Published in final edited form as:

IEEE Trans Biomed Eng. 2017 May ; 64(5): 980–989. doi:10.1109/TBME.2016.2585114.

Radiotherapy Planning Using an Improved Search Strategy in Particle Swarm Optimization

Arezoo Modiri, IEEE [Member], Xuejun Gu, Aaron M. Hagan, and Amit Sawant

Department of Radiation Oncology, the University of Texas Medical Center, Dallas, TX 75235

Abstract

Objective—Evolutionary stochastic global optimization algorithms are widely used in large-scale, non-convex problems. However, enhancing the search efficiency and repeatability of these techniques often requires well-customized approaches. This study investigates one such approach.

Methods—We use particle swarm optimization (PSO) algorithm to solve a 4-dimensional radiation therapy (RT) inverse planning problem, where the key idea is to use respiratory motion as an additional degree of freedom in lung cancer RT. The primary goal is to administer a lethal dose to the tumor target while sparing surrounding healthy tissue. Our iteratively adjusts radiation fluence-weights for all beam apertures across all respiratory phases. We implement three PSO-based approaches: conventionally-used unconstrained, hard-constrained and our proposed virtual search. As proof of concept, five lung cancer patient cases are optimized over ten runs using each PSO approach. For comparison, a dynamically penalized likelihood (DPL) algorithm- a popular RT optimization technique is also implemented and used.

Results—The proposed technique significantly improves the robustness to random initialization while requiring fewer iteration cycles to converge across all cases. DPL manages to find the global optimum in 2 out of 5 RT cases over significantly more iterations.

Conclusion—The proposed virtual search approach boosts the swarm search efficiency and, consequently, improves the optimization convergence rate and robustness for PSO.

Significance—RT planning is a large-scale, non-convex optimization problem, where finding optimal solutions in a clinically practical time is critical. Our proposed approach can potentially improve the optimization efficiency in similar time-sensitive problems.

Index Terms

Non-convex; optimization; PSO; radiotherapy

I. Introduction

Approximately two-thirds of all cancer patients receive some form of radiation therapy (RT) as part of their treatment regimen [1]. RT aims at administering a lethal radiation dose to the tumor while protecting normal tissue and surrounding organs at risk (OARs) [2]. The vast

majority of RT treatments are performed using gantry-mounted medical linear accelerators (linacs) which are capable of administering very high energy x-rays (6 – 18 MeV) from multiple directions or angles such that the beams converge on the tumor target and the dose to OARs and healthy tissue is distributed, thus limiting radiation-induced toxicity (see Fig. 1). In addition, radiation is administered over several daily sessions or "fractions" so as to allow normal cells to recover from radiation damage. To further limit the dose to normal structures, modern linacs are equipped with electronically controlled and programmable multileaf collimators (MLCs), comprising 120 – 180 tungsten plates or "leaves" arranged in two opposing banks. For each gantry angle, the MLC creates apertures and sculpts the beam shape according to the tumor profile as seen in the "beam's eye view" (BEV), a process known as 3D conformal radiotherapy (CRT), which is the focus of this work. An even more sophisticated delivery method, beyond the scope of this work, is intensity modulated radiotherapy (IMRT) where the MLC forms tiny apertures and delivers dose in "segments".

A critical part of administering radiation treatment is to create a "treatment plan" which specifies the number of beams, and the aperture shape, orientation and amount of irradiation for each beam so as to deliver the prescribed dose to the tumor target and maintain the dose to normal tissue and OARs within pre-established limits. As these limits become progressively tighter, the treatment planning becomes more challenging leading to inverse planning optimization techniques [4]. Here, we focus on optimization scenarios where radiation intensity is adjusted for each beam aperture. Depending on the RT technique, the number of adjustable parameters varies between tens to thousands. For instance, although typically 12 gantry angles (beams) are considered in both conventional CRT and IMRT, the number of apertures through which a beam is radiated increases from one in CRT to hundreds in IMRT [5, 6]. Arc-based therapies, as another example, employ large number of beams and deliver radiation on continuous rotations of the radiation source [7].

RT techniques can also be configured spatiotemporally, so-called 4D (3D+time) [8–10]. 4D planning uses time-dependent variations, such as organ motion and deformation, as an additional degree of freedom in optimization [11–14]. As the standard of care for thoracic and abdominal cancer patients, the respiratory cycle is sampled as a set of respiratory phases and computerized tomography (CT) scans are acquired at all the phases (so-called 4DCTs). Anatomical changes over respiratory cycle, identified in 4DCTs, are used in 4D inverse planning [15]. Because one separate plan per respiratory phase is created in 4D RT planning, the dimensionality of the optimization problem is scaled with the number of respiratory phases (typically between 6 and 10) compared to that of 3D.

While tumor irradiation coverage is the primary goal in RT planning, avoiding collateral toxicity in adjacent OARs is critical. Tumor coverage and OAR dose limitations are used to formulate the optimization objective (cost) function, which is generally non-convex. Different techniques have been used in the literature to simplify and solve this problem. For instance, some studies created and tackled a convex approximation of the actual non-convex problem [16]. Some others solved sequential convex optimization steps and applied adjustments according to the error in actual non-convex objective function error at intervals [17–19]. The adjustments were either applied manually through user interaction [18] or adaptively [19]. Other studies simplified the original optimization problem by reducing its

order using principal component analysis [20, 21]. Such simplifications make the optimization problem more tractable but also likely result in sub-optimal solutions.

Here, we employed a stochastic global optimization algorithm to explicitly tackle the non-convex RT planning optimization problem using particle swarm optimization (PSO). PSO searches the solution space using a cohort of agents called particles [22, 23]. These particles use inertia, cognition, social learning and random decision-making to move inside the solution space and to bypass discontinuities. Compared to other global evolutionary algorithms, such as genetic algorithms and neural networks, PSO is highly parallelizable, and relatively simple to implement and control [24]. These features led us choose PSO for 4D RT inverse planning. We also implemented dynamically penalized likelihood (DPL) algorithm, a popular optimization technique in RT planning [25–27], for comparison.

We investigated unconstrained and hard-constrained PSO-based optimization approaches. Further, we proposed and developed a virtual search approach to enhance the algorithm robustness to initialization and improve search efficiency. In the proposed approach, the objective function calculation was influenced using the most important objective term which, in the case of RT inverse planning, was the prescribed tumor coverage. Our proposed method borrowed the idea of projection-based algorithms to iteratively project solutions from infeasible region onto feasible region [28]. The developed virtual search approach was compared to unconstrained and hard-constrained PSO-based approaches as well as DPL technique highlighting robustness to random initialization and final optimum solution.

A preliminary version of this work was reported in an IEEE conference and AAPM abstracts [29–31]. Although the idea of virtual search was proposed in our preliminary reports, the algorithm was different. The previous method was based on virtually increasing the swarm population size, while the new algorithm keeps the swarm size identical and thus reduces computational complexity which is critical in RT planning. Also, in order to verify possible generalizability, here, optimization repeatability and robustness to random initialization are investigated not only for multiple RT cases but also for a benchmark problem (Rastigrin function).

The outline of this paper is as follows: sections II and III introduce the RT optimization problem scope and objective function characteristics, respectively. The three PSO-based approaches are explained in section IV. Section V compares the results of the introduced approaches for the benchmark Rastigrin problem. 5 stereotactic body radiation therapy (SBRT) lung cancer cases are introduced in section VI. The optimization results for the case studies using the PSO and the DPL methods are presented in section VII. For each PSO scenario, 10 runs were executed to validate the improved search efficiency of the proposed method.

II. Problem Scope

In this study, we investigated 4D CRT inverse planning and characterized the optimization problem in terms of non-convexity. In current clinical practice, respiration-induced variations of the tumor target with respect to the beam are typically accounted for by

introducing a motion-encompassing volume in 3D plan. This volume is called internal target volume (ITV) [32]. As shown in Fig. 2a, ITV is the union of all possible time-dependent gross tumor volumes (GTV - the visible extent of the tumor as assessed from pre-treatment imaging). Typically, an additional margin is added to the ITV in order to accommodate day-to-day patient positioning uncertainties, thus creating a planning target volume (PTV). A PTV generated over an ITV results in the irradiation of a significantly larger volume than that of the actual tumor. Thus, the conventional margin expansion strategy causes the irradiation of healthy surrounding tissues and OARs, which can lead to moderate-to-significant radiation toxicity [33]. This phenomenon becomes more important in hypofractionated RT, such as SBRT, where potent radiation doses are given in few fractions. In a 4D plan, however, the beam aperture shapes follow respiratory-phase-specific PTVs [34] as shown in Fig. 2b, and consequently, OAR-sparing is improved by suppressing the target volume expansion.

In an RT plan, the dose contribution of each beam aperture to corresponding volume elements (voxels) within the irradiated anatomical volume is scalable by a weight. That weight is the aperture's monitor unit (MU) weight, also called intensity-weight or fluence-weight. In this study, the optimization task is to adjust aperture weights across all respiratory phases. The goal is to satisfy clinical dose limitations related to healthy organs while first and foremost maintaining prescribed tumor coverage. More comprehensive planning methods, such as beamlet intensity optimization (fluence optimization [11]), are beyond the scope of this work. Also, while extendable to more complex planning techniques (IMRT and VMAT), we limited our present discussion to CRT, where there is only one aperture per beam.

III. Objective Function

Organ-based dose limitations, as well as prescribed tumor coverage are typically used to create the objective function for an RT planning optimization problem. Although organ-based dose limitations are patient specific, guidelines are reported in clinical protocols which are developed and updated by multi-institutional cooperative bodies such as the radiation therapy oncology group (RTOG). These protocols are based on clinical experiences of domain experts, i.e., radiation oncologists from multiple radiation oncology clinics (e.g., RTOG0236 for lung SBRT [35]).

In RT planning, tumor coverage is the primary goal. In most clinical scenarios, tumor coverage prescription is volume-based. For instance, in our institution, it is prescribed that 95% of tumor should receive 54 Gy dose in a 3-fraction SBRT plan. Inverse planning optimization aims at creating a dose distribution across 3D volume of patient body that satisfies the tumor coverage and spares OARs. One standard objective function designed for the purpose of inverse RT planning is the mean squared error (MSE) [4, 19, 36]. Inverse planning algorithm optimizes aperture weights to minimize dose error summed over CT image voxels. Such MSE objective function is formulated as the summation of normalized weighted MSE terms per structure (tumor and OARs) modelled as:

$$F = \sum_{i=1}^{N_{structures}} \frac{1}{N_{voxels_i}} \left(\sum_{f=1}^{N_{objective_i}^{up}} F_{i,f}^{up} + \sum_{f=1}^{N_{objective_i}^{low}} F_{i,f}^{low} + F_i^{max} \right) \quad (1)$$

$$\forall i \in structures, \forall f \in objective_i^{up/low},$$

$$F_{i,f}^{up} = \begin{cases} w_{i,f}^{up} \sum_{j=1}^{N_{voxels_i}} (D^j - C_{i,f}^{up})^2 \cdot H(D^j - C_{i,f}^{up}) & V_{i,f}^{HD} \geq V_{i,f}^{up} \\ 0 & \text{else} \end{cases}$$

$$F_{i,f}^{low} = \begin{cases} w_{i,f}^{low} \sum_{j=1}^{N_{voxels_i}} (D^j - C_{i,f}^{low})^2 \cdot H(D^j - C_{i,f}^{low}) & V_{i,f}^{LD} \leq V_{i,f}^{low} \\ 0 & \text{else} \end{cases}$$

$$F_i^{max} = w_i^{max} \sum_{j=1}^{N_{voxels_i}} (D^j - C_i^{max})^2 \cdot H(D^j - C_i^{max})$$

$$H(y) = \begin{cases} 1 & y \geq 0 \\ 0 & y < 0 \end{cases} \quad (2)$$

$$D = \sum_{k=1}^{N_{apertures}} DIR(\vartheta_k \times \mathbf{d}_k) \quad \vartheta_k \geq 0 \quad (3)$$

All parameters named with the letter N , in (1)–(3), denote the number of what is shown as subscript (e.g., $N_{structures}$ is the number of structures). Superscripts *up*, *low* and *max* in (1)–(2) correspond to volume-based upper dose bound, volume-based lower dose bound, and maximum allowable dose per voxel, respectively. For each structure, multiple upper and lower bounds (shown by summation over f in (1)) and one maximum voxel dose can be introduced. Each objective term in (1) is defined in (2). F^{up} , the upper dose bound, generates a penalty if the dose given to a certain volume within the structure (V^{up}) exceeds a certain value (C^{up}). F^{low} functions similarly for lower bound, considering V^{low} and C^{low} . V^{HD} is the volume within the structure that receives higher doses than C^{up} . Similarly, V^{LD} is the volume within the structure that receives lower doses than C^{low} . F^{max} , the maximum voxel dose objective, generates a penalty if the dose given to any voxel within a structure exceeds C^{max} .

Prioritizing factors (w^m , $m \in \{up, low, max\}$) weight the objective terms and are patient-specific and user-defined. D is the total 3D dose matrix. For radiating aperture k , the dose deposition matrix and weight are shown by d_k and ϑ_k , respectively. Here, d matrices are known, pre-calculated by a commercial dose calculation engine (Eclipse treatment planning system, Palo Alto). ϑ_k 's are the only independent unknown variables of this optimization problem and are adjusted by optimization algorithm. In 4D planning, scaled d matrices go through a deformable image registration (*DIR*) process, so that all phases are mapped on a reference phase for summation. The goal of optimization is to find optimal $\vartheta = [\vartheta_1, \vartheta_2, \dots, \vartheta_{N_{apertures}}]$ that minimizes F in (1).

The MSE function generally typifies a convex function; however, the volume-based terms, F^{up} and F^{low} , cause non-convexity by introducing discontinuities in the objective function. Note that RT optimization problem is degenerate [37]; i.e., optimization algorithm is merely guided by an objective function in which trade-offs can occur between organ-based objective terms. Thus, it is possible to get the same objective value from distinct dose distributions.

As a general rule, the choice of optimization technique and its related parameters depends on the shape of the objective function [38, 39]. To visualize the objective function shape in RT inverse planning, a simplified scenario is illustrated in Fig. 3, where only one aperture weight varies ($N_{apertures} = 1$) and only one phase exists; i.e., *DIR* is not required in (3). Three structures, called OAR_1 , OAR_2 , and PTV , each composed of nine voxels, were considered and three sets of arbitrarily selected priority weights (w 's) were tested. The parameters used in (1)–(2) creating Fig. 3 are shown in TABLE I. The initial voxel doses were chosen randomly in the range of (0, 1) for *OARs* and (0, 2) for *PTV*:

$$d_{OAR_i} = \text{rand}(3, 3), i \in \{1, 2\}$$

$$d_{PTV} = 2 \times \text{rand}(3, 3)$$

$$d_1 \leftarrow \text{concatenate } d_{OAR_1}, d_{OAR_2} \text{ and } d_{PTV}$$

$$D = \vartheta_1 \times d_1$$

The three curves in Fig. 3 exemplify how objective function in (1) varies with aperture weight. When extended to a larger number of apertures, discontinuities in the solution space can distract optimization algorithms. Therefore, gradient-based or simplified optimization techniques are less likely to yield global solutions.

IV. Optimization Technique

PSO is an evolutionary stochastic global optimization technique which imitates a biologically inspired organization of a swarm in its search activity. PSO was first introduced

by Eberhart and Kennedy in 1995 [22, 23], and soon after, used in benchmark optimization problems (e.g., Rosenbrock, Rastigrin, Griewank) and later applied to a variety of optimization-based applications [40]. Evolutionary algorithms with a stochastic nature, like PSO, have the ability to handle discontinuities such as those shown in Fig. 3 [41]. In addition, having a swarm of search agents instead of one evolving solution allows comprehensive exploration of solution space. Furthermore, high parallelizability inherent to PSO makes it easier to manage the computational complexity.

A comparison of global and local topologies of PSO is given by Bratton et al. [42] where a group of benchmark optimization problems are tested. From that study we surmise that our RT problem benefits from global PSO topology.

In global topology, each PSO particle is navigated by the following velocity function.

$$\mathbf{V}^t = \omega \mathbf{V}^{t-1} + c_1 \lambda_1^t \cdot (\mathbf{p}^{t-1} - \boldsymbol{\vartheta}^{t-1}) + c_2 \lambda_2^t \cdot (\mathbf{g}^{t-1} - \boldsymbol{\vartheta}^{t-1})$$

$$\boldsymbol{\vartheta}^t = \boldsymbol{\vartheta}^{t-1} + \mathbf{V}^t \quad (4)$$

where $\mathbf{V}$, $\mathbf{p}$, $\mathbf{g}$ and $\boldsymbol{\vartheta}$ denote the vectors of velocity, personal best position, global best position and current position, respectively. Note that $\boldsymbol{\vartheta}$, which represents the weight vector in (3), is, in fact, the particle position vector for the PSO. The iteration number is shown by a superscript (t). The following triad composes the velocity function: (i) inertia which dictates following the previous search direction; (ii) cognition which encourages moving toward the best individually-experienced solution (personal best) and (iii) social learning which encourages moving toward the best solution found by the entire swarm (global best). The significance of each velocity term is adjustable by ω , c_1 and c_2 . To add stochastic behavior, λ_1 and λ_2 are chosen to be random vectors with elements uniformly distributed in (0, 1). c_1 and c_2 were chosen to be equal to 2 and ω linearly decreased from 0.9 to 0.4 over iterations, as recommended in the literature [42, 43].

A. Swarm Population Size

There is no standard guideline for choosing the swarm population size. As a rule of thumb, the number of particles needs to be large enough to collect sufficient samples of the solution space. A smaller swarm needs larger number of iterations to explore and converge. Also, an excessively small swarm may not even get a chance to scan the entire solution space. A large swarm of particles can do a comprehensive exploration of the solution space but at the cost of an increased computational complexity ($O(n)$ with n being the swarm population size [23]). Bratton et al. demonstrated that a swarm population 20 was able to explore the solution space widely for most 30-dimensional benchmark problems, leading to a global optimum [42]; however, swarm population size depends on the solution space complexity. Note that with no volume-based constraint, the MSE objective function in (1)–(3) is a convex equation. The dose-volume constraints are the causes of discontinuities; hence, the

number of dose-volume constraints should be considered as the indicator of problem complexity and determiner of swarm population size.

B. Initialization

In RT optimization problem, an ideal paradigm for initial PSO particle positioning is to have one particle in the vicinity of every discontinuity in every dimension (see Fig. 3). The evolution trend of the objective function for such paradigm would exhibit relatively sharp migration to convergence after few exploration cycles. Nonetheless, such initial distribution pattern is not easy to configure. Mapping the solution space in an RT problem is challenging, especially when considering case-dependency and user-specificity and also potential large dimensionalities. Therefore, the original random uniform distribution of particles was used in this work.

C. Optimization Approaches

In most clinical scenarios, an RT plan which does not satisfy tumor coverage (i.e., covering 95% of tumor volume with 100% prescribed dose) is deemed unacceptable irrespective to OAR dose distribution. 95% tumor coverage is hence named critical objective in this paper. Two optimization approaches are typically followed for RT planning:

- Approach 1: Unconstrained optimization, where optimization objectives are summed as a single objective term and user-defined priority factors weigh each objective term.
- Approach 2: Hard-constrained optimization, where the critical objective is defined as a hard constraint. All other organ-based dose objectives are defined as explained in Approach 1.

Approach 1 is the most widely used RT optimization strategy [26, 36]. In this approach, it is the user's responsibility to prioritize the objectives and quantify the priority factors. There is no simple way to choose priority factors beyond trial and error. Approach 1 uses different but comparable priority factors for all structures and generally doesn't end up with satisfying the critical objective. To correct the final plan, the dose is normalized (scaled) upon optimization termination. However, this final normalization step is not consistent with the optimization process since it changes the absolute dose values.

Approach 2 seems to be the most straightforward strategy to satisfy the critical objective. However, hard constraints are known to impede convergence in optimization algorithms, specially evolutionary algorithms, by disturbing the continuity in optimization process [44]. To impose a hard constraint, we used penalty weighting of the objective function as introduced for PSO in the literature [45–47]; i.e., in order to remove an unfeasible solution from a particle's memory and preserve feasibility, a very large objective value (e.g., $1e^{26}$, for this study) was assigned when critical objective was violated. Note that the penalty assigned for violating a hard constraint creates a significantly larger objective value than that achievable by priority factor assignment in Approach 1.

We used our knowledge of the critical objective and explored a new avenue based on a simple synergistic combination of the previous two approaches.

- Approach 3: Virtual search optimization, where the critical objective is used to navigate the search agents.

We propose to bring the normalization step into optimization iterations. The normalization is performed prior to calculating the objective function in each iteration cycle and for each particle by introducing D'^t and NF^t as

$$D'^t = D^t \times NF^t$$

$$NF^t = D_{95}^{\text{prescribed}} / D_{95}^t \quad (5)$$

where $D_{95}^{\text{Prescribed}}$ and D_{95}^t denote the prescribed and the current (iteration t) doses received by equal or larger than 95% of tumor volume. NF^t is the normalization factor which scales the absolute dose values in matrix D so that the prescribed tumor coverage is achieved. D'^t represents a surrogate dose distribution related to a virtual image of ϑ^t (ϑ'^t) and is used to create a surrogate objective value for ϑ^t . The virtual images are projected copies of the infeasible particle positions inside feasible solution region.

In a simplified two-variable scenario, Fig. 4 compares actual and surrogate objective function values when DIR is not required (e.g., a 3D RT scenario). Each solid circle in Fig. 4a represents a particle position which is an array of two aperture weights ($\vartheta_{\text{aperture1}}$, $\vartheta_{\text{aperture2}}$). All solid circles bound by a single line represent vectors which are multiples of each other. Only one weight set on each line, shown as black circle, meets the critical objective. The black circles are the virtual images of their line-mate gray circles. Fig. 4b models the actual and surrogate objective values for the six particle positions shown on the solid line in Fig. 4a.

In order to avoid limiting PSO's exploration span, we choose to continue using ϑ^t in velocity calculation and position updating, while surrogate objective values using D'^t were used to quantify the goodness of each particle's position. Note that if the virtual images (ϑ'^t) were to replace the particle positions in velocity calculation, the solution space exploration would get limited to moving the black circles.

As compared to Approach 1, Approach 3 brings the normalization step into iteration cycles to guarantee that the critical objective is always satisfied. Also, as compared to Approach 2, Approach 3 preserves feasibility by creating projections of the solutions that violate the critical objective inside feasible solution space instead of discarding them. The objective function calculation step for the three otherwise identical approaches is demonstrated in Fig. 5.

Approach 3 is applicable to any problem where satisfying a critical objective is directly translatable to repositioning inside solution space. Note that RT optimization problem does not belong to the class of hyperplanes problems because the objective function value is not scaled with scaling the ϑ vector. Approach 3 enables carving the solution space to a smaller

one for the particles to explore because the particles don't see any improvement, in terms of objective value, if they move in the direction of the lines modeled in Fig. 4a.

Since our proposed modification mainly modifies the objective function calculation step, it can be applied to global optimization algorithms other than PSO.

In our preliminary study reported in [29], at each iteration, each infeasible particle position was replaced by a set of scaled positions (on a corresponding line in the simplified model of Fig. 4a), as if the swarm size was multiplied. Such approach was computationally expensive (PSO computation complexity is $O(n)$ with n being the swarm population size [23]). Here, the same core idea was implemented while keeping the swarm size fixed, yet, achieving global minimum.

V. Testing on a Benchmark

To evaluate the performance of our proposed approach in a typical optimization benchmark, we implemented 2D Rastigrin function:

$$f(z_1, z_2) = \sum_{i=1}^2 (z_i^2 - 10 \times \cos(2\pi z_i) + 10) \quad (6)$$

In our test, z_1 and z_2 changed between -5 and 5 . To create a variable space similar to that of RT planning problem shown in Fig. 4, a critical objective was added to Rastigrin optimization process, which confined the feasible region inside the solution space to a discontinuous curve shown in Fig. 6. The feasible region was modeled and formulized by:

$$\begin{cases} z_1^2 + z_2^2 = 1, z_1 > 0 & 0 \leq |\alpha| < \frac{\pi}{6} \\ z_1^2 + z_2^2 = 2.5^2, z_1 < 0 & \frac{\pi}{6} \leq |\alpha| < \frac{\pi}{3} \\ z_1^2 + z_2^2 = 5^2, z_1 > 0 & \frac{\pi}{3} \leq |\alpha| \leq \frac{\pi}{2} \end{cases}$$

$$\alpha = \arctan(z_2/z_1) \quad (7)$$

Fig. 6 is an example of a variable space where each point either belongs to the feasible region or has a unique multiple in the feasible region. We let each of the previously introduced PSO-based approaches solve this optimization problem while we already knew the optimum solution in the feasible region is $f(z_1, z_2) = 1$. The optimization algorithms were run 100 times over 200 iterations using 5 particles. The convergence trend for the three approaches is shown in Fig. 7, where bars show the variation range in every 10-iteration step.

For Approach 1, a critical objective, defined as $\Delta = 10 \times \sum_{i=1}^2 |z_i^{\text{feasible}} - z_i^{\text{infeasible}}|$, was added to $f(z_1, z_2)$. A feasible solution associated with an infeasible $z^{\text{infeasible}} = (z_1^{\text{infeasible}}, z_2^{\text{infeasible}})$ was defined as $\beta \times z^{\text{infeasible}}$ which satisfied (7) with β

being a real scaling factor. Note that in the feasible region, α is equal to zero; thus, keeping the optimum solution at $f(z_1, z_1) = 1$.

For Approach 2, a hard constraint was imposed by penalty weighting the objective function value for any infeasible solution ($f(z_1, z_2) = 50$).

For all three approaches, the feasible region was defined to be within the distance of 0.05 from the discontinuous line shown in Fig. 6 in either dimension. The termination criterion was set to maximum iteration cycles of 200.

Fig. 7 illustrates how a hard constraint makes it difficult for the PSO algorithm to converge. More importantly, Fig. 7 shows that Approach 3 enhances the optimization search efficiency by guaranteeing a faster convergence to the optimum solution. Approach 3 also demonstrates a convergence behavior which is more robust to random initialization as compared to the other two approaches.

Fig. 8 shows the optimization convergence trend of Approach 3 when virtual images of the particles replaced the actual particle positions in velocity function calculation (a) and when they didn't (b). To better highlight this phenomenon, swarm population was increased to 10. In Fig. 8, (a) converged to 1 at iteration 135 versus iteration 75 for (b), considering the convergence criterion of staying unchanged for 5 consecutive iterations.

VI. CASE STUDY

A. Method

We studied the three PSO approaches described in section IV.C for five CRT-SBRT lung cases and compared the results to those of the DPL optimization technique, which has been used for RT planning in the literature [26, 27]. We implemented DPL as explained in equation (5) in [26, 27], considering an exponent of $n_{DPL} = 1$.

Increasing n_{DPL} speeds up the DPL process by increasing the step size; however, it causes diverging oscillations after a certain problem-specific threshold. For each case, we increased n_{DPL} by steps of 0.5, starting from 1, and chose the largest n_{DPL} before diverging oscillations happen. Filtering term was excluded as in [26, 27] for 4D RT planning.

In order to have a predictable global optimum and an understandable comparison of the behavior of different algorithmic approaches:

- i. Only two structures, tumor (PTV) and one OAR (esophagus), were considered.
- ii. One single objective function was defined and used for all cases and all optimization approaches using parameters shown in TABLE II. The prescribed dose of 54 Gy to 95% of tumor volume was considered. A volume range of 94% (V^{up})–96% (V^{low}) was chosen in TABLE II for PTV critical objective violation bounds in order to account for DIR-induced uncertainties.
- iii. The termination criterion was set to maximum iterations of 30 for both PSO and DPL considering a PSO swarm size of 25.

- iv. Each PSO optimization was run 10 times to account for randomness.

The number of beams (apertures) per phase were [9, 10, 10, 11, 9] for [Patient1, Patient2, Patient3, Patient4, Patient5]; thus, the number of variables varied between 90 and 110. Fig. 9 shows the tumor site and beam configuration for Patient1.

An upper and a lower bound of 3.3 and 0 were considered for MU weights (ϕ). Maximum permissible MU weight was arbitrarily chosen as 3.3 to keep values close to a traditional gated RT plan where all required MUs are delivered over 3 phases ($10/3 = 3.3$). On delivery side, a zero-weighted aperture would be held OFF (beam-hold) and an aperture with large MU-weight would take multiple respiratory phases to be delivered. An absorbing wall boundary condition was considered for PSO, meaning that particles that “flew out” the bounds were neither discarded nor reflected but stopped and given new chances to fly back into the feasible space [43].

B. Data Preparation

For each patient, a 4D CT scan comprising ten CT volumes corresponding to ten respiratory phases was collected retrospectively. The target and normal organs were manually contoured on the CT volume corresponding to end exhalation (reference respiratory phase). The contours were propagated from the reference phase to the other phases using Velocity, a commercial image registration software tool (Varian Medical Systems, Palo Alto, CA). For each respiratory phase, a 3D plan was generated in Eclipse treatment planning system using the same beam configurations (gantry and couch angles) as the clinically delivered plan (e.g., Fig. 9). The dose deposition matrices (d 's in (3)), for all apertures in all phases, were exported from Eclipse and fed into the optimization algorithms. All particles in each PSO swarm were initialized randomly except for one, defined as $\phi = [1, 1, \dots, 1]$ (all ones), which represented a 4D plan simply made from uniting aforementioned initial 3D plans created in Eclipse. The all-ones solution was used as starting point for DPL as well. The aperture shapes were made conformal to individual-phase PTVs as illustrated in Fig. 2b. At each iteration cycle, the actual dose at each phase was multiplied by the corresponding ϕ . NiftyReg, an open-source package for deformable image registration, was then used to map the dose from different respiratory phases on the reference phase for summation [48].

C. Processing Platform

The optimization code was implemented in MATLAB R2013b and run using the parallel processing toolbox. The calculations were performed on a dual 8-core Intel® 3.1 GHz Xeon (R) CPU (32 processing cores with hyperthreading). While the parallel processing toolbox in the non-server version of MATLAB does not allow more than 12 task distributions (workers), introducing the following nested loops in functions improved parallelization on CPU.

```

for iteration = 1 : Niterations
    Initialize  $\phi$ ,  $V$ ,  $g$  and  $p$ 
    par for particles = 1 : Nparticles

```

```

–Function: Calculate Dose Matrix
for 1 :  $N_{phases}$ 
  for 1 :  $N_{beams}$ 
    Scale dose matrix  $d$  by its corresponding  $\Phi$ 
  end
  calculate summed doses over beams for each phase
end
par for 1 :  $N_{phases} - 1$ 
  Deform dose matrices
end
Calculate  $D$ 
–Function: Calculate  $F$ 
for  $i = 1 : N_{structures}$ 
  Calculate  $F^{up}$ ,  $F^{low}$  and  $F^{max}$ 
end
Calculate  $F$ 
end
Update  $g$  and  $p$ 
Calculate  $V$  and update  $\Phi$ 
end

```

VII. Case Study Results

Fig. 10 shows the dose volume histogram (DVH) curves for all three PSO approaches (solid lines) and the known initial solution (dashed lines). It is seen that the deviation range for the hard-constrained results are significantly larger compared to those of the unconstrained and virtual search approaches (similar to the observation for the benchmark problem in V). A smaller deviation range, despite random initialization, proves higher robustness and repeatability of an optimization approach.

Fig. 11 shows the averaged DVH curves over 10 runs for the PSO approaches compared to the DPL and the initial known solution. It is observed that, overall, the PSO approaches outperformed the DPL. Among the PSO approaches, the proposed Approach 3 was performing the best and the hard constrained Approach 2, the worst.

The objective function curves for the three PSO approaches over ten runs are shown in Fig. 12, where bars show the deviation ranges. In Approach 2, a very large objective value ($1e^{26}$) was assigned to the solutions violating the critical objective. Such high objective values were excluded from averaging if the entire swarm was found to be violating the critical objective in one iteration. For the first three cases, an objective value of zero was achievable, which means global best was known. For the other two cases, Approach 3 found the smallest objective value. It is clear that the virtual search approach not only achieved the global best faster and in fewer iteration cycles but also had smaller deviation ranges. Such qualities can

be utilized to speed up an RT optimization process. Fig. 13 compares average PSO objective values with those of the DPL confirming the superiority of Approach 3.

All algorithms were deemed to converge if they reached to 5% of the optimum solution or stayed within a 5% variation span for at least 5 iterations. Considering such criteria, 30 iterations didn't seem to be enough for the DPL to converge for Patient2 and Patient4. Thus, in Fig. 14, DPL was tested over 200 iterations (versus 30 iterations of the PSO). It appeared that DPL would eventually converge to the global minimum for those two cases, while, for the other three cases, it was trapped in a local minimum. For Patient3, even a small step size corresponding to $n_{DPL} = 1$ caused the DPL algorithm converge to a large objective value. Note that the results shown in Fig. 11 for Patient2 and Patient4 for the DPL are those of the 200 iterations.

VIII. Conclusion

While PSO is well-suited for large-scale, non-convex problems, optimization efficiency is improvable by using the known specifications of the problem at hand. In this paper, a new search approach, termed as the virtual search approach, was introduced and evaluated in radiotherapy inverse planning. In the proposed search approach, a critical objective (here, tumor coverage) was used to define the feasible region, so that the exploration effort of PSO particles could be efficiently reduced. The virtual search approach is a synergistic combination of a hard-constrained approach, where infeasible solutions are penalized (or withdrawn), and a conventionally-used unconstrained approach, where the final solution is artificially adjusted to meet the critical objective. We optimized five 4D CRT clinical plans for lung SBRT patients using the three approaches over ten runs. 90–110 variables (aperture weights) were optimized and it was shown that the virtual search approach was more robust to random initialization and faster in finding and converging to the global minimum. Also, we compared PSO approaches with those of the DPL method, which has been used in 4D RT planning. It was shown that DPL was trapped in local minima in 3 out of 5 cases and in the other two cases, significantly greater number of iterations was required to get to the global solution. We used a simplified RT problem, in this study, where only two structures were considered in order to keep the comparison tractable. In an actual large-scale RT problem, we believe that the efficiency improvement of virtual search PSO, compared to the other techniques investigated in this study, would be even more pronounced. Such fast and robust convergence features make the treatment planning strategy more feasible in a clinical workflow.

Acknowledgments

“This work was partially supported through research funding from National Institutes of Health (R01CA169102) and Varian Medical Systems, Palo Alto, CA, USA”.

The authors would like to thank Wayne Keranen from Varian Medical Systems and Dr. Jun Tan from UT Southwestern for their valuable help on Eclipse scripting and Dicom data management, respectively.

References

1. ASTRO. Fast Facts About Radiation Therapy. ASTRO Targeting Cancer Care; 2012.

2. Podgoršak, EB., Agency, IAE. Radiation Oncology Physics: A Handbook for Teachers and Students. International Atomic Energy Agency; 2005.
3. Varian Medical Systems, Products. <https://www.varian.com/oncology/products>.
4. Bortfeld T. Optimized planning using physical objectives and constraints. *Semin Radiat Oncol*. 1999 Jan; 9(1):20–34. [PubMed: 10196396]
5. Fraass BA, et al. Inverse-optimized 3D conformal planning: minimizing complexity while achieving equivalence with beamlet IMRT in multiple clinical sites. *Med Phys*. 2012 Jun; 39(6):3361–3374. [PubMed: 22755717]
6. Webb S. The physical basis of IMRT and inverse planning. *Br J Radiol*. 2003 Oct; 76(910):678–689. [PubMed: 14512327]
7. Teoh M, et al. Volumetric modulated arc therapy: a review of current literature and clinical use in practice. *The British Journal of Radiology*. 2011; 84(1007):967–996. [PubMed: 22011829]
8. Ma Y, et al. Inverse planning for four-dimensional (4D) volumetric modulated arc therapy. *Med Phys*. 2010 Nov; 37(11):5627–5633. [PubMed: 21158274]
9. Trofimov A, et al. Temporo-spatial IMRT optimization: concepts, implementation and initial results. *Phys Med Biol*. 2005 Jun 21; 50(12):2779–2798. [PubMed: 15930602]
10. Li G, et al. A novel four-dimensional radiotherapy planning strategy from a tumor-tracking beam's eye view. *Phys Med Biol*. 2012; 57(22):7579–7598. [PubMed: 23103415]
11. Nohadani O, et al. Motion management with phase-adapted 4D-optimization. *Phys Med Biol*. 2010 Sep 7; 55(17):5189–5202. [PubMed: 20714043]
12. Keall PJ, et al. Four-dimensional radiotherapy planning for DMLC-based respiratory motion tracking. *Med Phys*. 2005 Apr; 32(4):942–951. [PubMed: 15895577]
13. Rietzel E, et al. Four-dimensional image-based treatment planning: Target volume segmentation and dose calculation in the presence of respiratory motion. *Int J Radiat Oncol Biol Phys*. 2005 Apr 1; 61(5):1535–1550. [PubMed: 15817360]
14. Suh Y, et al. Four-dimensional IMRT treatment planning using a DMLC motion-tracking algorithm. *Phys Med Biol*. 2009 Jun 21; 54(12):3821–3835. [PubMed: 19478383]
15. Keall P. 4-dimensional computed tomography imaging and treatment planning. *Semin Radiat Oncol*. 2004 Jan; 14(1):81–90. [PubMed: 14752736]
16. Craft DL, et al. Improved planning time and plan quality through multicriteria optimization for intensity-modulated radiotherapy. *Int J Radiat Oncol Biol Phys*. 2012; 82(1):6.
17. Lougovski P, et al. Toward truly optimal IMRT dose distribution: inverse planning with voxel-specific penalty. *Technol Cancer Res Treat*. 2010 Dec; 9(6):629–636. [PubMed: 21070085]
18. Cotrutz C, Xing L. Using voxel-dependent importance factors for interactive DVH-based dose optimization. *Phys Med Biol*. 2002 May 21; 47(10):1659–1669. [PubMed: 12069085]
19. Zarepisheh M, et al. A DVH-guided IMRT optimization algorithm for automatic treatment planning and adaptive radiotherapy replanning. *Med Phys*. 2014 Jun; 41(6):061711. [PubMed: 24877806]
20. Kalantzis G, Apte A. A novel reduced-order prioritized optimization method for radiation therapy treatment planning. *IEEE Trans Biomed Eng*. 2014 Apr; 61(4):1062–1070. [PubMed: 24658231]
21. Lu R, et al. Reduced-order parameter optimization for simplifying prostate IMRT planning. *Phys Med Biol*. 2007 Feb 7; 52(3):849–870. [PubMed: 17228125]
22. Eberhart R, Kennedy J. A new optimizer using particle swarm theory. :39–43.
23. Eberhart, R., et al. *Swarm Intelligence*. The Morgan Kaufmann Series in Artificial Intelligence; 2001.
24. Calazan RM, et al. Parallel GPU-based implementation of high dimension Particle Swarm Optimizations. :1–4.
25. Llacer J. Inverse radiation treatment planning using the Dynamically Penalized Likelihood method. *Med Phys*. 1997 Nov; 24(11):1751–1764. [PubMed: 9394282]
26. Llacer J, et al. Comparative behaviour of the dynamically penalized likelihood algorithm in inverse radiation therapy planning. *Phys Med Biol*. 2001 Oct; 46(10):2637–2663. [PubMed: 11686280]
27. Tachibana H, Sawant A. Four-dimensional planning for motion synchronized dose delivery in lung stereotactic body radiation therapy. *Radiother Oncol*. 2016 Apr 30.

28. Bauschke HH, Borwein JM. On Projection Algorithms for Solving Convex Feasibility Problems. *SIAM Review*. 1996; 38(3):367–426.
29. Modiri, A., et al. Improved Swarm Intelligence Solution in Large Scale Radiation Therapy Inverse Planning; Accepted to be published in IEEEExpolre, Great Lake Biomedical Conference; 2015.
30. Modiri A, et al. TH-A-9A-02: BEST IN PHYSICS (THERAPY)-4D IMRT Planning Using Highly-Parallelizable Particle Swarm Optimization. *Medical Physics*. 2014; 41(6):543–543.
31. Modiri A, et al. SU-E-T-06: 4D Particle Swarm Optimization to Enable Lung SBRT in Patients with Central And/or Large Tumors. *Medical Physics*. 2015; 42(6):3331–3332.
32. Van Herk M. Errors and margins in radiotherapy. *Semin Radiat Oncol*. 2004 Jan; 14(1):52–64. [PubMed: 14752733]
33. Timmerman R, et al. Excessive toxicity when treating central tumors in a phase II study of stereotactic body radiation therapy for medically inoperable early-stage lung cancer. *J Clin Oncol*. 2006; 24:4833–4839. [PubMed: 17050868]
34. Murphy, MJ. Adaptive Motion Compensation in Radiotherapy. CRC Press; 2011.
35. RTOG0236. A Phase II Trial of Stereotactic Body Radiation Therapy (SBRT) in the Treatment of Patients with Medically Inoperable Stage I/II Non-Small Cell Lung Cancer. RADIATION THERAPY ONCOLOGY GROUP. 2004
36. A Practical Guide To Intensity-Modulated Radiation Therapy. Madison, Wisconsin: Medical Physics; 2003.
37. Alber M, et al. On the degeneracy of the IMRT optimization problem. *Med Phys*. 2002 Nov; 29(11):2584–2589. [PubMed: 12462725]
38. Michelucci, P. Handbook of Human Computation. Springer; 2013.
39. Nocedal, J., Wright, SJ. Numerical Optimization. Springer; 2000.
40. Fornarelli, G., Mescia, L. Swarm Intelligence for Electric and Electronic Engineering. IGI Global: Engineering Science Reference; 2013.
41. Gentle, JE., et al. Handbook of Computational Statistics. Springer; 2012.
42. Bratton D, Kennedy J. Defining a Standard for Particle Swarm Optimization. :120–127.
43. Jin NB, Rahmat-Samii Y. Advances in particle swarm optimization for antenna designs: Real-number, binary, single-objective and multiobjective implementations. *Ieee Transactions on Antennas and Propagation*. 2007 Mar; 55(3):556–567.
44. Deb, K. Multi-Objective Optimization using Evolutionary Algorithms. WILEY; 2001.
45. Parsopoulos, K., Vrahatis, MN. Particle Swarm Optimization Method for Constrained Optimization Problem. In: Sincak, P,Vascak, J,Kvasnicka, V., Pospichal, J., editors. Intelligent Technologies-Theory and Applications: New Trends in Intelligent Technologies. 2001. p. 214-220.
46. Trelea IC. The particle swarm optimization algorithm: convergence analysis and parameter selection. *Information Processing Letters*. 2003; 85(6):317–325.
47. Aguirre AH, et al. COPSO: Constrained Optimization via PSO algorithm. Center for Research in Mathematics (CIMAT). Technical report No. I-07-04/22-02-2007. 2007
48. Murphy K, et al. Evaluation of registration methods on thoracic CT: the EMPIRE10 challenge. *IEEE Trans Med Imaging*. 2011 Nov; 30(11):1901–1920. [PubMed: 21632295]

Fig. 1.

A gantry-mounted medical linear accelerator designed and manufactured by Varian Medical Systems ([3]). Radiotherapy beams delivered at three angles are shown.

Fig. 2.

(a) PTV in 3D ITV-based RT plan and (b) PTV in an individual respiratory phase for 4D planning.

Fig. 3.
Objective function for a simplified scenario using three different sets of arbitrarily selected priority factors (w 's in (2))

Fig. 4.

(a) shows particles randomly scattered in a 2-variable space. (b) shows the difference between actual and surrogate objective function values for the particles sitting on the bold solid line.

Fig. 5.

The block diagram of (a) Approach 1, (b) Approach 2 and (c) Approach 3.

Fig. 6.
Discontinuous curve modeling feasible region for $f(z_1, z_2)$.

Fig. 7.

Objective function convergence trend for (a) Approach 1, (b) Approach 2 and (c) Approach 3 applied to 2D Rastigrin function as introduced in (6) and (7). Bars show the deviation over 100 runs for 5 PSO particles.

Fig. 8. Objective function convergence for (a) Approach 3 when virtual images replaced the actual particle positions in velocity calculation and (b) when virtual images were only used for objective function calculation. Bars show the deviation over 100 runs for 10 particles.

Fig. 9.
9-beam clinical plan captured on a 3D view for Patient1.

Fig. 10.

DVH curves for the three PSO approaches for all 5 patients ran 10 times: Each column corresponds to a PSO approach as titled and the rows (a) to (e) correspond to Patient1 to Patient5. The one known initial solution is shown in dashed lines and the PSO results in solid lines. The PTV and esophagus curves are shown in blue and red, respectively.

Fig. 11.

DVH curves for the three PSO approaches, averaged over 10 runs, compared to the known initial solution and DPL solutions for (a) esophagus and (b) PTV: Each column corresponds to one patient. The initial known solution is shown in dashed lines while the DPL and the PSO solutions are shown in dotted and solid lines, respectively.

Fig. 12.

Objective function curves for the three PSO approaches for the 5 patients: Bars show the deviation ranges over 10 runs and solid lines show the average objective value over ten runs. Each column corresponds to a PSO approach as titled and rows (a) to (e) correspond to Patient1 to Patient5.

Fig. 13.

Objective function curves for the three PSO approaches averaged over 10 runs (solid lines), as well as the DPL method (dotted lines) for the 5 patients: Rows (a) to (e) correspond to Patient1 to Patient5.

Fig. 14. Objective function curves for (a) Patient2 and (b) Patient4 for the DPL method (dotted lines) over 200 runs compared to the virtual search PSO result (solid line) over 30 iterations

TABLE I

Test Parameters for Fig. 3

Test Structure	Test Parameter	Value
OAR_1	C^{up}	1Gy
	V^{up}	70%
OAR_2	C^{up}	2Gy
	V^{up}	60%
PTV	C^{up}	7.2Gy
	V^{up}	1%
	C^{low}	7Gy
	V^{low}	95%

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE II

Optimization Dose-Volume Constraints

	Maximum Point Dose		Upper Limit			Lower Limit		
	w^{max}	C^{max} (Gy)	w^{up}	V^{up} (%)	C^{up} (Gy)	w^{low}	V^{low} (%)	C^{low} (Gy)
Esophagus	1	15	1	2	10	-	-	-
PTV	-	-	1	1	64.80	1	99	48.60
			1	96	54.01	1	94	53.99