

Quantum-Inspired Algorithm for Radiotherapy Planning Optimization

Running Title: Quantum-Inspired RT Optimization

Julia M. Pakela^{1,2} BS, Huan-Hsin Tseng², PhD, Martha M. Matuszak² PhD, Randall K. Ten Haken² PhD, Daniel L. McShan² PhD, Issam El Naqa^{1,2} PhD

1. *Applied Physics Program, University of Michigan, Ann Arbor, MI, USA*
2. *Department of Radiation Oncology, University of Michigan, Ann Arbor, MI, USA*

Corresponding author:

Julia M. Pakela

Department of Radiation Oncology, Physics Division, University of Michigan

519 W. William St, Argus Bldg I, Ann Arbor, MI 48103-4943

Email: jpakela@umich.edu

Abstract

Purpose: Modern inverse radiotherapy treatment planning requires non-convex, 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.

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the [Version of Record](#). Please cite this article as [doi: 10.1002/MP.13840](https://doi.org/10.1002/MP.13840)

This article is protected by copyright. All rights reserved

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

43 **Results:** Our results indicate that QTA’s barrier-width function can be tuned to achieve faster
44 convergence rates. The QTA algorithm reached convergence up to 46.6% faster than SA for
45 beamlet intensity optimization and up to 26.8% faster for DAO. QTA and SA were ultimately
46 found to be equally insensitive to the initialization process, but the convergence rate of QTA was
47 found to be more sensitive to the complexity of the dose-volume constraints. The optimal
48 smoothing technique was found to be a combination of a Laplace-of-Gaussian (LOG) edge-
49 finding filter implemented as a penalty within the objective function and a two-dimensional
50 Savitzky-Golay filter applied to the final iteration; this achieved total monitor units more than
51 20% smaller than plans optimized by commercial treatment planning software.

52 **Conclusions:** We have characterized the performance of a stochastic, quantum-inspired
53 optimization algorithm, QTA, for radiotherapy treatment planning. This proof of concept study
54 suggests that QTA can be tuned to achieve faster convergence than SA; therefore, QTA may be a
55 good candidate for future knowledge-based or adaptive radiation therapy applications.

56 **Keywords:** IMRT, Simulated Annealing, Quantum Tunneling Optimization, Adaptive radiotherapy.

57

58 **1. Introduction**

59 Radiation therapy has been established as one of the primary modalities for cancer treatment,
60 used either exclusively or in combination with other techniques such as chemotherapy or
61 surgery.^{1,2} A critical challenge for radiation therapy (and all cancer therapies) is to deliver an
62 adequate dose to the tumor to ensure curative or palliative results while minimizing the dose
63 delivered to normal tissues. Intensity modulated radiation therapy (IMRT) is a type of external
64 beam radiation therapy in which each beam is subdivided into a grid of beamlets whose
65 intensities are determined by dynamic shielding via a multi-leaf collimator (MLC). Because
66 IMRT and other radiation therapy techniques which rely on dynamic intensity modulation (such
67 as Volumetric Arc Therapy (VMAT)) are capable of creating concave-shaped dose distributions,
68 they are particularly effective for challenging cases in which the tumor volume is irregular and
69 near critical organs at risk (OARs).^{3,4} The intensity modulations determined from this dynamic
70 shielding optimization are characterized by aperture or beamlet weights. The challenge of
71 calculating optimal weights for a treatment plan often represents a non-convex,⁵ large-scale
72 optimization problem that must be solved within a clinically reasonable timeframe. The ability to
73 quickly perform robust optimizations is particularly significant in online adaptive radiotherapy,
74 in which a patient's plan may be reoptimized several times during the treatment course to
75 account for changes such as tumor shrinkage or organ deformations.⁶

76
77 Quantum computing research is believed to hold promise for achieving computational speedup
78 for certain types of problems.⁷ In quantum computing, classical bits (whose two states are often
79 represented by 0 and 1) are replaced by quantum bits (qubits) which may exist in any linear
80 superposition of 0 and 1.⁸ This allows quantum computers to explore multiple solutions
81 simultaneously, and quantum algorithms can take advantage of this to achieve a significant
82 computational speedup.^{8,9} However, the direct use of quantum computers is still limited by
83 challenges related to creating a proper hardware environment where qubits are maintained in
84 quantum coherence⁷ and the number of qubits deployed is still limited (11-2,000¹⁰⁻¹⁵) to
85 effectively handle large scale optimization problems like planning optimization. On the other
86 hand, quantum-inspired algorithms also hold promise for achieving computational speedup of
87 complex optimization problems. Such algorithms are not necessarily quantum processes per se
88 (though some can be formulated to run on a quantum computer); rather, they are quantum
89 simulations designed to run on a classical computer.

90

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

98

99 While QA holds theoretical promise for certain problem classes with limited dimensionality,²¹ its
100 implementation on a classical computer is impractical for IMRT optimization¹⁹ and deployment
101 on a quantum computer is currently hindered by the limited number of qubits built into existing
102 hardware systems.^{19,21} To avoid these computational limitations, we have implemented another
103 quantum-inspired optimization scheme that models the exploration of higher-energy solutions
104 based on the probability of a particle tunneling through a one-dimensional potential energy
105 barrier. We refer to our algorithm as Quantum Tunnel Annealing (QTA) to distinguish it from
106 the classical QA algorithms described by de Falco and others.¹⁶⁻²² In this paper, we present a
107 proof-of-concept study that (1) demonstrates the behavior of QTA when applied to beamlet
108 intensity and direct aperture optimization for IMRT treatment planning, and (2) compares QTA
109 performance with that of SA as a representative benchmark of traditional optimization methods.

110

111 **2. Materials and Methods**

112 **2.1 Quantum Tunnel Annealing**

113 QTA works by modeling an optimization problem as a biased random walk over a fixed number
114 of iterations. During each iteration, a new potential solution (e.g., beamlet-weight vector) is
115 selected from within the neighborhood of the current solution. The energy associated with the
116 new potential solution, given by the objective function, is then calculated and compared against
117 that of the current solution. Potential solutions with lower energies are immediately accepted and
118 set as the current solution. A significant challenge associated with non-convex optimization
119 problems is that the algorithm can become stuck in a local minimum before it has a chance to
120 reach the globally optimal solution. To avoid this pitfall and ensure adequate exploration of the

121 solution space, QTA simulates quantum fluctuations, allowing the algorithm to accept a worse
 122 solution with some probability P . In this process, consider a quantum-particle with energy E ,
 123 traversing through a one-dimensional potential energy landscape, $V(x)$. The particle's
 124 wavefunction, $\Psi(x)$, obeys the time-independent Schrödinger equation:

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

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

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

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

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

134 with wave-numbers:

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

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

144

145 Hence, the probability of QTA accepting a worse solution can be redefined to be proportional to
146 $\exp\left(\frac{-w * \sqrt{V_{new} - V_{old}}}{\Gamma}\right)$, where Γ is the kinetic energy of the system (an annealing variable
147 synonymous with the temperature, T , in SA), V_i is the potential energy of the system at solution i
148 defined by the objective function, and w is the width of the barrier being tunneled through. This
149 barrier width is a dynamic parameter, which serves as an additional degree of freedom that is not
150 present in the SA formalism, as discussed in Section 2.2.

151

152 2.1.1 Calculation of barrier width

153 As stated in Section 2.1, the barrier width represents an additional degree of freedom, which
154 QTA can use to obtain an optimal solution in a shorter timeframe. The expected trend in the
155 barrier width's evolution over the course of the optimization can be derived from the following
156 argument: At the start of the optimization, energy barriers that the system encounters have finite
157 widths; as the system approaches its global minimum, the widths of any barriers encountered
158 would grow increasingly large.

159

160 In the interest of modeling the barrier width after a physical system in nature, one of the common
161 barrier width schedules tested was modeled after the growth rate of Gallium Arsenide (GaAs)
162 during the process of metal organic chemical vapor deposition (MOCVD). A typical MOCVD
163 setup consists of a reaction chamber and a substrate material on a heated platform. As the
164 substrate is heated by the platform, chemical reactions take place in the gas of the reaction
165 chamber, leading to the growth of thin films upon the surface of the substrate. In a horizontal-
166 type reaction chamber, the reactants are passed through the chamber horizontally. One of the
167 most common semiconductors grown using MOCVD is GaAs.²⁴ The growth of semiconductors
168 using MOCVD is a complex process influenced by many parameters. It was shown
169 experimentally that GaAs's growth rate is proportional to the square root of the gas velocity.²⁵
170 Given that kinetic energy is also proportional to the square root of velocity, we can express the
171 growth rate as:

172

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

173

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

176

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

177

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

179 The values corresponding to $w(t)$ were obtained using MATLAB's numerical integration
180 function, "integral()" and applying a proportionality factor (k); through trial and error, this was
181 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
182 displayed in Figure 2(a) and 2(b), respectively.

183

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

188

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

189

190 where:

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

191

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

199

200 **2.2 Simulated Annealing (SA)**

201 For comparison purposes, we used SA, a stochastic search algorithm, which was first introduced
202 for IMRT optimization by Webb in 1989.^{1,26} Like QTA, SA models the optimization problem as
203 a system which undergoes a biased random walk. Over the course of the random walk, the
204 system will always accept new solutions, which improve on the old solution. In order to avoid
205 getting trapped in local minima, the system accepts worse solutions with a probability
206 proportional to $\exp\left(\frac{-(V_{new} - V_{old})}{T}\right)$, where T is the temperature of the system that is annealed
207 (decreased) over the course of the algorithm search. Mathematically, SA was proven to converge
208 to a global optimal solution with minor assumptions on the cooling schedule and appropriate
209 conditions on irreducibility, aperiodicity and reversibility of the induced Markov chain.^{27,28}
210 Because SA has a long history of use in our clinic and the literature, it served as our benchmark
211 algorithm for evaluating the success of QTA.²⁹ The annealing schedule for T was identical to the
212 schedule used for the QTA annealing variable, Γ , and is defined in Equation 6. Note that while
213 the formalism of QTA shares many similarities with SA, the probability of accepting a worse
214 solution in QTA differs from SA in two key respects: (1) reduced dependence on the potential
215 energy difference between the current and new solution and (2) the presence of an additional
216 dynamic parameter in the barrier width. These differences provide QTA with more freedom to
217 explore the solution space.

218

219 **2.3 IMRT Case Selection**

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

223

224 *Case 1*, a 12-field three-dimensional IMRT liver plan, was selected as an “easy” test case to
225 confirm that both QTA and SA were performing properly. This case was not expected to pose a
226 significant challenge for either optimization algorithm because it featured a minimal amount of
227 overlap between the planning target volume (PTV) and the liver, and no overlap with other
228 structures. For simplicity, the structures selected for optimization from Case 1 were the PTV and
229 liver exclusive of the gross target volume (Liver – GTV) as shown in Figure 3a. Influence

230 matrices for these structures were calculated using built-in functions defined in the Eclipse
231 scripting application program interface (API). The voxel size used was 2mm and the beamlet
232 size was 5 mm x 5 mm, for a total of 158,720 voxels, 768 beamlets, and 1,602,504 nonzero
233 elements in the dose influence matrices.

234
235 *Case 2* served as a “challenge” case to determine if the additional degree of freedom associated
236 with QTA facilitated better results – such as plan quality, robustness, or speed – for more
237 clinically relevant and difficult optimization problems. Designed as a 5-field IMRT plan, Case 2
238 was selected because it had significant overlap between the PTV, stomach, and liver structures as
239 shown in Figure 3b. Because this was a proof of concept study, only a subset of structures from
240 the original treatment plan were included in our optimization. The structures were selected based
241 on the priority assigned to them in the original clinical treatment plan. In addition, the dose
242 volume histogram (DVH) constraints were also inspired by those used clinically. The influence
243 matrices for these structures (3mm voxel size, 2.5mm x 5mm beamlet size) were again calculated
244 using built-in functions available in the Eclipse scripting API. Case 2 contained 79,977 voxels,
245 4166 beamlets, and 1,558,612 nonzero elements in the dose influence matrices. Because Case 2
246 contained more than four times more beamlet weights, it also represented a more challenging
247 optimization problem than Case 1.

248
249 The DVH constraints used in the optimization of Case 1 and Case 2 can be viewed in Table 1
250 and Table 2, respectively.

251 252 **2.4 Objective function**

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

$$\begin{aligned} & \min_{\mathbf{b}} E(\mathbf{b}) \\ & \text{subject to } \mathbf{b} \geq 0 \end{aligned} \tag{9}$$

255 where:

256

This

$$E(\mathbf{b}) \tag{10}$$

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

257 and:

258

$$\mathbf{d}_n(\mathbf{b}) = I_n * \mathbf{b}. \tag{11}$$

259

260 The first term in the objective function represents the mean squared error between the prescribed
 261 dose, D_n , and the delivered dose, $d_n(\mathbf{b})$, for each structure n of N structures. $d_n(\mathbf{b})$ represents
 262 the dose delivered to each voxel in structure n , and is defined in Equation 2 as the product of the
 263 structure's influence matrix, I_n , and the beamlet-weight vector, \mathbf{b} . J_n is the number of voxels in
 264 structure n . The influence matrices for each structure were calculated using the Eclipse Scripting
 265 API's built-in "CalculateInfluenceMatrixToMemory()" function. The point cloud which was
 266 input into this influence matrix function was calculated using an in-house script that generates a
 267 normally distributed random set of point locations whose average distance is the cube root of the
 268 desired voxel size.

269

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

282

283 The second term in objective function represents a smoothing penalty which was implemented to
284 ensure the treatment plans could be delivered efficiently. In order to determine the optimal filter,
285 L , a series of QTA optimizations were performed on Case 2 using a number of different filter
286 types – including median, Savitzky-Golay (SG), plan intensity map variation (PIMV), and
287 Laplacian and Laplace of Gaussian (LOG) filters with kernels of sizes 3, 5, 7, 9, and 15,
288 respectively.³⁰ For the smoothing filters, a penalty value was defined as the squared difference
289 between the original and smoothed fluence map. For the PIMV-type filter, the square of the
290 PIMV value for each beam was used as the penalty. For the edge-finding filters of kernel size n ,
291 the filter kernel was convolved with the beamlet matrix \mathbf{B}_m (reshaped from the beamlet weight
292 vector) for each beam. The squared sum of the indices of the resulting matrix yielded a value
293 correlated to the degree of irregularity for each beamlet matrix. With the exception of the
294 Laplace filter and the PIMV filter, all filters tested were implemented using MATLAB built-in
295 functions. Each filter's performance was evaluated by visually inspecting fluence maps and
296 comparing the total number of MUs necessary after MLC leaf sequencing.

297

298 The third term in Equation 10 represents additional penalties based on DVH constraints
299 associated with each structure. The dose constraints and penalties, P_n , used in each case can be
300 viewed in Tables 1 and 2. For Case 1, simple Boolean conditions were used to assign penalty
301 values (for example, if 99% of the PTV receives < 33 Gy, add 100 to the DVH penalty). The
302 weighting factors α_n used in Case 1 were set to 1 for all structures. For the more challenging
303 Case 2, we found it necessary to adjust the calculation of the DVH penalty. Specifically, for Case
304 2, penalties for missed DVH constraints were assigned as the penalty value (listed in the last
305 columns of Tables 1 and 2) multiplied by the absolute difference between the DVH constraint
306 and the actual metric achieved. For example, if 99% of the PTV volume received ≥ 29 Gy, the
307 penalty for that constraint would be $(30-29) \times 100$. Because the constraint type is designated as
308 “lower”, no penalty is assigned if 99% of the PTV volume receives > 33 Gy. Finally, for Case 2,
309 α_{PTV} and $\alpha_{Liver-GTV}$ were set to 9 and 10, respectively.

310

311 **2.5 Extension to Influence-based Direct Aperture Optimization**

312 In addition to fluence map optimization, the objective function described in Section 2.4 can be
313 generalized to directly optimize apertures (defined by MLC leaf positions) and their weights
314 using a method known as influence-based direct aperture optimization (DAO).^{31,32} This is
315 accomplished by defining the fluence weights as a function of the MLC leaf segment positions
316 and aperture weights, which for small beamlets can be written as:

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

319
320
321 where l defines the MLC leaf positions, w_i is the weight assigned to aperture i , and $T(l)_i$ is a
322 transmission matrix whose values represent the fraction of each beamlet unobstructed by the
323 MLC leaf segments for aperture i .³¹

324

325 **2.6 Criteria for Convergence**

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

334

335 **2.7 Computing Environment**

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

340

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

345
346 The complete QTA algorithm for IMRT optimization is summarized in Figure 5. The maximum
347 possible number of iterations performed in each run was defined as $N = 5 * 10^5$. Because the
348 parameter N was used as a variable in both the annealing schedule (T or Γ) and the barrier width
349 schedule (w), its value was not altered over the course of the reported studies. Therefore, in order
350 to vary the actual number of iterations performed, an additional break parameter was defined
351 which forced the algorithm to end early at iteration $n = n^{\text{break}}$. This break parameter was
352 implemented both to shorten the duration of optimizations when it was clear an optimal solution
353 had been reached prior to N as well as to confirm that the convergence iteration numbers –
354 whose calculation was described in Section 2.6 – represented clinically acceptable plans.

355

356 3. Results

357 3.1 Case 1

358 Preliminary studies on a geometrically simple case, designated “Case 1”, confirmed that the
359 QTA and SA algorithms were performing properly. Figures 5(a) through 5(d) display the DVH
360 and potential energy (PE) trajectory results acquired by running the QTA and SA algorithms 20
361 times each for $N = 5 \times 10^5$ iterations and no premature breaks (i.e., $n^{\text{break}} > N$). Figures 5(e)
362 through 5(f) display representative dose distributions for QA and SA, respectively, which were
363 calculated in Eclipse using optimized beamlet-weights from the tenth run. For Case 1, the
364 incorporation of a Laplace edge-finding filter with a kernel size of 3 into the objective function
365 was found to yield sufficiently deliverable plans. Beamlet-weights generated from both QTA and
366 SA were found to consistently yield plans that satisfied the DVH constraints.

367
368 The DVH curves for QTA (5(a)) and SA (5(b)) indicate that for this case, QTA exhibited greater
369 stability over SA with respect to the quality of the final plan. SA converged to a solution with
370 worse PTV coverage 60% of the time. Figures 5(c) and 5(d) display the PE trajectories for the
371 QTA and SA runs, respectively. The PE trajectories for QTA indicate that QTA explored higher

372 energy solutions prior to sudden extreme drops around the ($n = 1 \times 10^5$)th iteration, whereas
373 SA featured a more linear decrease. The resulting dose distributions were found to be similar
374 between both algorithms and featured reasonable tumor coverage while minimizing the dose to
375 the surrounding normal tissues.

376

377 **3.2 Case 2**

378 **3.2.1 Refined Smoothing Filter**

379 In the pursuit of designing an objective function that can produce clinically acceptable and
380 deliverable plans, a comprehensive study (described in detail in Section 2.4) was performed to
381 determine the optimal measure of smoothness for use as a penalty in the objective function.
382 Smoothness was assessed qualitatively using the fluence maps and quantitatively using the total
383 MU required (summed over each beam). Figure 7(a) displays the optimized fluence map for one
384 of the Case 2 beams using a LOG filter within the objective function. The speckled appearance
385 of 7(a) suggests that smoothing within the objective function alone is not sufficient, and the MU
386 necessary for this plan was more than 20% larger than predicted for an Eclipse-optimized plan
387 which met the same DVH constraints. Adjustments to the size of the kernel and the type of filter
388 used within the objective function did not yield discernable improvement to fluence regularity or
389 total MU.

390

391 We also explored directly applying a smoothing filter to the beamlet weights outside of the
392 objective function. We found that the optimal smoothing process consisted of the 7x7 LOG filter
393 within the objective function, combined with a two-dimensional Savitzky-Golay filter applied to
394 the beamlet-weights during the final iteration of the algorithm. The optimized fluence map using
395 this refined smoothing filter is displayed in Figure 7(b) and appears markedly smoother than the
396 LOG-filter alone. This refined smoothing filter resulted in a total of 2877 MU, which was 34%
397 lower than the LOG filter alone and more than 20% lower than the Eclipse-optimized plan. The
398 plan quality, as gauged by DVH constraints, experienced only a slight reduction.

399

400 **3.2.2 Barrier Width Schedule Effect**

401 As discussed in Section 2.1.2, four different barrier width schedules were investigated for QTA.
402 One was inspired by the growth rate of GaAs in MOCVD, while the remaining three were

403 designed to allow for local fluctuations in the barrier width within a globally increasing trend.
404 Table 3 lists convergence rates calculated for QTA optimizations using the four barrier-width
405 schedules as well as optimizations for SA. Three of the four barrier widths tested yielded
406 convergence faster than SA. The optimal barrier width schedule was found to be the w^a function
407 with $w' = 1 \times 10^{-5}$, and it reached convergence in less than half the time required for SA.

408

409 3.2.3 Annealing Schedule Effect

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

419

420 3.2.4 Optimization Stability

421 QTA's stability was verified by performing a series of optimization tests using different starting
422 beamlet-weight vectors (10 runs per initial beamlet-weight vector tested, $N = 5 \times 10^5$ iterations,
423 which ran until $n^{break} = 2.5 \times 10^5$). As a comparison, SA optimizations were also performed
424 under the same conditions. The optimizations began with initial beamlet-weight vector values set
425 to 0, 11, and 20, respectively. These values represent the minimum, average, and maximum
426 fluence values expected for the optimized beamlet-weight vector. In addition, tests were also
427 conducted using an initial beamlet-weight vector whose values were randomly distributed over a
428 range from 0-20. In order to assess whether QTA is primarily advantageous later in the
429 annealing schedule after the algorithm has become stuck in local minima, additional tests were
430 performed on a hybrid SA-QTA algorithm, which ran SA for the first 5×10^4 iterations after
431 which the algorithm switched to QTA. The initial beamlet-weights used for the hybrid tests were
432 also randomly distributed over a range from 0-22.

433

434 Figure 9 displays the results for 10 QTA and SA optimizations using the randomly distributed
435 initial beamlet-weight vector. The DVH curves for QTA (9(a)) and SA (9(b)) suggest that both
436 reached final solutions with nearly identical dose coverage. This finding was found to hold for all
437 iterations regardless of the initial beamlet-weights used. The energy trajectories for QTA (9(c))
438 and SA (9(d)) are plotted on a Log scale to highlight differences in the shape of the curves. Like
439 Case 1, the QTA PE trajectories for Case 2 feature a region of rapid descent, located just after the
440 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
441 reach convergence. Figures 9(e) and 9(f) display the Eclipse-calculated dose distributions from
442 the tenth optimization for QTA and SA, respectively. The final dose distributions were found to
443 be nearly identical and exhibited reasonable dose coverage.

444
445 Table 4 displays the mean convergence rates (in seconds) for QTA, SA, and the hybrid SA-QTA
446 algorithm. QTA consistently exhibited faster convergence rates and had smaller standard
447 deviations than SA in all but one case ($b_0 = 11$). The convergence rates of the SA-QTA hybrid
448 algorithm were similar to the performance of SA.

449
450 The stability of QTA and SA was also assessed by making perturbations in the original dose
451 constraints. For each of these tests, a perturbation was made to a single constraint while all others
452 were held constant. Each optimization was run for $N = 5 \times 10^5$ iterations and stopped at n^{break}
453 $= 2.5 \times 10^5$. Table 5 summarizes the perturbations tested and the corresponding convergence
454 rates (in seconds). For all perturbation types, QTA exhibited faster convergence. However, the
455 percent difference in the perturbed convergence rates from the original convergence rate ranged
456 from 5.95%-43.7% for QTA and 4.1%-5.1% for SA, indicating that QTA may exhibit higher
457 sensitivity than SA.

458 459 **3.2.5 Aperture-Weight Optimization via Influence-Based DAO**

460 Influence-based DAO was performed on QTA and SA for 10 runs per initial beamlet-weight
461 vector tested, ($N = 5 \times 10^5$ iterations, which ran until $n^{break} = 2.5 \times 10^5$) using the fluence
462 approximation formalism described in Section 2.5. For these optimizations, leaf segment
463 information was extracted from a pre-existing Eclipse-optimized IMRT plan (with a total of 431
464 apertures) on Case 2, and aperture weights were optimized with the starting weight of each

465 aperture set to 0. Figure 10 displays the resulting cumulative DVHs (10(a) and 10(b)), potential
466 energy trajectories (10(c) and 10(d)), and representative dose distributions (10(d) and 10(e)) for
467 QTA and SA, respectively.

468
469 The DVHs displayed in Figures 10(a) and 10(b) indicate that QTA and SA achieved comparable
470 tumor coverage and OAR sparing. While the energy trajectories in Figures 10(c) and 10(d)
471 indicate that QTA exhibited more stochastic exploration of the solution space early on in the
472 optimization, QTA converged within $5,234 \pm 622.4$ (s) on average while SA had an average
473 convergence rate of $7,151 \pm 504.5$ (s). Figures 10(e) and 10(f) show that both algorithms also
474 produced similar dose distributions.

475

476 **4. Discussion**

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

480

481 Because it was more geometrically complex, Case 2 was used to characterize QTA's
482 performance. One of the ways QTA distinguishes itself from SA is that the probability of
483 accepting a worse solution during the course of the optimization is a function of the estimated
484 width of the potential energy barrier, providing an additional degree of freedom with which to
485 explore the solution space. We tested several expressions which were heuristically selected to
486 represent the barrier-width function. Adjusting the form of the barrier-width function did not
487 influence the quality of the final plan if the algorithm was allowed to run for its fully allotted
488 time. However, the form of the barrier-width function did influence how quickly the algorithm
489 reached convergence. The convergence results listed in Table 4 suggest that the barrier-width
490 function can be used as a tunable parameter to achieve faster convergence. While further tests are
491 warranted to determine an optimal expression for the barrier width, the majority of the functions
492 tested yielded faster convergence rates than SA.

493

494 The convergence rates of both algorithms were found to be dependent on the annealing schedule
495 chosen. For three of the five functions tested, SA failed to converge 30-80% of the time, while

496 QTA reached convergence for all five evaluated functions. In addition, QTA had faster mean
497 convergence rates for all five annealing functions tested. These results suggest that QTA is more
498 robust against the choice of annealing schedule. Another way to conceptualize this advantage is
499 to interpret QTA as having a modified annealing schedule in which the barrier width function
500 serves as an additional time-dependent, tunable parameter, coupled with a dampened dependence
501 on the energy difference between the current and new solution.

502
503 Testing the sensitivity of QTA with respect to changes in the initial beamlet-weights, b_0 , is useful
504 for determining whether the algorithm can reliably deliver clinically acceptable plans under
505 conditions where a “good” first guess is unknown. In initial beamlet-weight tests (described in
506 Section 3.2.4) we found that QTA consistently achieved faster convergence times over SA across
507 all variations of b_0 .

508
509 Unlike Case 1, Figures 8(a) and 8(b) suggest that both QTA and SA consistently achieved final
510 solutions of nearly identical plan quality for Case 2. These findings held even after varying the
511 initial starting guess. These results may seem surprising given that Case 2 represented the more
512 challenging case. The explanation lies in the difference between the objective functions used for
513 Cases 1 and 2, which are described in detail in Section 2.4. Case 1 penalties based on the DVH
514 constraints were assigned using Boolean conditions. Implementing the DVH constraint portion
515 of the objective function was found to be insufficient for Case 2 because it could not provide
516 sufficient PTV coverage without delivering an excessive dose to the organs at risk. Therefore,
517 when we began working on Case 2, it was necessary to adjust the objective function so that
518 penalties based on the DVH constraints were weighted more heavily as plan results strayed
519 farther from the objectives. The difference in results between Case 1 and Case 2 suggest that the
520 additional constraints applied to Case 2’s objective function narrowed the solution space
521 available to the algorithms. In light of this point, the combined results from both cases suggest
522 that QTA is more robust than SA to changes in the formulation of the objective function.

523
524 To assess QTA’s sensitivity to changes in treatment plan goals, a series of optimizations were
525 run for QTA and SA in which perturbations were made to the PTV dose prescription and to OAR

526 dose constraints. It was found that while QTA continued to achieve faster convergence rates,
527 those rates exhibited greater variation from the original, unperturbed convergence rate.

528

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

541

542 In order to further investigate the potential of QTA over SA, it is necessary to test additional
543 optimization formalisms with known ill-behavior. One such representative approach is to
544 estimate the aperture weights directly using the influence-based DAO approach described in
545 Section 2.5. DAO was evaluated on the more complex Case 2. The results from these
546 optimizations (presented in Section 3.2.5) indicate that on average QTA converged up to 26.8%
547 faster than SA. DAO is a more complex optimization problem than fluence optimization. While
548 the results of this study example may suggest that the performance gap between QTA and SA
549 seemingly becomes narrower, QTA still exhibits notable benefits over SA overall.

550

551 The limitations of this study are summarized as follows: Because only two patient cases were
552 considered, our knowledge of the algorithm's sensitivity to different cases is still limited. We
553 chose to only optimize the most challenging and critical structures in each case; for this reason,
554 the convergence times reported are not representative of a full treatment plan. In addition, the
555 expression used in QTA to define the probability of a particle tunneling through a potential
556 energy barrier contains weaknesses in its assumptions about the size of the annealing variable, Γ .

557 Due to these assumptions, while the formulation for QTA can be described as quantum-inspired,
558 it does not represent a true simulation of a quantum process. Nevertheless, QTA was found to
559 exhibit several qualities that suggest it might be an attractive candidate for applications which
560 necessitate rapid optimization of complex or challenging treatment plans. QTA consistently
561 performed faster than SA across multiple types of perturbations and yielded treatment plans of
562 equal quality. Furthermore, a hybrid SA-QTA algorithm was found to perform only slightly
563 better than SA alone, reinforcing the merit of the full QTA algorithm. The presence of an
564 additional degree of freedom represented by the barrier width schedule leaves open the
565 possibility that this parameter might be further fine-tuned to achieve even faster results.

566

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

578

579 In addition to further studies applying QTA to more challenging treatment problems, we would
580 also like to explore whether implementing QTA on a quantum computer could lead to greater
581 computational speedup. In their 2015 study, Nazareth and Spans reported on the first use of a
582 quantum annealing computer for IMRT beamlet weight optimization; they found that while SA
583 consistently produced higher-quality plans, optimizations performed on a quantum annealing
584 device (using a modified version of Tabu Search as the optimization algorithm) were >2.5 times
585 faster than SA.¹⁹ At the time of their study, the researchers were limited to a 512 qubit device,
586 which restricted the complexity of the treatment problems they could tackle. For reference, if the
587 beamlets in Case 1 were discretized using the same method used by Nazareth and Spans, 5,376

588 qubits would be required. In early 2019, the development of a 5,000 qubit commercial quantum
589 annealing computer was announced,³⁵ which would better allow QA to be scaled to higher
590 variable optimization problems but practical clinical application remains a subject for future
591 studies. We believe QTA would be an exciting candidate for quantum computing because it has
592 already shown promise over SA when run on a classical computer.

593

594 **5. Conclusions**

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

606

607 **References**

- 608 1. Bortfeld T. IMRT: a review and preview. *Phys Med Biol.* 2006;51(13):R363-379.
- 609 2. Citrin DE. Recent Developments in Radiotherapy. *N Engl J Med.* 2017;377(22):2200-
610 2201.
- 611 3. Cho B. Intensity-modulated radiation therapy: a review with a physics perspective.
612 *Radiat Oncol J.* 2018;36(1):1-10.
- 613 4. Bortfeld T, Schmidt-Ullrich R, Neve W, Wazer DE, SpringerLink. Image-Guided IMRT.
614 Berlin, Heidelberg: Springer Berlin Heidelberg,; 2006.
- 615 5. Zhang Y, Merritt M. Dose-volume-based IMRT fluence optimization: A fast least-
616 squares approach with differentiability. *Linear Algebra and its Applications.*
617 2008;428(5):1365-1387.

- 618 6. Lim-Reinders S, Keller BM, Al-Ward S, Sahgal A, Kim A. Online Adaptive Radiation
619 Therapy. *Int J Radiat Oncol Biol Phys*. 2017;99(4):994-1003.
- 620 7. Preskill J. Quantum computing and the entanglement frontier. *arXiv preprint*
621 *arXiv:12035813*. 2012.
- 622 8. Nielsen MA, Chuang IL. Quantum Computation and Quantum Information: 10th
623 Anniversary Edition. Cambridge: Cambridge University Press; 2010:
624 <http://esis.pace.edu/ctappert/cs837-18spring/QC-textbook.pdf>. Accessed 1/29/2019.
- 625 9. Grover LK. From Schrödinger's equation to the quantum search algorithm. *American*
626 *Journal of Physics*. 2001;69(7):769-777.
- 627 10. D-Wave Systems Inc. D-Wave Announces D-Wave 2000Q Quantum Computer and First
628 System Order. 2017, January 24; [https://www.dwavesys.com/press-releases/d-](https://www.dwavesys.com/press-releases/d-wave%20A0announces%C2%A0d-wave-2000q-quantum-computer-and-first-system-order)
629 [wave%20A0announces%C2%A0d-wave-2000q-quantum-computer-and-first-system-](https://www.dwavesys.com/press-releases/d-wave%20A0announces%C2%A0d-wave-2000q-quantum-computer-and-first-system-order)
630 [order](https://www.dwavesys.com/press-releases/d-wave%20A0announces%C2%A0d-wave-2000q-quantum-computer-and-first-system-order).
- 631 11. IBM. Quantum devices and simulators. [https://www.research.ibm.com/ibm-](https://www.research.ibm.com/ibm-q/technology/devices/)
632 [q/technology/devices/](https://www.research.ibm.com/ibm-q/technology/devices/).
- 633 12. IonQ Inc. A true quantum leap. 2018; <https://ionq.co/>.
- 634 13. Intel. Reinventing data processing with quantum computing.
635 <https://www.intel.com/content/www/us/en/research/quantum-computing.html>.
- 636 14. rigetti. QPU Specifications. 2019; <https://rigetti.com/qpu>.
- 637 15. Google AI. A Preview of Bristlecone, Google's New Quantum Processor. 2018, March 5;
638 <https://ai.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html>.
- 639 16. Apolloni B, Carvalho C, de Falco D. Quantum stochastic optimization. *Stochastic*
640 *Processes and their Applications*. 1989;33(2):233-244.
- 641 17. Farhi E, Goldstone J, Gutmann S, Lapan J, Lundgren A, Preda D. A quantum adiabatic
642 evolution algorithm applied to random instances of an NP-complete problem. *Science*.
643 2001;292(5516):472-475.
- 644 18. Mukherjee S, Chakrabarti BK. Multivariable optimization: Quantum annealing and
645 computation. *The European Physical Journal Special Topics*. 2015;224(1):17-24.
- 646 19. Nazareth DP, Spaans JD. First application of quantum annealing to IMRT beamlet
647 intensity optimization. *Phys Med Biol*. 2015;60(10):4137-4148.

- 648 20. Santoro GE, Martoňák R, Tosatti E, Car R. Theory of quantum annealing of an Ising spin
649 glass. *Science*. 2002;295(5564):2427-2430.
- 650 21. Morita S, Nishimori H. Mathematical foundation of quantum annealing. *Journal of*
651 *Mathematical Physics*. 2008;49(12):125210.
- 652 22. Kadowaki T, Nishimori H. Quantum annealing in the transverse Ising model. *Physical*
653 *Review E*. 1998;58(5):5355.
- 654 23. Zettili N. *Quantum Mechanics : Concepts and Applications*. New York, UNITED
655 KINGDOM: John Wiley & Sons, Incorporated; 2009.
- 656 24. Kuech TF. Metal-organic vapor phase epitaxy of compound semiconductors. *Materials*
657 *Science Reports*. 1987;2(1):1-49.
- 658 25. Leys MR, Veenliet H. A study of the growth mechanism of epitaxial GaAs as grown by
659 the technique of metal organic vapour phase epitaxy. *Journal of Crystal*
660 *Growth*. 55(1):145-153.
- 661 26. Webb S. Optimization of Conformal Radiotherapy Dose Distributions by Simulated
662 Annealing. *Physics in Medicine and Biology*. 1989;34(10):1349-1370.
- 663 27. Bertsimas D, Tsitsiklis J. Simulated Annealing. *Statistical Science*. 1993;8(1):10-15.
- 664 28. Hajek B. Cooling Schedules for Optimal Annealing. *Mathematics of Operations*
665 *Research*. 1988;13(2):311-329.
- 666 29. Kessler ML, McShan DL, Epelman MA, et al. Costlets: A Generalized Approach to Cost
667 Functions for Automated Optimization of IMRT Treatment Plans. *Optimization and*
668 *Engineering*. 2005;6(4):421-448.
- 669 30. Matuszak MM, Larsen EW, Fraass BA. Reduction of IMRT beam complexity through
670 the use of beam modulation penalties in the objective function. *Medical Physics*.
671 2007;34(2):507-520.
- 672 31. MacFarlane M, Hoover DA, Wong E, Goldman P, Battista JJ, Chen JZ. A fast inverse
673 direct aperture optimization algorithm for intensity-modulated radiation therapy. *Medical*
674 *Physics*. 2019;46(3):1127-1139.
- 675 32. Härdenmark B, Liander A, Rehbinder H, Löf J, Robinson D. P3IMRT: Direct machine
676 parameter optimization. *Pinnacle White Paper*. 2004(4535):983.
- 677 33. Luo Y, McShan DL, Matuszak MM, et al. A multiobjective Bayesian networks approach
678 for joint prediction of tumor local control and radiation pneumonitis in nonsmall-cell lung

- 679 cancer (NSCLC) for response-adapted radiotherapy. *Medical Physics*. 2018;45(8):3980-
680 3995.
- 681 34. Tseng H-H, Luo Y, Cui S, Chien J-T, Ten Haken RK, Naqa IE. Deep reinforcement
682 learning for automated radiation adaptation in lung cancer. *Medical Physics*.
683 2017;44(12):6690-6705.
- 684 35. D-Wave Systems Inc. D-Wave Previews Next-Generation Quantum Computing Platform.
685 2019, February; [https://www.dwavesys.com/press-releases/d-wave-previews-next-](https://www.dwavesys.com/press-releases/d-wave-previews-next-generation-quantum-computing-platform)
686 [generation-quantum-computing-platform](https://www.dwavesys.com/press-releases/d-wave-previews-next-generation-quantum-computing-platform).

688 Acknowledgements

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

694 Conflict of Interest Statement

695 The authors have no conflicts of interest to report.

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

700 **Figure 2.** (a) Barrier width rate extrapolated from metal organic chemical vapor deposition
701 (MOCVD) studies by Leys and Veenvliet.²⁵ (b) Barrier width function calculated via numerical
702 integration of (a). (c) Additional width functions explored in this study.

704 **Figure 3.** CT scans show contours for structures optimized for Case 1 (a) and Case 2 (b),
705 respectively. Case 1 features a PTV that is roughly spherical in shape and far from major organs.
706 (with the exception of the liver). Case 2 features a PTV with convex geometrical features and
707 close proximity to both the liver and the stomach.

708
709 PTV: planning target volume

710

711 **Figure 4.** Process of calculating the convergence iteration number from a representative QTA
712 optimization. The energy gradient (middle) is calculated from the saved energy history (right).
713 From this gradient a MAM of width 100 was calculated (left). The black vertical line is plotted at
714 the maximum iteration number j for $\max_j(|\text{MAM}(j)| > c^{\text{tol}})$.

715

716 MAM: moving average mean; QTA: quantum tunnel annealing

717

718 **Figure 5.** Quantum Tunnel Annealing (QTA) algorithm for intensity modulated radiation
719 therapy (IMRT) optimization.

720

721 **Figure 6.** Optimization results for QTA and SA applied to Case 1. Figures 6(a) and 6(b) display
722 DVH curves for 10 separate optimizations using QTA and SA, respectively. Figures 6(c) and
723 6(d) display the PE trajectories for the 10 QTA and SA optimizations. Figures 6(e) and 6(f)
724 display representative dose distributions calculated in Eclipse using fluence values from the 10th
725 QTA and SA optimization.

726

727 DVH: dose volume histogram; PE: potential energy; QTA: quantum tunnel annealing; SA:
728 simulated annealing

729

730 **Figure 7.** (a) displays the fluence map results for a single beam in Case 2 resulting from the
731 QTA optimization without refined smoothing. (b) displays the fluence map results from QTA
732 optimization with refined smoothing.

733

734 QTA: quantum tunnel annealing;

735

736 **Figure 8.** (a) displays the annealing schedule functions tested for QTA and SA. Note that T4 was
737 the annealing schedule used for all remaining studies. (b) displays box and whisker plots of the
738 convergence results for QTA and SA, respectively, for each annealing schedule.

739

740 QTA: quantum tunnel annealing; SA: simulated annealing

741

742 **Figure 9.** DVH bands, PE trajectories, and representative dose distributions for stochastic
743 optimizations ($N = 10$, 500000 iterations) with the initial beamlet-weight vector set to random
744 values uniformly distributed between 0 and 22 on a challenging SBRT liver case for QTA ((a),
745 (c), and (e)) and SA ((b), (d), and (f)), respectively.

746

747 DVH: dose volume histogram; PE: potential energy; QTA: quantum tunnel annealing; SA:
748 simulated annealing; SBRT: stereotactic body radiation therapy;

749

750 **Figure 10.** DVH bands, PE trajectories, and representative dose distributions for stochastic
751 optimizations ($N = 10$, 250000 iterations) of aperture weights on a challenging SBRT liver case
752 for QTA ((a), (c), and (e)) and SA ((b), (d), and (f)), respectively.

753

754 DVH: dose volume histogram; PE: potential energy; QTA: quantum tunnel annealing; SA:
755 simulated annealing; SBRT: stereotactic body radiation therapy;

Author Manuscript

Table 1. DVH constraints applied to objective function for Case 1.

Case 1 DVH Constraints					
Structure	Constraint Type	Limit	Volume (%)	Dose [Gy]	Penalty
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 target volume; Gy: Gray; PTV: planning target volume

Table 2. DVH constraints applied to objective function for Case 2.

Case 2 DVH Constraints					
Structure	Constraint Type	Limit	Volume (%)	Dose [Gy]	Penalty
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

Author Manuscript

Table 3. Convergence times (in seconds) for QTA with different barrier width schedules as well as SA.

Case 2 Convergence Results QTA Barrier Width Testing	
Algorithm type	Convergence (s)
SA	1062.5
QTA, $w^a, w' = 1 \times 10^{-5}$	528.6
QTA, $w^a, w' = 1 \times 10^{-7}$	637.2
QTA, $w^a, w' = 1 \times 10^{-9}$	1762
QTA, MOCVD	874.2

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

Table 4. Mean convergence times (in seconds) for QTA and SA with perturbations to the initial beamlet-weight values.

Case 2 Convergence Results for Initial Beamlet-Weight Testing			
b_0	QTA (s)	SA (s)	SA-QTA hybrid (s)
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

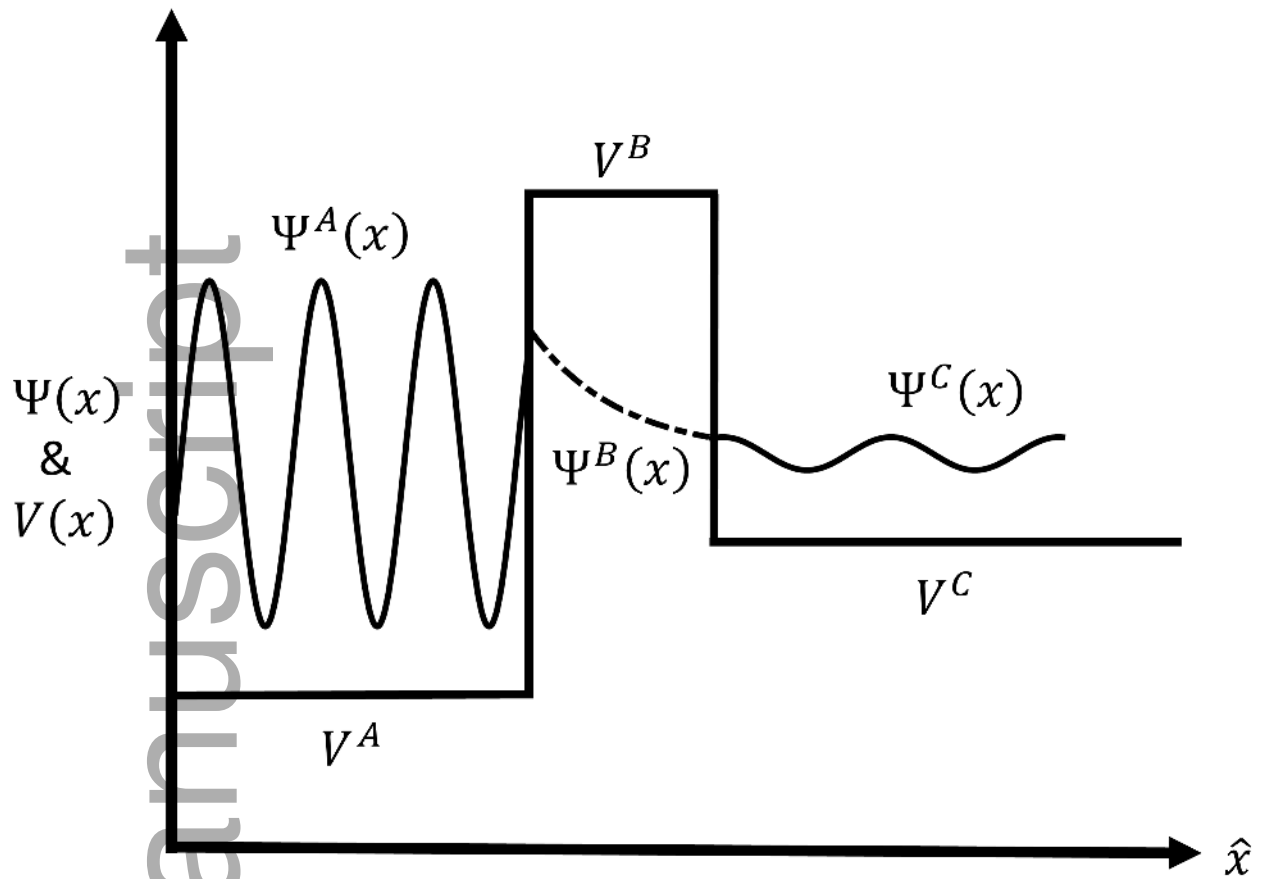
Table 5. Parameter changes and convergence times (in seconds) for QTA and SA with perturbations to the original dose constraints.

Case 2 Convergence Results for Parameter Testing					
Parameter	Organ	from	to	QTA (s)	SA (s)
DVH	Stomach	Max dose = 28Gy	Max dose = 18Gy	788.8	1038.9
DVH	Liver	Mean dose = 4Gy	Mean dose = 2Gy	951.6	1108.2

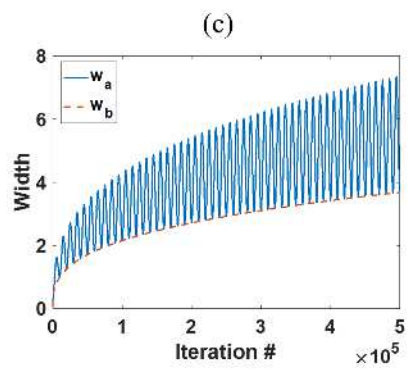

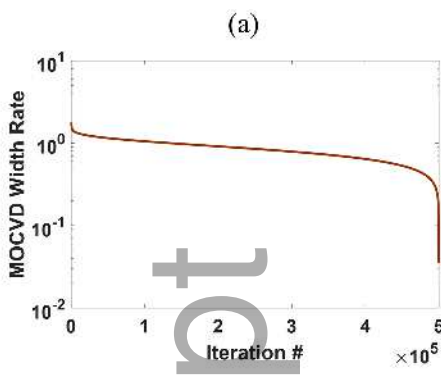
DVH and Dose	PTV	Target dose = 33Gy	Target Dose = 43Gy	701.2	1028.6
Original	Original	N/A	N/A	661.8	1083.3

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


Author Manuscript




mp_13840_f1.tif



mp_13840_f2.tif




(a)



(b)

mp_13840_f3.tif

Author Manuscript



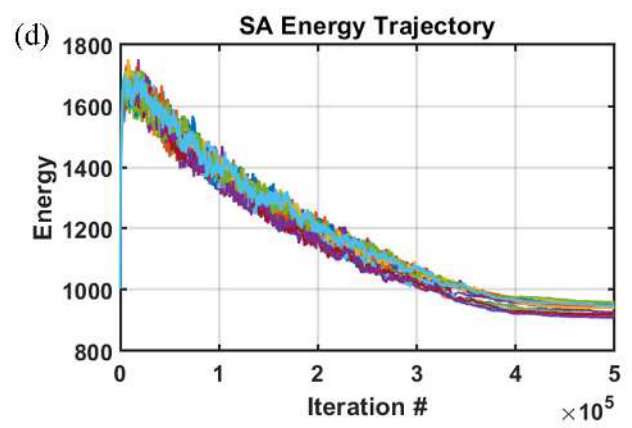
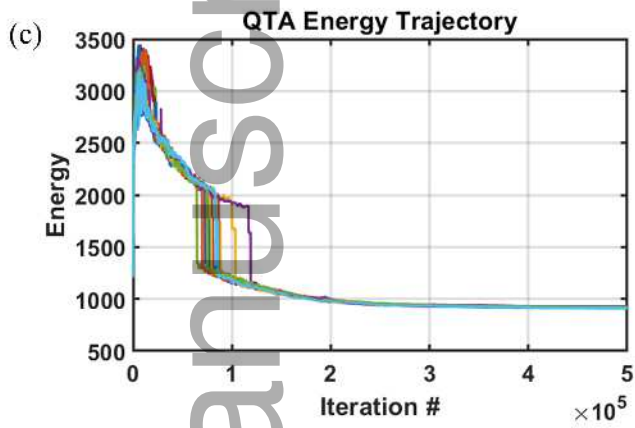




mp_13840_f4.tif


Author Manuscript

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  $\exp\left(\frac{-w\sqrt{E(b_n)-E(b_{n-1})}}{\mathcal{T}}\right) > rand$  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
```



mp_13840_f6.tif

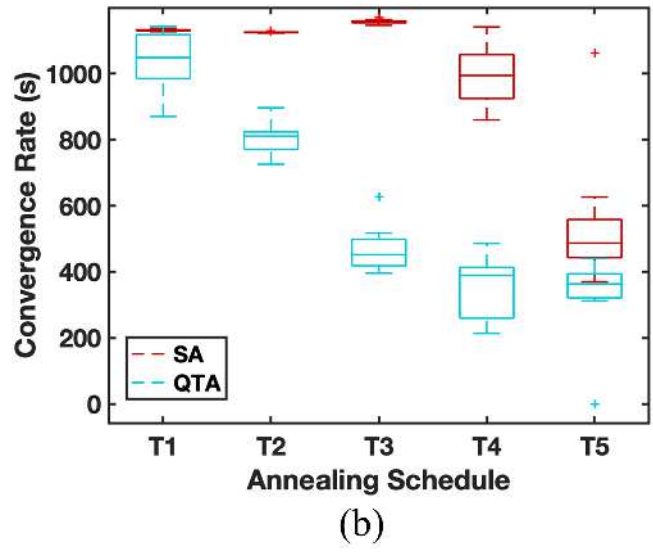
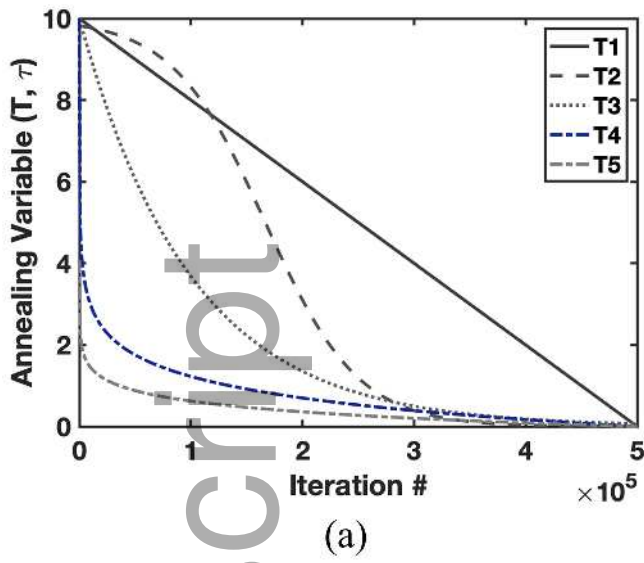


(a)




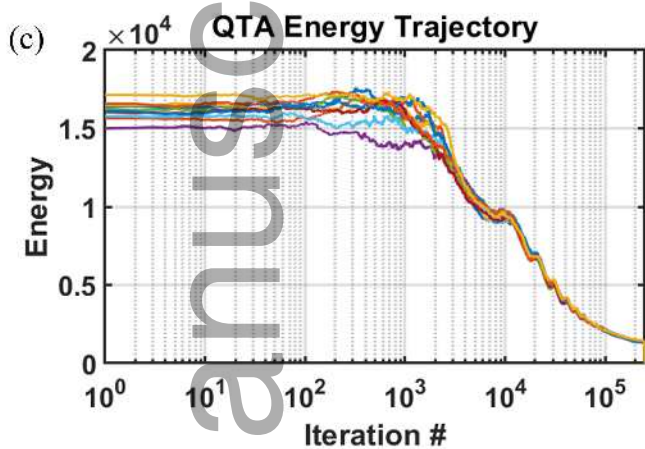


(b)

mp_13840_f7.tif



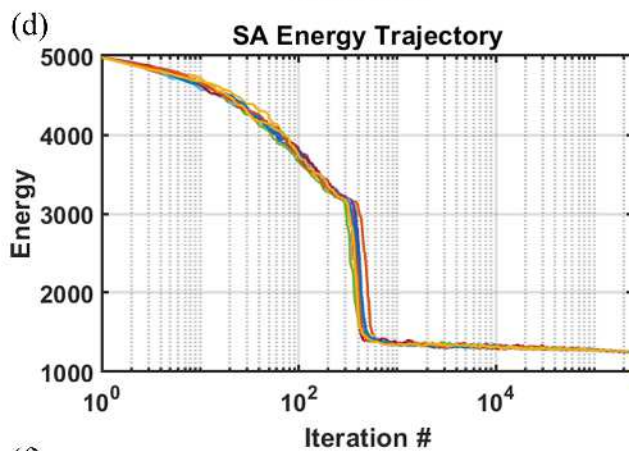

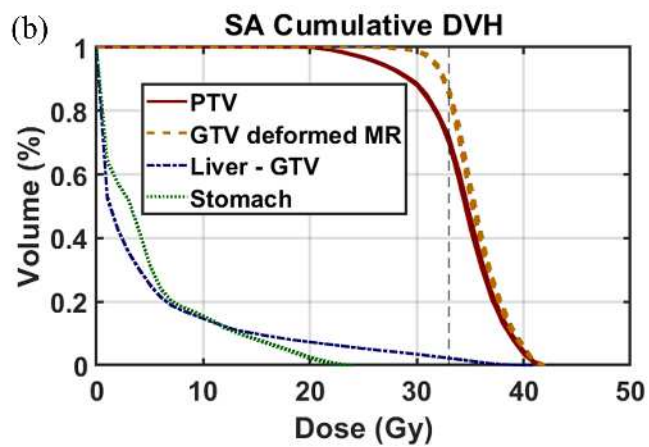

Author Manuscript



mp_13840_f8.tif



mp_13840_f9.tif



mp_13840_f10.tif