

Quantum-inspired algorithm for radiotherapy planning optimization

Julia M. Pakela^{a)}

Applied Physics Program, University of Michigan, Ann Arbor, MI, USA

Department of Radiation Oncology, University of Michigan, Ann Arbor, MI, USA

Huan-Hsin Tseng, Martha M. Matuszak, Randall K. Ten Haken, and Daniel L. McShan

Department of Radiation Oncology, University of Michigan, Ann Arbor, MI, USA

Issam El Naqa

Applied Physics Program, University of Michigan, Ann Arbor, MI, USA

Department of Radiation Oncology, University of Michigan, Ann Arbor, MI, USA

(Received 11 April 2019; revised 9 September 2019; accepted for publication 9 September 2019; published 7 November 2019)

Purpose: Modern inverse radiotherapy treatment planning requires nonconvex, large-scale optimizations that must be solved within a clinically feasible timeframe. We have developed and tested a quantum-inspired, stochastic algorithm for intensity-modulated radiotherapy (IMRT): quantum tunnel annealing (QTA). By modeling the likelihood probability of accepting a higher energy solution after a particle tunneling through a potential energy barrier, QTA features an additional degree of freedom (the barrier width, w) not shared by traditional stochastic optimization methods such as Simulated Annealing (SA). This additional degree of freedom can improve convergence rates and achieve a more efficient and, potentially, effective treatment planning process.

Methods: To analyze the character of the proposed QTA algorithm, we chose two stereotactic body radiation therapy (SBRT) liver cases of variable complexity. The “easy” first case was used to confirm functionality, while the second case, with a more challenging geometry, was used to characterize and evaluate the QTA algorithm performance. Plan quality was assessed using dose-volume histogram-based objectives and dose distributions. Due to the stochastic nature of the solution search space, extensive tests were also conducted to determine the optimal smoothing technique, ensuring balance between plan deliverability and the resulting plan quality. QTA convergence rates were investigated in relation to the chosen barrier width function, and QTA and SA performances were compared regarding sensitivity to the choice of solution initializations, annealing schedules, and complexity of the dose-volume constraints. Finally, we investigated the extension from beamlet intensity optimization to direct aperture optimization (DAO). Influence matrices were calculated using the Eclipse scripting application program interface (API), and the optimizations were run on the University of Michigan’s high-performance computing cluster, Flux.

Results: Our results indicate that QTA’s barrier-width function can be tuned to achieve faster convergence rates. The QTA algorithm reached convergence up to 46.6% faster than SA for beamlet intensity optimization and up to 26.8% faster for DAO. QTA and SA were ultimately found to be equally insensitive to the initialization process, but the convergence rate of QTA was found to be more sensitive to the complexity of the dose-volume constraints. The optimal smoothing technique was found to be a combination of a Laplace-of-Gaussian (LOG) edge-finding filter implemented as a penalty within the objective function and a two-dimensional Savitzky–Golay filter applied to the final iteration; this achieved total monitor units more than 20% smaller than plans optimized by commercial treatment planning software.

Conclusions: We have characterized the performance of a stochastic, quantum-inspired optimization algorithm, QTA, for radiotherapy treatment planning. This proof of concept study suggests that QTA can be tuned to achieve faster convergence than SA; therefore, QTA may be a good candidate for future knowledge-based or adaptive radiation therapy applications. © 2019 American Association of Physicists in Medicine [<https://doi.org/10.1002/mp.13840>]

Key words: adaptive radiotherapy, IMRT, quantum tunneling optimization, simulated annealing

1. INTRODUCTION

Radiation therapy has been established as one of the primary modalities for cancer treatment, used either exclusively or in combination with other techniques such as chemotherapy or surgery.^{1,2} A critical challenge for radiation therapy (and all

cancer therapies) is to deliver an adequate dose to the tumor to ensure curative or palliative results while minimizing the dose delivered to normal tissues. Intensity-modulated radiation therapy (IMRT) is a type of external beam radiation therapy in which each beam is subdivided into a grid of beamlets whose intensities are determined by dynamic shielding via a

multileaf collimator (MLC). Because IMRT and other radiation therapy techniques which rely on dynamic intensity modulation [such as volumetric arc therapy (VMAT)] are capable of creating concave-shaped dose distributions, they are particularly effective for challenging cases in which the tumor volume is irregular and near critical organs at risk (OARs).^{3,4} The intensity modulations determined from this dynamic shielding optimization are characterized by aperture or beamlet weights. The challenge of calculating optimal weights for a treatment plan often represents a nonconvex,⁵ large-scale optimization problem that must be solved within a clinically reasonable timeframe. The ability to quickly perform robust optimizations is particularly significant in online adaptive radiotherapy, in which a patient's plan may be reoptimized several times during the treatment course to account for changes such as tumor shrinkage or organ deformations.⁶

Quantum computing research is believed to hold promise for achieving computational speedup for certain types of problems.⁷ In quantum computing, classical bits (whose two states are often represented by 0 and 1) are replaced by quantum bits (qubits) which may exist in any linear superposition of 0 and 1.⁸ This allows quantum computers to explore multiple solutions simultaneously, and quantum algorithms can take advantage of this to achieve a significant computational speedup.^{8,9} However, the direct use of quantum computers is still limited by challenges related to creating a proper hardware environment where qubits are maintained in quantum coherence⁷ and the number of qubits deployed is still limited (11–2000^{10–15}) to effectively handle large-scale optimization problems like planning optimization. On the other hand, quantum-inspired algorithms also hold promise for achieving computational speedup of complex optimization problems. Such algorithms are not necessarily quantum processes per se (though some can be formulated to run on a quantum computer); rather, they are quantum simulations designed to run on a classical computer.

The idea of incorporating quantum-inspired techniques into stochastic algorithms was first proposed by de Falco et al. in 1989.¹⁶ A few years later, Kadowaki and Nishimori demonstrated the use of Quantum Annealing (QA) on an Ising model of atomic spins by applying a transverse field, which was annealed to 0° and numerically solving the time-dependent Schrödinger equation for small systems; they found that the probability of reaching the ground state was consistently higher for QA than Simulated Annealing (SA). Many studies have since ensued that have demonstrated QA's potential for a variety of problems.^{16–20}

While QA holds a theoretical promise for certain problem classes with limited dimensionality,²¹ its implementation on a classical computer is impractical for IMRT optimization¹⁹ and deployment on a quantum computer is currently hindered by the limited number of qubits built into existing hardware systems.^{19,21} To avoid these computational limitations, we have implemented another quantum-inspired optimization scheme that models the exploration of higher energy solutions based on the probability of a particle tunneling through a one-dimensional (1D) potential energy barrier. We refer to our algorithm as quantum tunnel annealing (QTA) to

distinguish it from the classical QA algorithms described by de Falco and others.^{16–22} In this paper, we present a proof-of-concept study that (a) demonstrates the behavior of QTA when applied to beamlet intensity and direct aperture optimization for IMRT treatment planning, and (b) compares QTA performance with that of SA as a representative benchmark of traditional optimization methods.

2. MATERIALS AND METHODS

2.A. Quantum tunnel annealing

Quantum tunnel annealing works by modeling an optimization problem as a biased random walk over a fixed number of iterations. During each iteration, a new potential solution (e.g., beamlet-weight vector) is selected from within the neighborhood of the current solution. The energy associated with the new potential solution, given by the objective function, is then calculated and compared against that of the current solution. Potential solutions with lower energies are immediately accepted and set as the current solution. A significant challenge associated with nonconvex optimization problems is that the algorithm can become stuck in a local minimum before it has a chance to reach the globally optimal solution. To avoid this pitfall and ensure adequate exploration of the solution space, QTA simulates quantum fluctuations, allowing the algorithm to accept a worse solution with some probability P . In this process, consider a quantum particle with energy E , traversing through a 1D potential energy landscape, $V(x)$. The particle's wavefunction, $\Psi(x)$, obeys the time-independent Schrödinger equation:

$$H\Psi(x) = E\Psi(x), \quad (1)$$

where the Hamiltonian operator, H , is a function of the particle's potential V and kinetic energy T :

$$H = T + V. \quad (2)$$

Figure 1 illustrates such a particle encountering a potential energy barrier (denoted V^B). The particle's wavefunction prior, during, and after encountering the barrier can be expressed as:

$$\Psi(x) = \begin{cases} Ae^{ikx} + A'e^{-ikx}, & \text{in region A} \\ Be^{kx} + B'e^{-kx}, & \text{in region B} \\ Ce^{ik'x}, & \text{in region C} \end{cases}, \quad (3)$$

$$V(x) = \begin{cases} V^A, & \text{in region A} \\ V^B, & \text{in region B} \\ V^C, & \text{in region C} \end{cases}$$

with wave-numbers²³:

$$k = \sqrt{\frac{2m}{\hbar^2}(\Gamma - V^A)}, \quad \kappa = \sqrt{\frac{2m}{\hbar^2}(V^B - \Gamma)}, \quad (4)$$

$$k' = \sqrt{\frac{2m}{\hbar^2}(\Gamma - V^C)}.$$

A positive exponent represents the particle traveling to the right, and a negative exponent represents the particle traveling to the left. Thus, $A(A')$ and $B(B')$ represent the




FIG. 1. Figurative illustration of a particle (represented by its wave-function, Ψ) tunneling through a potential energy barrier (in region B) in a one-dimensional energy landscape.

amplitudes for the incident (reflected) waves in regions A and B, respectively, and C is the amplitude of the wave transmitted through the barrier. The probability of tunneling through the barrier is given by the transmission coefficient $T = \frac{k'}{k} \left| \frac{C}{A} \right|^2$. This value has suggested by Mukherjee and Chakrabarti to be on the order of $e^{-\frac{w\sqrt{V^B-V^A}}{\Gamma}}$ using a Wentzel–Kramers–Brillouin (WKB) approximation.^{18,23}

Hence, the probability of QTA accepting a worse solution can be redefined to be proportional to $\exp\left(\frac{-w*\sqrt{V_{new}-V_{old}}}{\Gamma}\right)$, where Γ is the kinetic energy of the system (an annealing variable synonymous with the temperature, T, in SA), V_i is the potential energy of the system at solution i defined by the objective function, and w is the width of the barrier being tunneled through. This barrier width is a dynamic parameter, which serves as an additional degree of freedom that is not present in the SA formalism, as discussed in Section 2.B.

2.A.1. Calculation of barrier width

As stated in Section 2.A, the barrier width represents an additional degree of freedom, which QTA can use to obtain an optimal solution in a shorter timeframe. The expected trend in the barrier width's evolution over the course of the optimization can be derived from the following argument: At the start of the optimization, energy barriers that the system encounters have finite widths; as the system approaches its global minimum, the widths of any barriers encountered would grow increasingly large.

In the interest of modeling the barrier width after a physical system in nature, one of the common barrier width schedules tested was modeled after the growth rate of Gallium Arsenide (GaAs) during the process of metal organic chemical vapor deposition (MOCVD). A typical MOCVD setup consists of a reaction chamber and a substrate material on a heated platform. As the substrate is heated by the platform, chemical reactions take place in the gas of the reaction

chamber, leading to the growth of thin films upon the surface of the substrate. In a horizontal-type reaction chamber, the reactants are passed through the chamber horizontally. One of the most common semiconductors grown using MOCVD is GaAs.²⁴ The growth of semiconductors using MOCVD is a complex process influenced by many parameters. It was shown experimentally that GaAs's growth rate is proportional to the square root of the gas velocity.²⁵ Given that kinetic energy is also proportional to the square root of velocity, we can express the growth rate as:

$$\frac{dw}{dt}(t) \propto \sqrt[4]{\Gamma(t)}, \quad (5)$$

where t represents the annealing time defined as the iteration number, and Γ is the kinetic energy of the annealing system, defined in this study as:

$$\Gamma(t) = 10 \times \left(1 - \frac{\log(t)}{\log(N)}\right), \quad (6)$$

with N defined as the total number of iterations performed during the optimization.

The values corresponding to $w(t)$ were obtained using MATLAB's numerical integration function, "integral()" and applying a proportionality factor (k); through trial and error, this was found to work well with $k = 1 \times 10^{-5}$. Both $\frac{dw}{dt}$ and $w(t)$ with $t \leq N = 5 \times 10^5$ iterations are displayed in Figs. 2(a) and 2(b), respectively.

Because QTA occasionally accepts worse solutions, it stands to reason that the barrier width does not grow continuously but rather experiences local width fluctuations combined with a globally increasing trend. Therefore, in addition to the MOCVD-inspired barrier width schedule, another schedule was also tested, defined as:

$$w^a(t) = w^b(t) \left(\sin^2 \frac{50\pi t}{N} + 1 \right), \quad (7)$$

where:

$$w^b(t) = 10 \times \sqrt[3]{[3](w' * t)} \quad (8)$$

with $w' > 0$ used as a tunable parameter to control how quickly the width increases over the course of the optimization. The form of w^a was chosen to introduce more local variations in the barrier width schedule in addition to the global trend of increasing width at a decreasing rate. This was done by coupling a fractional power function (given by w^b) with a sinusoidal function. A squared sine function was chosen to ensure that the width was always at least as large as the global trend. For an annealing schedule where $N = 5 \times 10^5$, the period of 10,000 corresponded to 10 full cycles during the search time.

2.B. Simulated annealing (SA)

For comparison purposes, we used SA, a stochastic search algorithm, which was first introduced for IMRT optimization by Webb in 1989.^{1,26} Like QTA, SA models the optimization




FIG. 2. (a) Barrier width rate extrapolated from metal organic chemical vapor deposition studies by Leys and Veenvliet.²⁵ (b) Barrier width function calculated via numerical integration of (a). (c) Additional width functions explored in this study. [Color figure can be viewed at wileyonlinelibrary.com]

problem as a system which undergoes a biased random walk. Over the course of the random walk, the system will always accept new solutions, which improve on the old solution. In order to avoid getting trapped in local minima, the system accepts worse solutions with a probability proportional to $\exp\left(\frac{-(V_{new}-V_{old})}{T}\right)$, where T is the temperature of the system that is annealed (decreased) over the course of the algorithm search. Mathematically, SA was proven to converge to a global optimal solution with minor assumptions on the cooling schedule and appropriate conditions on irreducibility, aperiodicity, and reversibility of the induced Markov chain.^{27,28} Because SA has a long history of use in our clinic and the literature, it served as our benchmark algorithm for evaluating the success of QTA.²⁹ The annealing schedule for T was identical to the schedule used for the QTA annealing variable, Γ , and is defined in Eq. (6). Note that while the formalism of QTA shares many similarities with SA, the probability of accepting a worse solution in QTA differs from SA in two key respects: (a) reduced dependence on the potential energy difference between the current and new solution and (b) the presence of an additional dynamic parameter in the barrier width. These differences provide QTA with more freedom to explore the solution space.

2.C. IMRT case selection

To analyze the performance of our quantum-inspired algorithm, we compared QTA and SA on two stereotactic body radiation therapy (SBRT) liver cases chosen from the University of Michigan Radiation Oncology Department's clinical database.

Case 1, a 12-field three-dimensional IMRT liver plan, was selected as an “easy” test case to confirm that both QTA and SA were performing properly. This case was not expected to pose a significant challenge for either optimization algorithm because it featured a minimal amount of overlap between the planning target volume (PTV) and the liver, and no overlap with other structures. For simplicity, the structures selected for optimization from Case 1 were the PTV and liver exclusive of the gross tumor volume (Liver - GTV) as shown in Fig. 3(a). Influence matrices for these structures were calculated using built-in functions defined in the Eclipse scripting application program interface (API). The voxel size used was

2 mm and the beamlet size was 5 mm \times 5 mm, for a total of 158 720 voxels, 768 beamlets, and 1 602 504 nonzero elements in the dose influence matrices.

Case 2 served as a “challenge” case to determine if the additional degree of freedom associated with QTA facilitated better results — such as plan quality, robustness, or speed — for more clinically relevant and difficult optimization problems. Designed as a five-field IMRT plan, Case 2 was selected because it had a significant overlap between the PTV, stomach, and liver structures as shown in Fig. 3(b). Because this was a proof of concept study, only a subset of structures from the original treatment plan were included in our optimization. The structures were selected based on the priority assigned to them in the original clinical treatment plan. In addition, the dose-volume histogram (DVH) constraints were also inspired by those used clinically. The influence matrices for these structures (3-mm voxel size, 2.5 mm \times 5 mm beamlet size) were again calculated using built-in functions available in the Eclipse scripting API. Case 2 contained 79 977 voxels, 4166 beamlets, and 1 558 612 nonzero elements in the dose influence matrices. Because Case 2 contained more than four times more beamlet weights, it also represented a more challenging optimization problem than Case 1.

The DVH constraints used in the optimization of Case 1 and Case 2 can be viewed in Tables I and II, respectively.

2.D. Objective function

The objective function used for both SA and QTA IMRT optimization is defined by:

$$\begin{aligned} \min_b E(\mathbf{b}) \\ \text{subject to } \mathbf{b} \geq 0 \end{aligned} \quad (9)$$

where:

$$\begin{aligned} E(\mathbf{b}) = & \sum_{n=1}^N \frac{\lambda_n \|D_n - d_n(\mathbf{b})\|^2}{J_n} + \beta \\ & \sum_{m=1}^M \sum_{ij} |(L \times B_m)_{ij}|^2 \\ & + \sum_{n=1}^N \alpha_n P_n(\mathbf{b}, DVH \text{ constraints}) \end{aligned} \quad (10)$$




FIG. 3. Computed tomography scans show contours for structures optimized for Case 1 (a) and Case 2 (b), respectively. Case 1 features a PTV that is roughly spherical in shape and far from major organs. (with the exception of the liver). Case 2 features a PTV with convex geometrical features and close proximity to both the liver and the stomach. PTV: planning target volume [Color figure can be viewed at wileyonlinelibrary.com]

TABLE I. DVH constraints applied to objective function for Case 1.

Structure	Constraint type	Limit	Volume (%)	Dose (Gy)	Penalty
Case 1 DVH constraints					
PTV	DVH point	Lower	100	29.7	50
	DVH point	Lower	95	30	50
	DVH point	Upper	0	60	100
	Max dose range	N/A	N/A	[30, 42]	100
Liver - GTV	DVH point	Upper	0	42	100

DVH: dose volume histogram; GTV: gross tumor volume; Gy: Gray; PTV: planning target volume.

TABLE II. DVH constraints applied to objective function for Case 2.

Structure	Constraint type	Limit	Volume (%)	Dose (Gy)	Penalty
Case 2 DVH constraints					
PTV	DVH point	Lower	99	33	100
	DVH point	Lower	95	30	200
	DVH point	Lower	100	28	200
	DVH point	Upper	0	48	160
GTV deformed MR	DVH point	Lower	100	43	100
	DVH point	Upper	0	48	160
Liver - GTV	Mean	N/A	N/A	4	50
Stomach	DVH point	Upper	0.001	28	150
Stomach PRV	DVH point	Upper	0.003	25	300

DVH: dose volume histogram; GTV Deformed MR: gross tumor volume deformed from magnetic resonance imaging; Gy: Gray; PRV: planning organ at risk volume; PTV: planning target volume.

and:

$$\mathbf{d}_n(\mathbf{b}) = \mathbf{I}_n * \mathbf{b}. \quad (11)$$

The first term in the objective function represents the mean squared error between the prescribed dose, \mathbf{D}_n , and the

delivered dose, $\mathbf{d}_n(\mathbf{b})$, for each structure n of N structures. $\mathbf{d}_n(\mathbf{b})$ represents the dose delivered to each voxel in structure n , and is defined in Eq. (2) as the product of the structure's influence matrix, \mathbf{I}_n , and the beamlet-weight vector, \mathbf{b} . J_n is the number of voxels in structure n . The influence matrices for each structure were calculated using the Eclipse Scripting API's built-in "CalculateInfluenceMatrixToMemory()" function. The point cloud which was input into this influence matrix function was calculated using an in-house script that generates a normally distributed random set of point locations whose average distance is the cube root of the desired voxel size.

For an influence matrix \mathbf{I} , matrix element \mathbf{I}_{ij} is defined as the dose contribution to voxel i from beamlet j . Any given beamlet is expected to contribute primarily to the voxels it overlaps with and their nearest neighbors. However, due to scattering effects Eclipse-generated influence matrices contain no nonzero values; they contain a subset of elements whose values are orders of magnitude smaller than the largest values in the matrix — corresponding to a beamlet's contribution to a distant voxel. To facilitate faster optimization, a tolerance value was defined below which influence values were deemed negligibly small and reset to 0. This allowed for the influence matrices to be saved as sparse matrices, reducing calculation times. An acceptable tolerance value was determined by trial and error to be 0.015. We loaded fluence vectors that were optimized using filtered influence matrices into the Eclipse scripting API, performed MLC leaf sequencing and dose-volume calculation, and compared the resulting DVH histograms with those produced in-house.

The second term in objective function represents a smoothing penalty which was implemented to ensure the treatment plans could be delivered efficiently. In order to determine the optimal filter, L , a series of QTA optimizations were performed on Case 2 using a number of different filter types — including median, Savitzky-Golay (SG), plan intensity map variation (PIMV), and Laplacian and Laplace of Gaussian (LOG) filters with kernels of sizes 3, 5, 7, 9, and 15, respectively.³⁰ For the smoothing filters, a penalty value was defined as the squared difference between the original

and smoothed fluence map. For the PIMV-type filter, the square of the PIMV value for each beam was used as the penalty. For the edge-finding filters of kernel size n , the filter kernel was convolved with the beamlet matrix \mathbf{B}_m (reshaped from the beamlet weight vector) for each beam. The squared sum of the indices of the resulting matrix yielded a value correlated to the degree of irregularity for each beamlet matrix. With the exception of the Laplace filter and the PIMV filter, all filters tested were implemented using MATLAB built-in functions. Each filter's performance was evaluated by visually inspecting fluence maps and comparing the total number of MUs necessary after MLC leaf sequencing.

The third term in Eq. (10) represents additional penalties based on DVH constraints associated with each structure. The dose constraints and penalties, P_n , used in each case can be viewed in Tables I and II. For Case 1, simple Boolean conditions were used to assign penalty values (e.g., if 99% of the PTV receives <33 Gy, add 100 to the DVH penalty). The weighting factors α_n used in Case 1 were set to 1 for all structures. For the more challenging Case 2, we found it necessary to adjust the calculation of the DVH penalty. Specifically, for Case 2, penalties for missed DVH constraints were assigned as the penalty value (listed in the last columns of Tables I and II) multiplied by the absolute difference between the DVH constraint and the actual metric achieved. For example, if 99% of the PTV volume received ≥ 29 Gy, the penalty for that constraint would be $(30 - 29) \times 100$. Because the constraint type is designated as “lower,” no penalty is assigned if 99% of the PTV volume receives >33 Gy. Finally, for Case 2, α_{PTV} and $\alpha_{Liver-GTV}$ were set to 9 and 10, respectively.

2.E. Extension to influence-based direct aperture optimization

In addition to fluence map optimization, the objective function described in Section 2.D can be generalized to directly optimize apertures (defined by MLC leaf positions) and their weights using a method known as influence-based direct aperture optimization (DAO).^{31,32} This is accomplished by defining the fluence weights as a function of the MLC leaf

segment positions and aperture weights, which for small beamlets can be written as:

$$b(l, w) = \sum_i T(l)_i \times w_i, \quad (12)$$

where l defines the MLC leaf positions, w_i is the weight assigned to aperture i , and $T(l)_i$ is a transmission matrix whose values represent the fraction of each beamlet unobstructed by the MLC leaf segments for aperture i .³¹

2.F. Criteria for convergence

In order to compare QTA and SA's performance in a faithful manner, it is necessary to develop a quantitative method for defining convergence. For both optimization methods, the energy at each iteration t was saved in a vector, $\mathbf{E}(t)$. The gradient of $\mathbf{E}(t)$ was calculated numerically in MATLAB. From this gradient, a moving average mean (MAM) with width 100 was then calculated. A tolerance value, c^{tol} , was selected by trial and error, and the largest index position, j — for which $|\text{MAM}(j)| > c^{\text{tol}}$ — was identified. The convergence point for the algorithm was then defined as iteration $j + 1$. Figure 4 displays the process of finding the j (and thus $j + 1$) from $\mathbf{E}(t)$. An appropriate value for c^{tol} was found to be 0.1.

2.G. Computing environment

All beamlet-weight optimizations described in this paper were performed using MATLAB scripts with GPU acceleration on the University of Michigan's High-Performance Computing Linux-based cluster, Flux (central processing unit (CPU): Intel Haswell, graphics processing unit (GPU): Nvidia K40). Each job was submitted with 2 CPU cores (4 GB/core) and 1 GPU.

MLC leaf sequencing and dose volume calculations used for final plan visualization were performed using clinical software (Varian Medical Systems, Inc. Eclipse Treatment Planning System: Varian Leaf Motion Calculator Version 13.6.23, Anisotropic Analytical Algorithm Version 15.5.11).

The complete QTA algorithm for IMRT optimization is summarized in Fig. 5. The maximum possible number of iterations performed in each run was defined as $N = 5 \times 10^5$.




FIG. 4. Process of calculating the convergence iteration number from a representative QTA optimization. The energy gradient (middle) is calculated from the saved energy history (right). From this gradient a MAM of width 100 was calculated (left). The black vertical line is plotted at the maximum iteration number j for $\max_j(|\text{MAM}(j)|) > c^{\text{tol}}$. MAM: moving average mean; QTA: quantum tunnel annealing. [Color figure can be viewed at wileyonlinelibrary.com]

Algorithm 1: QTA

```

Initialize:  $b_{solution} \leftarrow b_0$ ;
 $V(b) := \sum_{n=1}^N \lambda_n \|D_n - d_n(b)\|^2 + \beta \sum_{m=1}^M \sum_{ij} |(L * B_m)_{ij}|^2 + \sum_{n=1}^N \alpha_n P_n(b, DVHconstraints)$ ;
 $d(b) := I * b$ ;
 $w \leftarrow w_0$ ;
 $\mathcal{T} \leftarrow \mathcal{T}_0$ ;
for  $I_{ij}$  in  $I$  do
  if  $I_{ij} < tol$  then
    |  $I_{ij} = 0$ 
  end
end
for  $n \leftarrow 1$  to  $N_{max}$  do
   $\mathcal{T}_n \leftarrow \mathcal{T}_0(1 - \log(n)/\log(N))$ ;
   $w_n \leftarrow w_{n+1}$ ;
   $b_n \leftarrow \max(0, b_n + r_{neighbor})$ ;
  if  $n = n_{break}$  then
    |  $b_{solution} \leftarrow SG_{filter}(b_{solution})$ ;
    | for  $b_i$  in  $b_{solution}$  do
    |   if  $b_i < 0$  then
    |     |  $b_i = 0$ 
    |   end
    | end
    | break;
  end
  if  $V(b_n) - V(b_{n-1}) < 0$  then
    |  $b_{solution} \leftarrow b_n$ ;
  else
    | if  $\exp(\frac{-w\sqrt{E(b_n) - E(b_{n-1})}}{\mathcal{T}}) > rand$  then
    |   |  $b_{solution} \leftarrow b_n$ ;
    | end
  end
  if  $n = N$  then
    |  $b_{solution} \leftarrow SG_{filter}(b_{solution})$ ;
    | for  $b_i$  in  $b_{solution}$  do
    |   if  $b_i < 0$  then
    |     |  $b_i = 0$ 
    |   end
    | end
    | break;
  end
end
end

```

FIG. 5. Quantum tunnel annealing algorithm for intensity-modulated radiation therapy optimization.

Because the parameter N was used as a variable in both the annealing schedule (\mathcal{T} or Γ) and the barrier width schedule (w), its value was not altered over the course of the reported studies. Therefore, in order to vary the actual number of iterations performed, an additional break parameter was defined which forced the algorithm to end early at iteration $n = n^{break}$. This break parameter was implemented both to shorten the duration of optimizations when it was clear an optimal solution had been reached prior to N as well as to confirm that the convergence iteration numbers — whose calculation was described in Section 2.F — represented clinically acceptable plans.

3. RESULTS

3.A. Case 1

Preliminary studies on a geometrically simple case, designated “Case 1”, confirmed that the QTA and SA algorithms were performing properly. Figures 6(c) (a) through 6(d) display the DVH and potential energy (PE) trajectory results acquired by running the QTA and SA algorithms 20 times

each for $N = 5 \times 10^5$ iterations and no premature breaks (i.e., $n^{break} > N$). Figures 6(e) through 6(f) display representative dose distributions for QA and SA, respectively, which were calculated in Eclipse using optimized beamlet weights from the tenth run. For Case 1, the incorporation of a Laplace edge-finding filter with a kernel size of 3 into the objective function was found to yield sufficiently deliverable plans. Beamlet-weights generated from both QTA and SA were found to consistently yield plans that satisfied the DVH constraints.

The DVH curves for QTA [Fig. 6(a)] and SA [Fig. 6(b)] indicate that for this case, QTA exhibited greater stability over SA with respect to the quality of the final plan. SA converged to a solution with worse PTV coverage 60% of the time. Figures 6(c) and 6(d) display the PE trajectories for the QTA and SA runs, respectively. The PE trajectories for QTA indicate that QTA explored higher energy solutions prior to sudden extreme drops around the ($n = 5 \times 10^5$)th iteration, whereas SA featured a more linear decrease. The resulting dose distributions [Fig. 6(e) and Fig. 6(f)] were found to be similar between both algorithms and featured reasonable tumor coverage while minimizing the dose to the surrounding normal tissues.

3.B. Case 2

3.B.1. Refined smoothing filter

In the pursuit of designing an objective function that can produce clinically acceptable and deliverable plans, a comprehensive study (described in detail in Section 2.D) was performed to determine the optimal measure of smoothness for use as a penalty in the objective function. Smoothness was assessed qualitatively using the fluence maps and quantitatively using the total MU required (summed over each beam). Figure 7(a) displays the optimized fluence map for one of the Case 2 beams using a LOG filter within the objective function. The speckled appearance of 7(a) suggests that smoothing within the objective function alone is not sufficient, and the MU necessary for this plan was more than 20% larger than predicted for an Eclipse-optimized plan which met the same DVH constraints. Adjustments to the size of the kernel and the type of filter used within the objective function did not yield discernable improvement to fluence regularity or total MU.

We also explored directly applying a smoothing filter to the beamlet weights outside of the objective function. We found that the optimal smoothing process consisted of the 7×7 LOG filter within the objective function, combined with a two-dimensional (2D) Savitzky–Golay filter applied to the beamlet-weights during the final iteration of the algorithm. The optimized fluence map using this refined smoothing filter is displayed in Fig. 7(b) and appears markedly smoother than the LOG-filter alone. This refined smoothing filter resulted in a total of 2877 MU, which was 34% lower than the LOG filter alone and more than 20% lower than the Eclipse-optimized plan. The plan quality, as gauged by DVH constraints, experienced only a slight reduction.

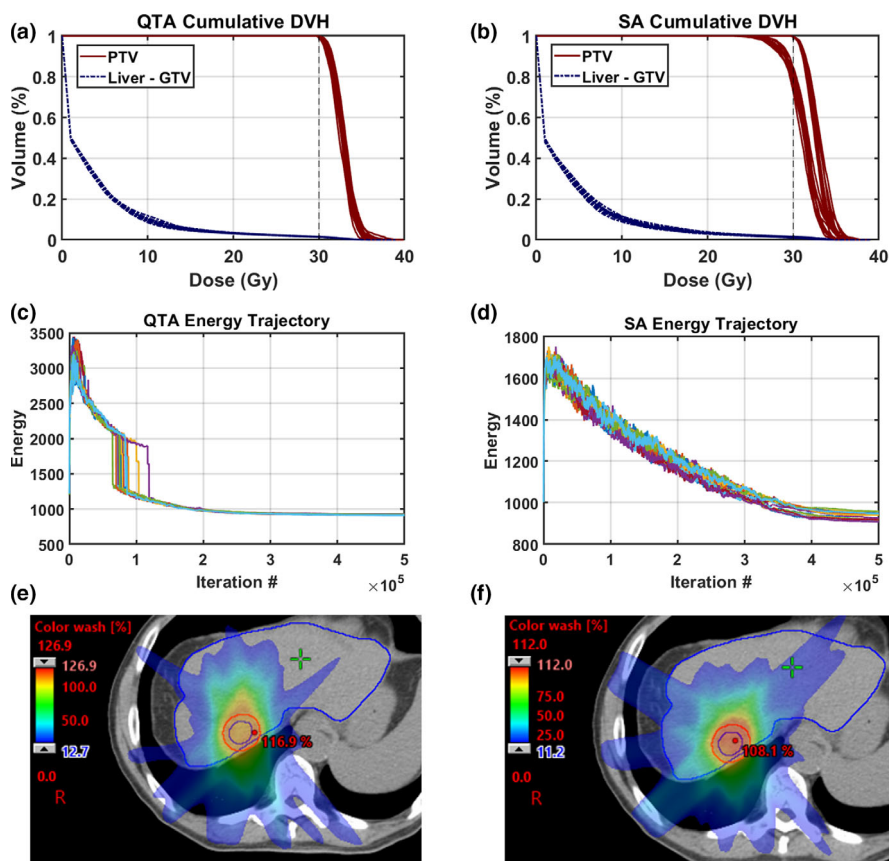


FIG. 6. Optimization results for QTA and SA applied to Case 1. (a) and (b) display DVH curves for 10 separate optimizations using QTA and SA, respectively. (c) and (d) display the PE trajectories for the 10 QTA and SA optimizations. (e) and (f) display representative dose distributions calculated in Eclipse using fluence values from the 10th QTA and SA optimization. DVH: dose volume histogram; PE: potential energy; QTA: quantum tunnel annealing; SA: simulated annealing. [Color figure can be viewed at wileyonlinelibrary.com]

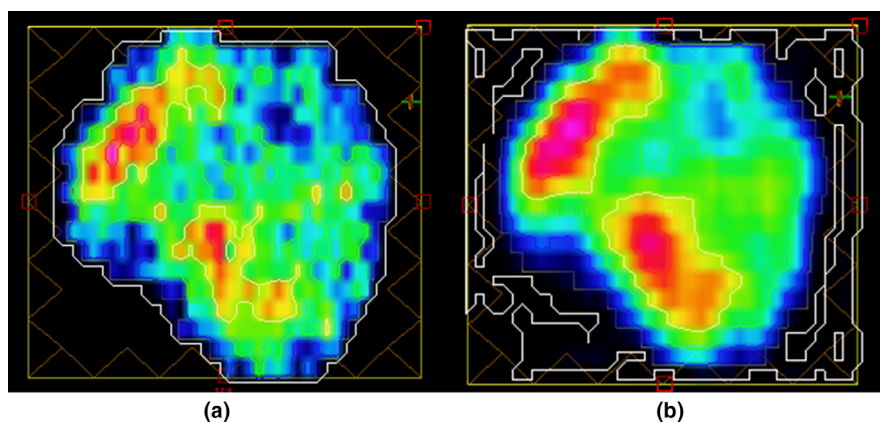


FIG. 7. (a) displays the fluence map results for a single beam in Case 2 resulting from the QTA optimization without refined smoothing. (b) displays the fluence map results from QTA optimization with refined smoothing. QTA: quantum tunnel annealing. [Color figure can be viewed at wileyonlinelibrary.com]

3.B.2. Barrier width schedule effect

As discussed in Section 2.A.1, four different barrier width schedules were investigated for QTA. One was inspired by the growth rate of GaAs in MOCVD, while the remaining three were designed to allow for local fluctuations in the barrier width within a globally increasing

trend. Table III lists convergence rates calculated for QTA optimizations using the four barrier-width schedules as well as optimizations for SA. Three of the four barrier widths tested yielded convergence faster than SA. The optimal barrier width schedule was found to be the w^a function with $w' = 1 \times 10^{-5}$, and it reached convergence in less than half the time required for SA.

TABLE III. Convergence times (in seconds) for QTA with different barrier width schedules as well as SA.

Algorithm type	Convergence (s)
Case 2 convergence results QTA barrier width testing	
SA	1062.5
QTA, $w^{\text{eff}}, w' = 1 \times 10^{-5}$	528.6
QTA, $w^{\text{eff}}, w' = 1 \times 10^{-7}$	637.2
QTA, $w^{\text{eff}}, w' = 1 \times 10^{-9}$	1762
QTA, MOCVD	874.2

MOCVD: metal organic chemical vapor deposition; QTA: quantum tunnel annealing; SA: simulated annealing.

3.B.3. Annealing schedule effect

Each algorithm's sensitivity to the choice of annealing schedule was assessed by comparing their performance across five different functions [shown in Fig. 8(a)]: T1, a linear function; T2, a sigmoidal function; T3, an exponential function; T4, a logarithmic function; and T5, a power law function with fractional exponent. Figure 8(b) displays box and whisker plots of the convergence rates for QTA and SA, respectively, for each annealing schedule. For schedules T1, T2, and T3, SA failed to reach convergence prior to the breakpoint at $n^{\text{break}} = 2.5 \times 10^5$, resulting in the tight spread of data for SA at these schedules. QTA exhibited lower average convergence rates for all five annealing schedules. Note that for this paper, T4 [defined by Eq. (6)] served as the default annealing schedule.

3.B.4. Optimization stability

Quantum tunnel annealing's stability was verified by performing a series of optimization tests using different starting beamlet-weight vectors (10 runs per initial beamlet-weight vector tested, $N = 5 \times 10^5$ iterations, which ran until $n^{\text{break}} = 2.5 \times 10^5$). As a comparison, SA optimizations were also performed under the same conditions. The optimizations began with initial beamlet-weight vector values set to 0, 11, and 20, respectively. These values represent the minimum,

average, and maximum fluence values expected for the optimized beamlet-weight vector. In addition, tests were also conducted using an initial beamlet-weight vector whose values were randomly distributed over a range from 0 to 20. In order to assess whether QTA is primarily advantageous later in the annealing schedule after the algorithm has become stuck in local minima, additional tests were performed on a hybrid SA-QTA algorithm, which ran SA for the first 5×10^4 iterations after which the algorithm switched to QTA. The initial beamlet-weights used for the hybrid tests were also randomly distributed over a range from 0 to 22.

Figure 9 displays the results for 10 QTA and SA optimizations using the randomly distributed initial beamlet-weight vector. The DVH curves for QTA [9(a)] and SA [9(b)] suggest that both reached final solutions with nearly identical dose coverage. This finding was found to hold for all iterations regardless of the initial beamlet-weights used. The energy trajectories for QTA [9(c)] and SA [9(d)] are plotted on a Log scale to highlight differences in the shape of the curves. Like Case 1, the QTA PE trajectories for Case 2 feature a region of rapid descent, located just after the 10^3 iteration. All QTA and SA runs (for $b_0 = 0, 1, 20$, and rand) required $\geq 1.1 \times 10^5$ iterations to reach convergence. Figures 9(e) and 9(f) display the Eclipse-calculated dose distributions from the tenth optimization for QTA and SA, respectively. The final dose distributions were found to be nearly identical and exhibited reasonable dose coverage

Table IV displays the mean convergence rates (in seconds) for QTA, SA, and the hybrid SA-QTA algorithm. QTA consistently exhibited faster convergence rates and had smaller standard deviations than SA in all but one case ($b_0 = 11$). The convergence rates of the SA-QTA hybrid algorithm were similar to the performance of SA.

The stability of QTA and SA was also assessed by making perturbations in the original dose constraints. For each of these tests, a perturbation was made to a single constraint while all others were held constant. Each optimization was run for $N = 5 \times 10^5$ iterations and stopped at $n^{\text{break}} = 2.5 \times 10^5$. Table V summarizes the perturbations tested and the corresponding convergence rates (in seconds). For all perturbation types, QTA exhibited faster convergence.

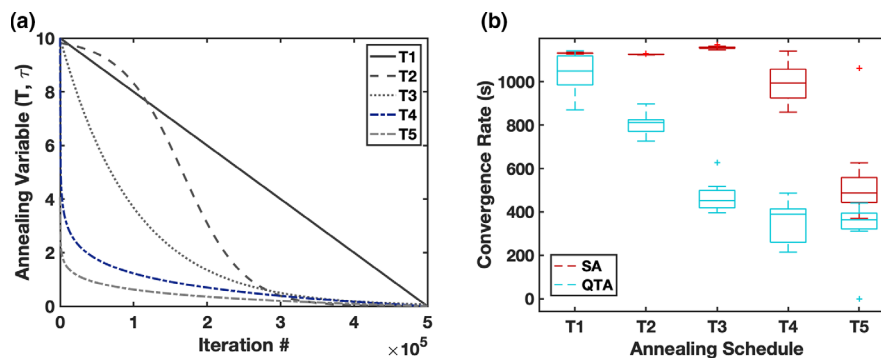


FIG. 8. (a) Displays the annealing schedule functions tested for QTA and SA. Note that T4 was the annealing schedule used for all remaining studies. (b) displays box and whisker plots of the convergence results for QTA and SA, respectively, for each annealing schedule. QTA: quantum tunnel annealing; SA: simulated annealing. [Color figure can be viewed at wileyonlinelibrary.com]

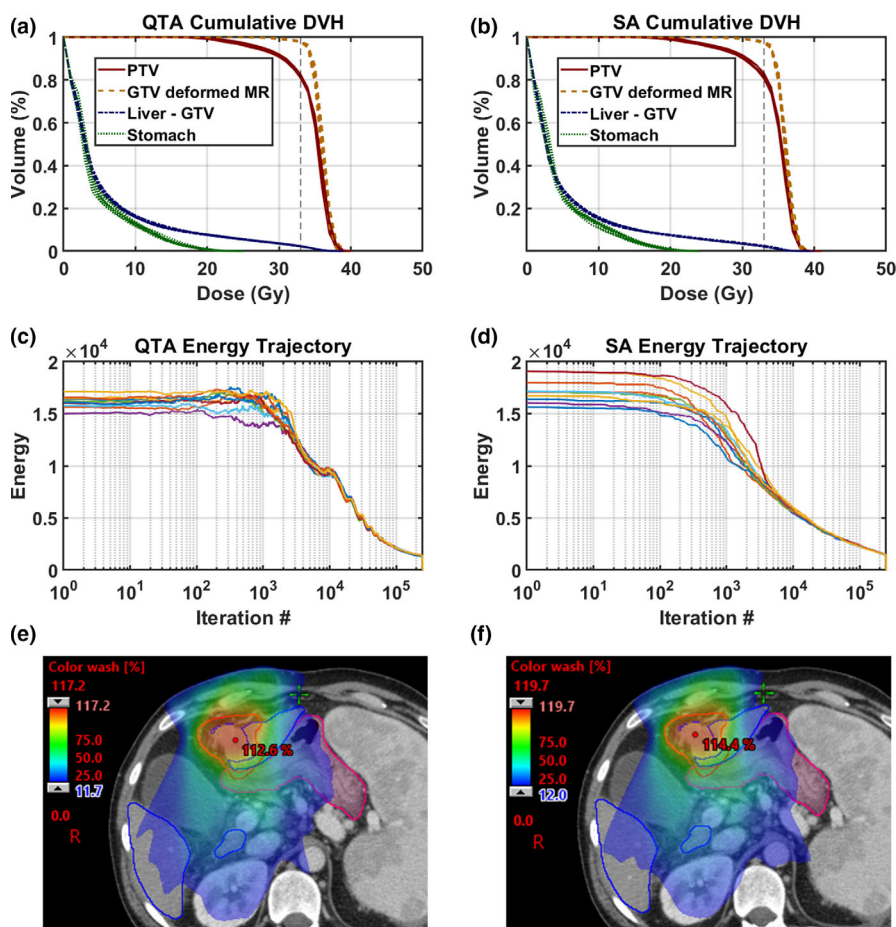


FIG. 9. DVH bands, PE trajectories, and representative dose distributions for stochastic optimizations ($N = 10, 500\,000$ iterations) with the initial beamlet-weight vector set to random values uniformly distributed between 0 and 22 on a challenging SBRT liver case for QTA [(a), (c), and (e)] and SA [(b), (d), and (f)], respectively. DVH: dose-volume histogram; PE: potential energy; QTA: quantum tunnel annealing; SA: simulated annealing; SBRT: stereotactic body radiation therapy. [Color figure can be viewed at wileyonlinelibrary.com]

However, the percent difference in the perturbed convergence rates from the original convergence rate ranged from 5.95%–43.7% for QTA and 4.1%–5.1% for SA, indicating that QTA may exhibit higher sensitivity than SA.

3.B.5. Aperture-weight optimization via influence-based DAO

Influence-based DAO was performed on QTA and SA for 10 runs per initial beamlet-weight vector tested, $N = 5 \times 10^5$ iterations, which ran until $n^{break} = 2.5 \times 10^5$ using the fluence approximation formalism described in Section 2.E. For these optimizations, leaf segment information was extracted from a pre-existing Eclipse-optimized IMRT plan (with a total of 431 apertures) on Case 2, and aperture weights were optimized with the starting weight of each aperture set to 0. Figure 10 displays the resulting cumulative DVHs [10(a) and 10(b)], potential energy trajectories [10(c) and 10(d)], and representative dose distributions [10(d) and 10(e)] for QTA and SA, respectively.

The DVHs displayed in Figs. 10(a) and 10(b) indicate that QTA and SA achieved comparable tumor coverage and OAR sparing. While the energy trajectories in Figs. 10(c) and 10(d)

indicate that QTA exhibited more stochastic exploration of the solution space early on in the optimization, QTA converged within 5234 ± 622.4 (s) on average while SA had an average convergence rate of 7151 ± 504.5 (s). Figures 10(e) and 10(f) show that both algorithms also produced similar dose distributions.

4. DISCUSSION

The optimization results from Case 1 confirmed that both algorithms were capable of delivering clinically acceptable results. QTA was found to be more stable than SA with regard to the quality of the final solution it converged to, as SA converged to a worse solution 60% of the time.

Because it was more geometrically complex, Case 2 was used to characterize QTA's performance. One of the ways QTA distinguishes itself from SA is that the probability of accepting a worse solution during the course of the optimization is a function of the estimated width of the potential energy barrier, providing an additional degree of freedom with which to explore the solution space. We tested several expressions which were heuristically selected to represent the barrier-width function. Adjusting the form of the barrier-

TABLE IV. Mean convergence times (in s) for QTA and SA with perturbations to the initial beamlet-weight values.

b_0	QTA (s)	SA (s)	SA-QTA hybrid (s)
Case 2 convergence results for initial beamlet-weight testing			
0	637.9 ± 63.2	982.4 ± 96.3	N/A
11	644.8 ± 84.4	987.4 ± 82.1	N/A
22	693.2 ± 75.9	1103 ± 84.0	N/A
Randomly distributed	611.0 ± 72.4	996.4 ± 103.0	953.1 ± 65.5

QTA: quantum tunnel annealing; SA: simulated annealing.

TABLE V. Parameter changes and convergence times (in seconds) for QTA and SA with perturbations to the original dose constraints.

Parameter	Organ	From	To	QTA (s)	SA (s)
Case 2 convergence results for parameter testing					
DVH	Stomach	Max dose = 28 Gy	Max dose = 18 Gy	788.8	1038.9
DVH	Liver - GTV	Mean dose = 4 Gy	Mean dose = 2 Gy	951.6	1108.2
DVH and dose	PTV	Target dose = 33 Gy	Target dose = 43 Gy	701.2	1028.6
Original	Original	N/A	N/A	661.8	1083.3

DVH: dose volume histogram; GTV: gross tumor volume; Gy: Gray; PTV: planning target volume; QTA: quantum tunnel annealing; SA: simulated annealing.

width function did not influence the quality of the final plan if the algorithm was allowed to run for its fully allotted time. However, the form of the barrier-width function did influence how quickly the algorithm reached convergence. The convergence results listed in Table IV suggest that the barrier-width function can be used as a tunable parameter to achieve faster convergence. While further tests are warranted to determine an optimal expression for the barrier width, the majority of the functions tested yielded faster convergence rates than SA.

The convergence rates of both algorithms were found to be dependent on the annealing schedule chosen. For three of the five functions tested, SA failed to converge 30–80% of the time, while QTA reached convergence for all five evaluated functions. In addition, QTA had faster mean convergence rates for all five annealing functions tested. These results suggest that QTA is more robust against the choice of annealing schedule. Another way to conceptualize this advantage is to interpret QTA as having a modified annealing schedule in which the barrier width function serves as an additional time-dependent, tunable parameter, coupled with a dampened dependence on the energy difference between the current and new solution.

Testing the sensitivity of QTA with respect to changes in the initial beamlet-weights, b_0 , is useful for determining whether the algorithm can reliably deliver clinically acceptable plans under conditions where a “good” first guess is

unknown. In initial beamlet-weight tests (described in Section 3.B.4), we found that QTA consistently achieved faster convergence times over SA across all variations of b_0 .

Unlike Case 1, Figs. 8(a) and 8(b) suggest that both QTA and SA consistently achieved final solutions of nearly identical plan quality for Case 2. These findings held even after varying the initial starting guess. These results may seem surprising given that Case 2 represented the more challenging case. The explanation lies in the difference between the objective functions used for Cases 1 and 2, which are described in detail in Section 2.D. Case 1 penalties based on the DVH constraints were assigned using Boolean conditions. Implementing the DVH constraint portion of the objective function was found to be insufficient for Case 2 because it could not provide sufficient PTV coverage without delivering an excessive dose to the organs at risk. Therefore, when we began working on Case 2, it was necessary to adjust the objective function so that penalties based on the DVH constraints were weighted more heavily as plan results strayed farther from the objectives. The difference in the results between Case 1 and Case 2 suggests that the additional constraints applied to Case 2’s objective function narrowed the solution space available to the algorithms. In light of this point, the combined results from both cases suggest that QTA is more robust than SA to changes in the formulation of the objective function.

To assess QTA’s sensitivity to changes in treatment plan goals, a series of optimizations were run for QTA and SA in which perturbations were made to the PTV dose prescription and to OAR dose constraints. It was found that while QTA continued to achieve faster convergence rates, those rates exhibited greater variation from the original, unperturbed convergence rate.

For Case 2, it was found that implementing smoothing only within the objective function was insufficient for producing plans with clinically deliverable fluence maps. This is likely due to the algorithms’ stochastic nature and the fact that Case 2 contained more than four times the number of beamlets as Case 1. Ultimately, a refined smoothing technique was developed which combined a LOG filter — used to define an irregularity penalty in the objective function, with a 2D SG filter that was applied to each beamlet map during the final iteration. The resulting fluence maps for these plans had total MU values which were more than 20% less than those for an Eclipse-optimized plan. It is perhaps unorthodox to include a smoothing filter outside of the objective function, as this can compromise plan quality.³⁰ However, we found that implementing the SG filter during the optimization’s final iteration had only a small impact on plan quality, and all plans generated using this technique were comparable in quality to plans generated using Eclipse-based IMRT optimization.

In order to further investigate the potential of QTA over SA, it is necessary to test additional optimization formalisms with known ill-behavior. One such representative approach is to estimate the aperture weights directly using the influence-based DAO approach described in Section 2.E. DAO was evaluated on the more complex Case 2. The results from

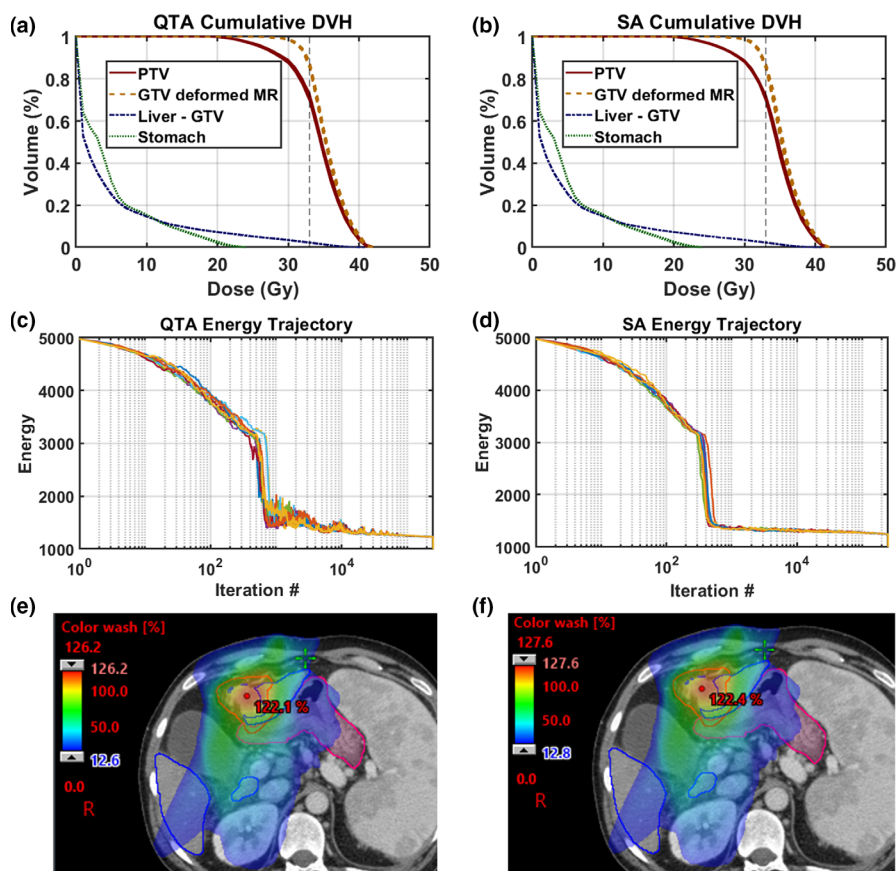


FIG. 10. DVH bands, PE trajectories, and representative dose distributions for stochastic optimizations ($N = 10, 250\,000$ iterations) of aperture weights on a challenging SBRT liver case for QTA [(a), (c), and (e)] and SA [(b), (d), and (f)], respectively. DVH: dose-volume histogram; PE: potential energy; QTA: quantum tunnel annealing; SA: simulated annealing; SBRT: stereotactic body radiation therapy. [Color figure can be viewed at wileyonlinelibrary.com]

these optimizations (presented in Section 3.B.5) indicate that on average QTA converged up to 26.8% faster than SA. DAO is a more complex optimization problem than fluence optimization. While the results of this study example may suggest that the performance gap between QTA and SA seemingly becomes narrower, QTA still exhibits notable benefits over SA overall.

The limitations of this study are summarized as follows: Because only two patient cases were considered, our knowledge of the algorithm's sensitivity to different cases is still limited. We chose to only optimize the most challenging and critical structures in each case; for this reason, the convergence times reported are not representative of a full treatment plan. In addition, the expression used in QTA to define the probability of a particle tunneling through a potential energy barrier contains weaknesses in its assumptions about the size of the annealing variable, Γ . Due to these assumptions, while the formulation for QTA can be described as quantum inspired, it does not represent a true simulation of a quantum process. Nevertheless, QTA was found to exhibit several qualities that suggest it might be an attractive candidate for applications which necessitate rapid optimization of complex or challenging treatment plans. QTA consistently performed faster than SA across multiple types of perturbations and yielded treatment plans of equal quality. Furthermore, a hybrid SA-

QTA algorithm was found to perform only slightly better than SA alone, reinforcing the merit of the full QTA algorithm. The presence of an additional degree of freedom represented by the barrier width schedule leaves open the possibility that this parameter might be further fine-tuned to achieve even faster results.

The results of this study suggest that the extra degree of freedom associated with QTA's barrier-width schedule allowed for the algorithm to be better "tuned" to converge at faster rates than SA. Natural future directions for this work include performing QTA optimizations on full IMRT treatment plans, as well as including VMAT plans, which represent a larger optimization problem. Based on QTA's computational speedup and ability to escape local minima, it may be a useful tool for computationally demanding adaptive radiotherapy applications. Finally, QTA would be a valuable tool for implementing more complex (typically nonconvex) objectives based on biological optimization objections combining imaging and molecular biomarkers with dose-response functions derived via multiple outcome and utility modeling methods,^{33,34} which as of now are hindered in clinical implementation by a lack of efficient and robust optimization techniques.

In addition to further studies applying QTA to more challenging treatment problems, we would also like to explore

whether implementing QTA on a quantum computer could lead to greater computational speedup. In their 2015 study, Nazareth and Spaans reported on the first use of a quantum annealing computer for IMRT beamlet weight optimization; they found that while SA consistently produced higher quality plans, optimizations performed on a quantum annealing device (using a modified version of Tabu Search as the optimization algorithm) were >2.5 times faster than SA.¹⁹ At the time of their study, the researchers were limited to a 512 qubit device, which restricted the complexity of the treatment problems they could tackle. For reference, if the beamlets in Case 1 were discretized using the same method used by Nazareth and Spaans, 5376 qubits would be required. In early 2019, the development of a 5000-qubit commercial quantum annealing computer was announced,³⁵ which would better allow QA to be scaled to higher variable optimization problems but practical clinical application remains a subject for future studies. We believe QTA would be an exciting candidate for quantum computing because it has already shown promise over SA when run on a classical computer.

5. CONCLUSIONS

In this study, we have explored the behavior of a novel algorithm inspired by quantum tunneling, QTA, for IMRT beamlet-weight optimization on two SBRT liver cases. We compared QTA's performance with classical SA, an algorithm which has historically been used for this application. On the easier case, QTA exhibited greater stability than SA. On the challenging case, when allowed to run for the fully allotted number of iterations, both algorithms performed well and exhibited stability with respect to plan quality. With regard to the differences observed between Case 1 and Case 2, it is worth noting that the primary benefit of QTA in a more constrained solution space is the speedup at which it converges, while in a larger (i.e., less constrained) solution space, QTA appears to achieve both faster convergence and plans of more robust quality. Extension to DAO is demonstrated to be feasible with similar performance suggestion potential application of QTA for VMAT-type applications as well.

ACKNOWLEDGMENTS

This research was supported financially in part by the University of Michigan Regents' Fellowship and by National Institutes of Health (NIH) grant R37-CA222215 and R01-CA233487. This research was also supported in part through computational resources and services provided by Advanced Research Computing at the University of Michigan, Ann Arbor.

CONFLICT OF INTEREST

The authors have no conflict of interest to report.

^{a)}Author to whom correspondence should be addressed. Electronic mail: jpakela@umich.edu.

REFERENCES

- Bortfeld T. IMRT: a review and preview. *Phys Med Biol.* 2006;51:R363–R379.
- Citrin DE. Recent developments in radiotherapy. *N Engl J Med.* 2017;377:2200–2201.
- Cho B. Intensity-modulated radiation therapy: a review with a physics perspective. *Radiat Oncol J.* 2018;36:1–10.
- Bortfeld T, Schmidt-Ullrich R, Neve W, Wazer DE. *Image-Guided IMRT.* Berlin, Heidelberg: Springer; 2006.
- Zhang Y, Merritt M. Dose-volume-based IMRT fluence optimization: a fast least-squares approach with differentiability. *Linear Algebra Appl.* 2008;428:1365–1387.
- Lim-Reinders S, Keller BM, Al-Ward S, Sahgal A, Kim A. Online adaptive radiation therapy. *Int J Radiat Oncol Biol Phys.* 2017;99:994–1003.
- Preskill J. Quantum computing and the entanglement frontier. *arXiv preprint arXiv:12035813*; 2012.
- Nielsen MA, Chuang IL. *Quantum Computation and Quantum Information: 10th Anniversary Edition.* Cambridge: Cambridge University Press; 2010: <http://csis.pace.edu/ctappert/cs837-18spring/QC-textbook.pdf>. Accessed 1/29/2019.
- Grover LK. From Schrödinger's equation to the quantum search algorithm. *Am J Phys.* 2001;69:769–777.
- D-Wave Systems Inc. D-Wave Announces D-Wave 2000Q Quantum Computer and First System Order; 2017 January;24; <https://www.dwavesys.com/press-releases/d-wave%C2%A0announces%C2%A0d-wave-2000q-quantum-computer-and-first-system-order>
- IBM. Quantum devices and simulators. <https://www.research.ibm.com/ibm-q/technology/devices/>
- IonQ Inc. A true quantum leap; 2018. <https://ionq.co/>
- Intel. Reinventing data processing with quantum computing. <https://www.intel.com/content/www/us/en/research/quantum-computing.html>
- rigetti. QPU Specifications; 2019. <https://rigetti.com/qpu>
- Google AI. A Preview of Bristlecone. Google's New Quantum Processor; 2018 March 5. <https://ai.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html>
- Apolloni B, Carvalho C, de Falco D. Quantum stochastic optimization. *Stoch Proc Appl.* 1989;33:233–244.
- Farhi E, Goldstone J, Gutmann S, Lapan J, Lundgren A, Preda D. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science.* 2001;292:472–475.
- Mukherjee S, Chakrabarti BK. Multivariable optimization: quantum annealing and computation. *Eur Phys J-Spec Top.* 2015;224:17–24.
- Nazareth DP, Spaans JD. First application of quantum annealing to IMRT beamlet intensity optimization. *Phys Med Biol.* 2015;60:4137–4148.
- Santoro GE, Martoňák R, Tosatti E, Car R. Theory of quantum annealing of an Ising spin glass. *Science.* 2002;295:2427–2430.
- Morita S, Nishimori H. Mathematical foundation of quantum annealing. *J Math Phys.* 2008;49:125210.
- Kadowaki T, Nishimori H. Quantum annealing in the transverse Ising model. *Phys Rev E.* 1998;58:5355.
- Zettili N. *Quantum Mechanics: Concepts and Applications.* New York, NY: John Wiley & Sons, Incorporated; 2009.
- Kuech TF. Metal-organic vapor phase epitaxy of compound semiconductors. *Mater Sci Rep.* 1987;2:1–49.
- Leys MR, Veenliet H. A study of the growth mechanism of epitaxial GaAs as grown by the technique of metal organic vapour phase epitaxy. *J Cryst Growth.* 55, 145–153.
- Webb S. Optimization of conformal radiotherapy dose distributions by simulated annealing. *Phys Med Biol.* 1989;34:1349–1370.
- Bertsimas D, Tsitsiklis J. Simulated annealing. *Stat Sci.* 1993;8:10–15.
- Hajek B. Cooling schedules for optimal annealing. *Mathe Oper Res.* 1988;13:311–329.
- Kessler ML, McShan DL, Epelman MA, et al. Costlets: a generalized approach to cost functions for automated optimization of IMRT treatment plans. *Optimization Eng.* 2005;6:421–448.

30. Matuszak MM, Larsen EW, Fraass BA. Reduction of IMRT beam complexity through the use of beam modulation penalties in the objective function. *Med Phys.* 2007;34:507–520.
31. MacFarlane M, Hoover DA, Wong E, Goldman P, Battista JJ, Chen JZ. A fast inverse direct aperture optimization algorithm for intensity-modulated radiation therapy. *Med Phys.* 2019;46:1127–1139.
32. Härdenmark B, Liander A, Rehbinder H, Löf J, Robinson D. P3IMRT: direct machine parameter optimization. Pinnacle White Paper; 2004:983.
33. Luo Y, McShan DL, Matuszak MM, et al. A multiobjective Bayesian networks approach for joint prediction of tumor local control and radiation pneumonitis in nonsmall-cell lung cancer (NSCLC) for response-adapted radiotherapy. *Med Phys.* 2018;45:3980–3995.
34. Tseng H-H, Luo Y, Cui S, Chien J-T, Ten Haken RK, Naqa IE. Deep reinforcement learning for automated radiation adaptation in lung cancer. *Med Phys.* 2017;44:6690–6705.
35. D-Wave Systems Inc. D-Wave Previews Next-Generation Quantum Computing Platform; 2019 February. <https://www.dwavesys.com/press-releases/d-wave-previews-next-generation-quantum-computing-platform>