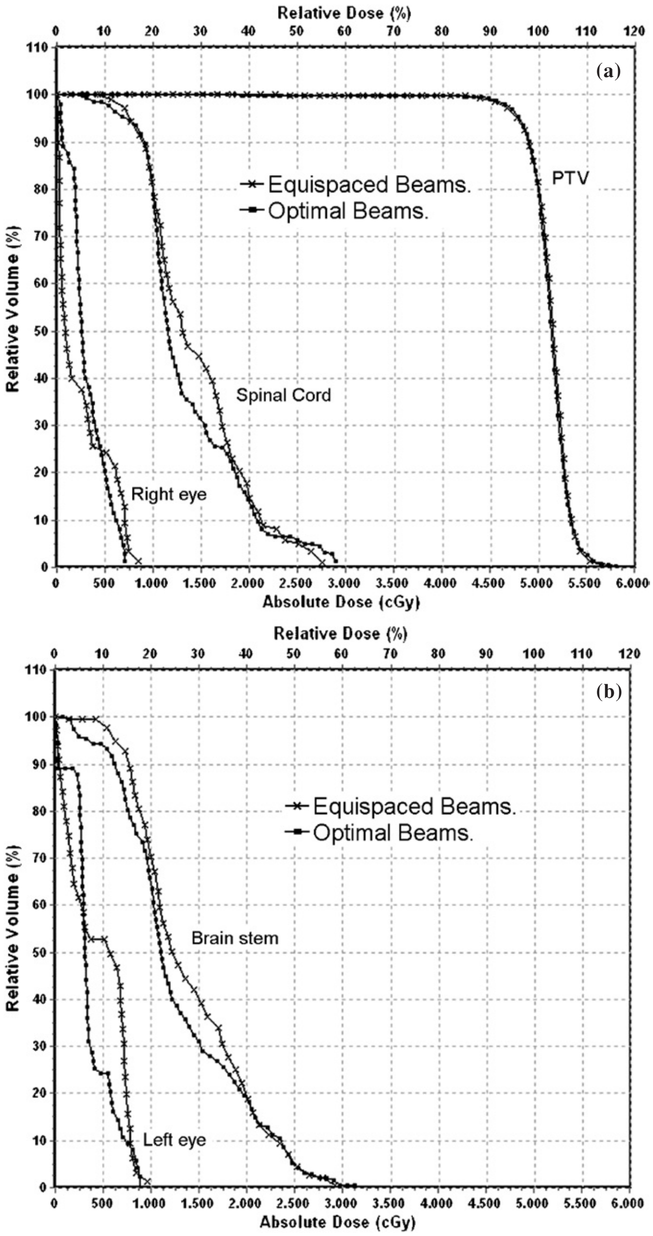

Figure 11. DVH curves of the volumes: (a) PTV, spinal cord and right eye, and (b) brain stem and left eye of the manual plan and the optimal plan for the head-and-neck case.

eyes, the brain stem and the spinal cord. The dose prescription to the PTV is 5000 cGy, which is normalized to 100%. The defined constraints for all the volumes are tabulated in table 4.

Seven coplanar beams are used for both of the PSO- and GA-based algorithms, and the dose distributions are compared with that of a manual plan with a set of seven equispaced coplanar beams with angles of 0°, 51°, 103°, 154°, 206°, 257° and 309°. Similarly, the optimization runs 20 times for both of the algorithms in order to get a fair comparison.

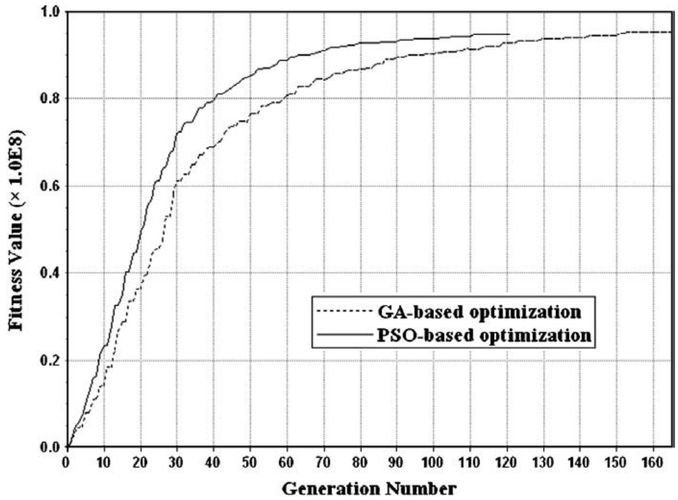

Figure 12. The fitness value versus generation number curves of the PSO- and GA-based optimizations for the head-and-neck case. Note that the fitness values are those of the current best individual (plan).

Table 4. The constraints for the volumes of the head-and-neck case shown in figure 10.

Volume	Constraints
PTV	Prescription: 5000 cGy
Left eye	Max dose: 1000 cGy
Right eye	Max dose: 1000 cGy
Brain stem	Max dose: 3000 cGy; the volume receiving doses higher than 2000 cGy: $<20\%$
Spinal cord	Max dose: 3000 cGy; the volume receiving doses higher than 2000 cGy: $<20\%$

Excitingly, all the runs of the two algorithms find the same optimal beam angles: 20°, 55°, 95°, 130°, 180°, 255° and 320°. The IMRT dose distributions at the sagittal plane are shown in figure 10 for the plans with manually selected and computer-optimized beams. From these dose distributions we can see that the PTV in the optimal plan is more tightly surrounded by high doses with a higher conformity, and the brain stem and the spinal cord being visible at this sagittal plane are spared high doses for larger volumes, compared to the dose distributions of the manual plan. The improvement in dose distribution resulting from the beam angle optimization is demonstrated again by the DVH comparisons shown in figure 11, from which we can clearly observe that, besides the brain stem and spinal cord, both eyes receive high doses with less volumes without sacrificing the dose distribution of the PTV.

Table 5 lists statistics on the optimizations of the two algorithms. From the table we can see that all the runs are finished within 27 min, a clinically acceptable computation time. Again, the table shows that the PSO-based algorithm is more efficient than the GA-based one. The mean computation time is reduced from 24 min 39 s to 19 min 12 s, about 22% of the computation time is saved by the PSO-based algorithm compared to the GA-based algorithm.

The corresponding fitness value versus generation number curves are illustrated in figure 12 for one GA-based run and one PSO-based run for this head-and-neck case. Similarly, the selected two runs are the fastest runs among the 20 runs in the PSO- and GA-based

Table 5. Performance comparisons of the PSO- and GA-based optimizations for the head-and-neck case.

Optimization algorithm	Run times	Minimum computation time	Maximum computation time	Mean computation time	Success times	Success rate[a]
GA based	20	23 min 16 s	26 min 55 s	24 min 39 s	20	1.0
PSO based	20	18 min 41 s	20 min 8 s	19 min 12 s	20	1.0

[a] The fraction of the number of optimization runs in which the optimal result was found.

optimizations. Again, this figure shows that the convergence of the PSO algorithm is faster than GA in the context of the beam angle optimization problem.

4. Discussion and conclusions

In this study, a particle swarm optimization (PSO) based solution is proposed for the problem of beam angle optimization in IMRT planning. The so-called BASPSO technique combines the PSO and the CG algorithms together. Summarily, the main idea of BASPSO is to interpret each particle as a solution to the beam angle selection problem and to let these particles explore the search space. The optimization results of a simulated case and two clinical tumour cases (a prostate case and a head-and-neck case) demonstrate that the presented PSO-based algorithm is valid and efficient and can provide the optimal beam angles within a clinically acceptable computation time. The results of the two clinical cases show that the optimal plan can noticeably improve the dose distributions compared to the manual plan, even the differences are not too large between the optimal and manual angles (e.g., the prostate case).

Furthermore, the preliminary comparisons indicate that, as a whole, the PSO-based algorithm seems to outperform, or at least compete with, the GA-based algorithm in the context of the computation time and the robustness. However, it should be pointed out that the computation time is not absolutely comparable, because different evolution mechanisms are adopted in the PSO and the GA algorithms, and the performances of these two algorithms are dominated by different parameters with different physical or mathematical meanings. Despite all this, the results show that PSO is attractive and competitive with GA. We now consider a thorough comparison of PSO with GA before concluding that one algorithm absolutely outperforms another. Specially, the robustness of PSO to control parameters has been regarded as a major strong point over GA (Kennedy and Spears 1998). This study has not concentrated on the robustness property of PSO; though it has been preliminarily revealed during our experimental tests to determine the PSO parameters. The robustness of PSO over GA is an important focus to be systemically investigated in the coming work.

The optimization of beam angles for IMRT planning is important but challenging work because of the extensive computation (Webb 2000). As mentioned earlier, the common strategy for beam angle optimization in IMRT is to treat the beam angle selection and the beam intensity map optimization as two iteratively implemented separate steps. At present, the algorithms for the beam intensity map optimization are relatively mature, and gradient-based algorithms are commonly used. In fact, the computation burden should be attributed to the necessary global search methods, which is one of the main working directions to improve the optimization efficiency. This paper has concentrated on the latter issue, and the beam intensity maps are optimized using the available CG algorithm. Of course, it would be interesting to attempt to simultaneously optimize both the beam angles and the intensity maps using PSO, which would be investigated in our future works.

To our knowledge, the presented study is the first attempt at proposing a PSO algorithm for the problem of radiation treatment planning optimization. Though the results of our pilot study are appealing, it is worth mentioning that the current study is very preliminary for the PSO-based beam angle optimization applied in the IMRT planning. There is still much work, requiring further insight, to be carried out for the successful and efficient application in clinical practice, such as the more suitable mapping of the radiotherapy optimization problem to PSO search space, robustly adaptive selection of PSO parameters, and so on. In addition, further studies involving a larger number of prostate, head-and-neck and lung cancer cases with more complex structures are required to confirm the general applicability of the proposed approach. In fact, there is much further work needed before the proposed PSO-based optimization algorithm becomes an efficient routine tool in IMRT practice.

The PSO algorithm is based on a simplified model of the swarm theory, in which the birds or particles use their own experience together with the swarm experience in order to find their food or nest. From the socio-cognitive viewpoint, this means that the mind and, as a consequence, the intelligence, are social attributes. A feature of the PSO paradigm that distinguishes it from other evolutionary computation paradigms is its reliance on the individual's memory and the group's finding. The characteristic of memory makes the knowledge of good solutions to be retained by all particles. Whereas in GA, the valuable information contained in those preferable chromosomes would potentially be destroyed or lost after the crossover operation (Gen and Cheng 2000). Furthermore, the constructive cooperation between all particles in a generation provides more chances for particles in the swarm to share preferable information between them compared to GA, in which the information sharing only happens between two parent individuals in a generation. We think that these features of PSO would be the partial explanation for its better convergence compared to the GA.

As an emerging population-based heuristic optimization algorithm, PSO has gained increasing popularity in recent years due to its interesting philosophy and the optimization effectiveness. So far PSO appears to be a very useful technique for solving hard global optimization (GO) problems, and a good alternative in the cases where other techniques fail. The development of PSO is still ongoing, and there are many unsolved issues in PSO such as the thorough convergence analysis and the mathematical validation of particle swarm theory. Also, further research is required to fully comprehend the dynamics and the potential limits of this technique (Kennedy 1998, Parsopoulos and Vrahatis 2002, Hu *et al* 2004, Kennedy 2004).

An interesting point that should be mentioned is that the standard PSO begins by initializing a random swarm. Unless there is a prior knowledge about the parameter space, the initial particles are typically distributed uniformly about the space to facilitate a global search. Some studies have shown that a specially designed initialization helps the algorithm explore the search space more efficiently (Hu *et al* 2004). It is expected that the performance of PSO for beam angle optimization could be improved by a problem-dependent initialization with the guidance of some prior knowledge such as the computer-calculated beam scores (Gokhale *et al* 1994, Pugachev and Xing 2002) and expert knowledge (Xing *et al* 1999, Li *et al* 2004).

It would also be interesting to study the feasibility and effectiveness of these hybrid algorithms by combining PSO and GA or SA. The apparent mutual exclusivity philosophies between PSO and other heuristic algorithms (such as the social context of PSO versus the biological context in GA) leaves open the possibility of integrating the two techniques. Now an important research trend on PSO is to merge or combine PSO with the other evolutionary computation techniques (Eberhart and Shi 2004, Hu *et al* 2004). Especially, the hybridization of PSO with GA has been increasingly studied because both of them are population based. There are indications that in several instances the hybridized approach outperforms either CA or PSO working individually (Boeringer and Werner 2004, Coello Coello *et al* 2004). It is expected that better performances could be achieved for the beam angle optimization problem by fine-tuning the parameters of PSO itself and by hybridizing PSO with other EAs, which is one aspect that we would like to explore in the future.

Acknowledgments

The work reported here is supported by a grant from the 973 Project of China (grant no 2003CB716100), a grant from NSFC of China (grant no 90208003), a grant from the Doctor Training Fund of the Ministry of Education (MOE) of China, and a grant from TRAPOYT of China. The work is also partially supported by Topslane Inc. The authors would like to thank Wenyan Chen and Jiancheng Zheng of Topslane for their helpful discussions and assistance. The authors would also like to thank the anonymous referees for their insightful comments.

References

Abido M A 2002 Optimal power flow using particle swarm optimization *Electr. Power Energy Syst.* **26** 563–71
Alber M *et al* 2002 On the degeneracy of the IMRT optimization problem *Med. Phys.* **29** 2584–9
Boeringer D W and Werner D H 2004 Particle swarm optimization versus genetic algorithms for phased array synthesis *IEEE Trans. Antennas Propag.* **52** 771–9
Borgers C and Quinto E T 1999 On the non-uniqueness of optimal radiation treatment plans *Inverse Problems* **15** 1115–38
Bortfeld T and Schlegel W 1993 Optimization of beam orientation in radiation therapy: some theoretical considerations *Phys. Med. Biol.* **38** 291–304
Carlisle A and Dozier G 2001 An off-the-shelf PSO *Proc. the Particle Swarm Optimization Workshop* pp 1–6
Cedeno W and Agrafiotis D K 2003 Using particle swarms for the development of QSAR models based on K-nearest neighbor and kernel regression *J. Comput. Aided Mol. Des.* **17** 255–63
Chang H, Ratnaweera A, Halgamuge S K and Watson H C 2004 Particle swarm optimization for protein motif discovery *Genet. Program. Evol. Mach.* **5** 203–14
Chvetsova A V, Calvetti D, Sohn J W and Kinsella T J 2005 Regularization of inverse planning for intensity-modulated radiotherapy *Med. Phys.* **32** 501–14
Clerc M and Kennedy J 2002 The particle swarm—explosion, stability, and convergence in a multidimensional complex space *IEEE Trans. Evol. Comput.* **6** 58–73
Coello Coello C A, Pulido G T and Lechuga M S 2004 Handling multiple objectives with particle swarm optimization *IEEE Trans. Evol. Comput.* **8** 256–79
Das S *et al* 2003 Beam orientation selection for intensity-modulated radiation therapy based on target equivalent uniform dose maximization *Int. J. Radiat. Oncol. Biol. Phys.* **55** 215–24
Deasy J O 1997 Multiple local minima in radiotherapy optimization problems with dose–volume constraints *Med. Phys.* **24** 1157–61
De Jong K, Potter M and Spears W 1997 Using problem generators to explore the effects of epistasis *Proc. the 7th Int. Conf. on Genetic Algorithms* pp 338–45
Djajaputra D, Wu Q, Wu Y and Mohan R 2003 Algorithm and performance of a clinical IMRT beam-angle optimization system *Phys. Med. Biol.* **48** 3191–212
Eberhart R C and Kennedy J 1995 A new optimizer using particle swarm theory *Proc. 6th Int. Symp. on Micro Machine and Human Science (Nagoya, Japan)* (Piscataway, NJ: IEEE) pp 39–43
Eberhart R C and Shi Y 1998a Comparison between genetic algorithms and particle swarm optimization *Proc. 7th Ann. Conf. on Evolutionary Programming* (Berlin: Springer) pp 611–15
Eberhart R C and Shi Y 1998b Evolving artificial neural networks *Proc. Int. Congress on Neural Networks and Brain (Beijing, China)* pp PL5–PL13
Eberhart R C and Shi Y 2000 Comparing inertia weights and constriction factors in particle swarm optimization *Proc. Congress on Evolutionary Computation* vol 1 (Piscataway, NJ: IEEE) pp 84–8
Eberhart R C and Shi Y 2001 Particle swarm optimization: developments, applications and resources *Proc. Congress on Evolutionary Computation (Hawaii)* pp 81–6
Eberhart R C and Shi Y 2004 Guest Editorial special issue on particle swarm optimization *IEEE Trans. Evol. Comput.* **8** 201–3

Gaede S, Wong E and Rasmussen H 2004 An algorithm for systematic selection of beam directions for IMRT *Med. Phys.* **31** 376–88

Gao Y and Xie S 2004 A blind source separation algorithm using particle swarm optimization *Proc. the IEEE 6th Circuits and Systems Symposium (Shanghai, China)* pp 297–300

Gen M and Cheng R 2000 *Genetic Algorithms and Engineering Optimization* (New York: Wiley)

Gokhale P, Hussein E M A and Kulkarni N 1994 Determination of beam orientation in radiotherapy planning *Med. Phys.* **21** 393–400

Goldberg D E 1989 *Genetic Algorithms in Search, Optimization, and Machine Learning* (Reading, MA: Addison-Wesley)

Haas O C L, Burnham K J and Mills J A 1998 Optimization of beam orientation in radiotherapy using planar geometry *Phys. Med. Biol.* **43** 2179–93

Hou Q, Wang J, Chen Y and Galvin J M 2003 Beam orientation optimization for IMRT by a hybrid method of genetic algorithm and the simulated dynamics *Med. Phys.* **30** 2360–76

Hu X, Shi Y and Eberhart R C 2004 Recent advances in particle swarm *Proc. IEEE Int. Conf. on Evolutionary Computation* pp 90–7

Jeraj R, Wu C and Mackie T R 2003 Optimizer convergence and local minima errors and their clinical importance *Phys. Med. Biol.* **48** 2809–27

John Cho B C, Mijnheer B J and Bartelink H 2004 Determining optimal two-beam axial orientations for heart sparing in left-sided breast cancer patients *Med. Phys.* **31** 111–21

Kennedy J 1997 The particle swarm: social adaptation of knowledge *Proc. IEEE Congress on Evolutionary Computation* pp 303–8

Kennedy J 1998 The behavior of particles *Proc. 7th Ann. Conf. on Evolutionary Programming* (Berlin: Springer) pp 581–91

Kennedy J 2004 Probability and dynamics in the particle swarm *Proc. IEEE Congress on Evolutionary Computation* vol 1 (Piscataway, NJ: IEEE) pp 340–7

Kennedy J and Eberhart R C 1995 Particle swarm optimization *Proc. IEEE Int. Conf. on Neural Networks* (Piscataway, NJ: IEEE) pp 1942–48

Kennedy J and Eberhart R C 1997 A discrete binary version of the particle swarm algorithm *Proc. Conf. Systems, Man, Cybernetics* (Piscataway, NJ: IEEE) pp 4104–8

Kennedy J and Mendes R 2002 Population structure and particle swarm performance *Proc. IEEE Congress on Evolutionary Computation* pp 1671–6

Kennedy J and Spears W M 1998 Matching algorithms to problems: an experimental test of the particle swarm and some genetic algorithms on the multimodal problem generator *Proc. IEEE Congress on Evolutionary Computation* pp 78–83

Krink T 2002 *Ant Systems and Particle Swarm Optimization* http://www.evalife.dk/ToEC2002/ transparencies/ lecture08b.pdf

Laskari E C, Parsopoulos K E and Vrahatis M N 2002 Particle swarm optimization for integer programming *Proc. IEEE Congress on Evolutionary Computation* pp 1582–7

Li Y, Yao J and Yao D 2004 Automatic beam angle selection in IMRT planning using genetic algorithm *Phys. Med. Biol.* **49** 1915–32

Llacer J, Deasy J O, Bortfeld T R, Solberg T D and Promberger C 2003 Absence of multiple local minima effects in intensity modulated optimization with dose–volume constraints *Phys. Med. Biol.* **48** 183–210

Meedt G, Alber M and Nusslin F 2003 Non-coplanar beam direction optimization for intensity-modulated radiotherapy *Phys. Med. Biol.* **48** 2999–3019

Messerschmidt L and Engelbrecht A P 2004 Learning to play games using a PSO-based competitive learning approach *IEEE Trans. Evol. Comput.* **8** 280–8

Meyer J *et al* 2004 Application of influence diagrams to prostate intensity-modulated radiation therapy plan selection *Phys. Med. Biol.* **49** 1637–53

Parsopoulos K E and Vrahatis M N 2002 Recent approaches to global optimization problems through particle swarm optimization *Nat. Comput.* **1** 235–306

Pugachev A *et al* 2001 Role of beam orientation optimization in intensity-modulated radiation therapy *Int. J. Radiat. Oncol. Biol. Phys.* **50** 551–60

Pugachev A and Xing L 2001 Computer-assisted selection of coplanar beam orientations in intensity-modulated radiation therapy *Phys. Med. Biol.* **46** 2467–76

Pugachev A and Xing L 2002 Incorporating prior knowledge into beam orientation optimization in IMRT *Int. J. Radiat. Oncol. Biol. Phys.* **54** 1565–74

Pugachev A, Boyer A L and Xing L 2000 Beam orientation optimization in intensity-modulated radiation treatment planning *Med. Phys.* **27** 1238–45

Rahmat-Samii Y 2003 Genetic Algorithm (GA) and particle swarm optimization (PSO) in engineering electromagnetics *17th Int. Conf. on Applied Electromagnetics and Communications* pp 1–5

Rasmussen T K and Krink T 2003 Improved hidden Markov model training for multiple sequence alignment by a particle swarm optimization—evolutionary algorithm hybrid *Biosystems* **72** 5–17

Renders J M and Flasse S P 1996 Hybrid methods using genetic algorithms for global optimization *IEEE Trans. Syst. Man Cybern.-B* **26** 243–58

Robinson J and Rahmat-Samii Y 2004 Particle swarm optimization in electromagnetics *IEEE Trans. Antennas Propag.* **52** 397–407

Rowbottom C G, Nutting C M and Webb S 2001 Beam-orientation optimization of intensity-modulated radiotherapy: clinical application to parotid gland tumors *Radiother. Oncol.* **59** 169–77

Rowbottom C G, Webb S and Oldham M 1998 Improvement in prostate radiotherapy from the customization of beam directions *Med. Phys.* **25** 1171–9

Salman A, Ahmad I and Al-Madani S 2002 Particle swarm optimization for task assignment problem *Microprocess. Microsyst.* **26** 363–71

Schreibmann E, Lahanas M, Xing L and Baltas D 2004 Multiobjective evolutionary optimization of the number of beams, their orientations and weights for intensity-modulated radiation therapy *Phys. Med. Biol.* **49** 747–70

Schreibmann E and Xing L 2004 Feasibility study of beam orientation class-solutions for prostate IMRT *Med. Phys.* **31** 2863–70

Shi Y and Eberhart R C 1998a Parameter selection in particle swarm optimization *Evolutionary Programming VII: Proc. EP* (New York: Springer) pp 591–600

Shi Y and Eberhart R C 1998b A modified particle swarm optimizer *Proc. the IEEE Int. Conf. on Evolutionary Computation* (Piscataway, NJ: IEEE) pp 69–73

Soderstrom S and Brahme A 1992 Selection of suitable beam orientation in radiation therapy using entropy and Fourier transform measures *Phys. Med. Biol.* **37** 911–24

Soderstrom S and Brahme A 1995 Which is the most suitable number of photon beam portals in coplanar radiation therapy? *Int. J. Radiat. Oncol. Biol. Phys.* **33** 151–9

Souza W D, Meyer R R and Shi L 2004 Selection of beam orientations in intensity-modulated radiation therapy using single-beam indices and integer programming *Phys. Med. Biol.* **49** 3465–81

Spirou S V and Chui C S 1998 A gradient inverse planning algorithm with dose-volume constraints *Med. Phys.* **25** 321–33

Stein J, Mohan R, Wang X H, Bortfeld T, Wu Q W, Preiser K, Ling C C and Schlegel W 1997 Number and orientations of beams in intensity-modulated radiation treatments *Med. Phys.* **24** 149–60

Ting T, Rao M V C, Loo C K and Ngu S S 2003 Solving unit commitment problem using hybrid particle swarm optimization *J. Heuristics* **9** 507–20

Trelea I C 2003 The particle swarm optimization algorithm: convergence analysis and parameter selection *Inf. Process. Lett.* **85** 317–25

van den Bergh F and Engelbrecht A P 2004 A cooperative approach to particle swarm optimization *IEEE Trans. Evol. Comput.* **8** 225–39

Wachowiak M P *et al* 2004 An approach to multimodal biomedical image registration utilizing particle swarm optimization *IEEE Trans. Evol. Comput.* **8** 289–301

Wang X *et al* 2004 Development of methods for beam angle optimization for IMRT using an accelerated exhaustive search strategy *Int. J. Radiat. Oncol. Biol. Phys.* **60** 1325–37

Webb S 2000 *Intensity-Modulated Radiation Therapy* (Bristol: Institute of Physics Publishing)

Wu Q and Mohan R 2002 Multiple local minima in IMRT optimization based on dose–volume criteria *Med. Phys.* **29** 1514–27

Xing L *et al* 1999 A medical knowledge based system for the selection of beam orientations in intensity modulated radiation therapy (IMRT) *Int. J. Radiat. Oncol. Biol. Phys.* **45** 246–7

Zhang X *et al* 2004 Speed and convergence properties of gradient algorithms for optimization of IMRT *Med. Phys.* **31** 1141–52