

Brachytherapy Implant Evaluation Using New Isodose Surface Metrics

Alexandra Christensen

A senior thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of
Bachelor of Science

J. Adam Cunha, Research Advisor
John S. Colton, Departmental Advisor

Department of Physics and Astronomy
Brigham Young University

Copyright © 2020 Alexandra Christensen

All Rights Reserved

ABSTRACT

Brachytherapy Implant Evaluation Using New Isodose Surface Metrics

Alexandra Christensen
Department of Physics and Astronomy, BYU
Bachelor of Science

Brachytherapy is the use of ionizing radiation at short distances to treat disease, delivering the appropriate dose via an implant through which radionuclides can travel in or on the body. The geometry of implant placement is critical to achieving brachytherapy plans that do not overdose the target region. Implants are typically designed based on the radiation oncologist's experience, and although some groups have created parameters for quantitative implant comparison, most lack simplicity and objectivity. This project describes an investigation of alternative indices that can be used for brachytherapy implant evaluation. The surface area to volume ratio of isodose surfaces and the number of non-contiguous isodose surfaces were calculated as functions of dose contour value. Several simple geometric target volumes were created to test the ability of these new metrics to differentiate between implant configurations. The surface area to volume ratio demonstrated the highest discernment power. A higher surface area to volume ratio correlates with the configuration expected to represent a more optimal implant. This is particularly useful when comparing implants with similar target coverage and can be used to evaluate implants with geometries for which experience is limited. This approach is sensitive to differences in implant geometry and has the simplicity of being a single number that can be used for comparison or as a parameter in optimization.

Keywords: radiotherapy, brachytherapy, dose-volume histogram, isodose, dose grid, implant

ACKNOWLEDGMENTS

I would like to thank my advisor Dr. Cunha for believing in my potential as a scientist and for introducing me to the world of medical physics, as well as Dr. Hsu for showing great enthusiasm for this project and always helping to push it forward. I would also like to thank Dr. Colton for lending his time to help me navigate the physics program and to help me improve this thesis. Finally, I would like to thank my husband John, who has given me constant love and support as I developed my passion for medical physics.

Contents

Table of Contents	iv
List of Figures	1
List of Tables	1
1 Introduction	3
1.1 Brachytherapy Overview	3
1.2 Implants in HDR Brachytherapy	4
1.3 The Dose-Volume Histogram	5
1.3.1 Standard Form	5
1.3.2 Anderson's Natural DVH	7
1.4 Isodose surface observations	8
1.5 Project Overview	9
2 Computational Methods	10
2.1 Surface Area to Volume Ratio	10
2.1.1 Surface Area Calculation	11
2.1.2 Volume Calculation	13
2.1.3 Bubble Index	15
2.2 Direct Structure Count	15
2.3 Validation	16
2.4 Testing	19
3 Results and Conclusions	23
3.1 Outcome of Simulations	23
3.2 Directions for Future work	28
3.3 Conclusions	29
Bibliography	31
Index	33

List of Figures

1.1	Dose-Volume Histogram for a hypothetical target slice	6
1.2	Simple series of isodose surfaces at several dose contour values	9
2.1	Example of a voxel in a brachytherapy dose grid	11
2.2	Example of a voxel in a brachytherapy dose grid being intersected by an isodose surface	13
2.3	Diagram of the approximation used to calculate the volume of an isodose surface contained in one voxel	14
2.4	Plots of the surface area, volume, and Bubble Index of isodose surfaces in a trial dose grid. Algorithm outputs are compared to ideal outputs for validation	18
2.5	Algorithm output for the counting of individual isodose surfaces compared to 3D renders of isodose surfaces for validation	19
3.1	BI curves for the 1 cm cubic target	24
3.2	BI curves for the 2 cm × 1 cm × 1 cm rectangular prism target	24
3.3	BI curves for the triangular prism target	25
3.4	BI curves for the 5 cm cubic target	26
3.5	BI curves for the 3 mm × 3 mm × 4 mm rectangular prism target	27

List of Tables

2.1	Diagrams for all targets created and tested in this thesis	21
3.1	Results for the 1 cm cubic target	23
3.2	Results for the 2 cm × 1 cm × 1 cm rectangular prism target	24
3.3	Results for the triangular prism target	25
3.4	Results for the 5 cm cubic target	26
3.5	Results for the 3 mm × 3 mm × 4 mm rectangular prism target	26

Chapter 1

Introduction

1.1 Brachytherapy Overview

Radiation therapy (or radiotherapy) is the practice of using ionizing radiation to cure disease. The goal is to deliver a prescribed dose of ionizing radiation to a target disease while minimizing dose to healthy tissue. Dose is measured in "grays" (Gy) and describes the absorption of energy from ionizing radiation into matter. One gray is equal to the absorption of 1 joule by 1 kilogram of matter [1].

Brachytherapy is a branch of radiotherapy in which the dose of ionizing radiation is delivered near or inside the patient [1] [2]. This is in contrast to external beam radiation therapy, where beams of particles are delivered to the patient from a clinical linear accelerator. In brachytherapy, the radiation dose is delivered via steel capsules containing a decaying radionuclide typically a few cubic millimeters in size. These capsules are referred to as "sources." For high dose rate (HDR) brachytherapy, the only therapy with which this study is concerned, the sources are delivered robotically at predetermined locations and for designated times, and are removed after the procedure [2].

A general outline of HDR brachytherapy treatment provides context for the new research presented in this thesis. First, a patient receives a diagnostic image displaying the anatomy surrounding the disease area. From this information, an implant arrangement is made by a radiation oncologist. The implant is the set of applicators and/or catheters which will be traversed by the brachytherapy source. A radiation oncologist then chooses an implant arrangement based on previous experience and successful implants for similar disease in the literature. For example, many oncologists place an implant based on achieving sufficient "peripheral loading," meaning that catheters are biased toward the edge of the disease area to compensate for increased dose in the center due to the overlapping contribution from each catheter. After the implant is physically placed, the patient is imaged again, and the both the implant and anatomical structures are digitized in treatment planning software. Typically, optimization algorithms are used to calculate the appropriate "dwell positions" and "dwell times," the locations and of brachytherapy sources and the duration they will stay in these locations in order to achieve sufficient coverage of the disease target while sparing healthy tissue.

Once the treatment planning software calculates the optimal dwell positions and times, a 3D scalar dose grid can be constructed. The dose grid shows the dose that the patient will receive at any point in space. This is the main source of information that this study uses in order to create a new way to evaluate HDR brachytherapy implants. After optimization is completed and dose is calculated, the patient can be treated by the physical insertion of brachytherapy sources.

1.2 Implants in HDR Brachytherapy

The HDR brachytherapy implant is of particular interest in this study. Because the implant limits the available space for dwell positions, plan efficiency is dependent on the implant placement. Efficiency in this case is the implant's ability to provide the desired dose to the designated target without "wasting" dose by creating excessively-dosed areas, or "hotspots." While it may seem

effective to overdose a disease area, hotspots increase the chances of secondary diseases due to brachytherapy; thus dose homogeneity is desirable [3] [4].

Implants are typically placed either through standardized applicators or by interstitial needles [2]. The decision of which applicator to use or precisely where to place interstitial needles is made by the operating radiation oncologist and is primarily based on experience of both the doctor herself and the collective experience of brachytherapy professionals [5]. While certain implant configurations do correlate with positive outcomes, the fact that they rely on the oncologist's experience means that there may be more advantageous configurations that are not being considered.

Various studies have attempted to evaluate implant efficiency quantitatively, but what is desired as the result of this study is a parameter simple enough to be inserted into an optimization algorithm. Many previous attempts have fallen short of this goal either by being too complicated or relying on subjective parameters [4]. If implant configurations could be determined via optimization instead of experience, new implant configurations could be developed and there could be a higher degree of customization of individual patients' implants.

1.3 The Dose-Volume Histogram

1.3.1 Standard Form

The dose-volume histogram (DVH) is a graphical tool used to evaluate dose distribution in radiotherapy treatment [1]. A DVH is created by dividing an organ or a therapy target into unit volumes or "voxels" and sorting each voxel into a bin according to the dose it receives. Brachytherapy plans are typically evaluated using the cumulative DVH (from here on abbreviated simply as DVH), which presents the total volume of an organ or target that receives a given minimum dose. In practice, several indices along the DVH curve are used for plan evaluation, such as the volume receiving 100% of the prescription dose (V_{100}). These indices simplify the DVH to focus on only one dose




Figure 1.1 (Left) An example of a dose-volume histogram for a two-dimensional slice of a brachytherapy target region, where each voxel is assigned a dose. (Right) The black dots represent the DVH for the Target Region. Axes are normalized to show the percentage of a prescription dose on the x axis and the percentage of total target volume on the y axis. The red line represents the ideal DVH, meaning that the every voxel in the target would receive a complete prescription dose and no more. In reality, such an ideal is difficult to achieve, and some part of the target will typically receive a higher dose than the prescription. Note that any given point on the DVH represents not the exact volume receiving a given dose, but the volume that receives *at least* that dose.

level at a time, which facilitates comparison between plans. An example of a DVH is shown in Fig. 1.1.

Although a DVH is an essential and informative part of radiotherapy, some information is lost in its creation. First, the DVH does not communicate any spatial information concerning dose distribution, because the DVH is concerned only with the total number of voxels at each dose [6]. Thus a plan in which small pockets of high dose are scattered throughout the target volume will appear similar to a plan in which there is one large mass receiving that dose. Differentiating between two such plans may be crucial to avoiding dose hotspots in some cases.

Another shortcoming of the DVH, which is relevant to brachytherapy in particular, is its inability to account for the inverse square effect. In brachytherapy, unlike in external beam therapy, dose falls off proportionally to

$$d = \frac{S}{r^2}, \quad (1.1)$$

where d is the dose at a point that is a distance r away from a radiation source with strength S . Thus

brachytherapy produces steep dose gradients in and near the target volume. The inverse square effect results in the highest doses being attributed to only a small proportion of voxels near the brachytherapy source, whereas lower doses are attributed to the much larger number of distant voxels as the inverse square function flattens with increasing r . When a plan is evaluated using a regular DVH, this effect obscures dose variations happening near radiation sources since the volume receiving high doses is often relatively small compared to the entire volume of interest. Yet, the structure of high dose regions are of interest as they describe the degree of excess dose and therefore overall implant efficiency.

Additionally, the dosimetric indices evaluated on a typical DVH, such as V_{100} , are not sufficient to holistically describe a brachytherapy plan since these indices select only a handful of dose values to investigate. Thus, the ability to describe transitions across dose levels is limited. To more effectively describe implant efficiency in a typical brachytherapy plan, a metric which can describe dose distribution but which accounts for the aforementioned failures of the standard DVH is needed.

1.3.2 Anderson's Natural DVH

To suppress inverse square law effects, Anderson proposed a "natural" dose volume histogram (NDVH) [7]. In an NDVH, voxels are sorted into bins according to the associated dose as in a regular DVH, but an algebraic normalization then takes place in order to adjust for distance from the radiation source, thus un-obscuring variations occurring at high doses. The algebraic normalization was constructed so that the NDVH curve of a single source is constant across all dose levels. Using this method, Anderson demonstrated the ability to assess the quality of idealized source configurations. While Anderson's NDVH does account for inverse square law effects, it still does not account for spatial distribution of dose.

In practice, the NDVH is difficult to use and cannot easily be translated to a simple parameter that can be inserted into an optimization algorithm [4]. A publication by Fournier-Bidoz et al. [8]

demonstrates that the use of the NDVH requires complicated analysis of the NDVH curve, resulting in multiple parameters, whose comparative importance is ambiguous.

Like Anderson's NDVH, the research presented in this study is an attempt to provide more information about implant quality than a standard DVH by seeking new parameters. Rather than relying solely on the total volume receiving a given dose, this study seeks information that lies in the structure of isodose surfaces in the dose grid.

1.4 Isodose surface observations

A 3D dose distribution can be visualized as a series of "isodose surfaces" outlined by dose points of equal value. Points in the dose grid inside an isodose surface must have dose values higher than the dose that defines the surface, because these points will be closer to brachytherapy sources. Isodose surfaces are analogous to topological contour lines on a geographic map which outline points of equal altitude in two dimensions; likewise, isodose surface contours outline points of equal dose value in three dimensions. As the defining dose value changes, the isodose surfaces morph in size and shape. The isodose surface defined by 100% of the prescription should roughly conform to a target volume in a successful plan, while higher dose isodose surfaces would shrink around the brachytherapy sources, and lower dose isodose surfaces spread outward from the target volume, again similar to how high altitude contour lines on a map would shrink along a mountain peak and spread out from a mountain base. Fig. 1.2 shows how this topological transition might appear for a plan with two sources, shown as black points. In this example, the single, contiguous isodose surface breaks into two, smaller isodose surfaces around each separate source at 250% of the prescribed dose (an arbitrary value in this case). We propose that a plan that undergoes this topological transition at a lower dose is a more efficient plan, since the separation of individual isodose surfaces limits the available volume for regions of undesirable high doses. Even if two




Figure 1.2 Isodose surface renderings around two brachytherapy radiation sources, shown as black points. Each contour is defined by points in space with equal ionizing dose value. High doses have smaller isodose surface contours that are closer to the sources.

implants result in plans that have similar DVHs, an investigation of isodose surface transitions will yield the more advantageous implant placement. Some measurement of isodose surface shape must be used to quantify these transitions through a dose range of interest. Such an indicator of implant quality would be significant, since the implant placement limits possible dwell positioning in planning and can determine the prevalence of dose hotspots.

1.5 Project Overview

This thesis presents the development and testing of an algorithm that evaluates implant efficiency using isodose surfaces. I begin by describing the new metrics used in this study to gain further information about implant efficiency. Sections 2.1.1 - 2.3 include a detailed description of the algorithm to calculate these metrics using a brachytherapy dose grid. The methods used to create implant simulations and test our metrics' ability to evaluate implant efficiency are found in 2.4. Section 3.1 describes the simulations made and their results. Directions for future work are included in Section 3.2. Conclusions concerning the usefulness of the new metrics used in this study are discussed in Section 3.3

Chapter 2

Computational Methods

2.1 Surface Area to Volume Ratio

The goal of this study is to create and test new metrics for implant evaluation based on isodose surface geometry. The "bubble index" (BI) is defined at a given isodose contour value as the ratio of the combined surface area of all isodose surfaces at that dose to the volume contained within them¹. BI is used in this study to evaluate implant efficiency because it is unaffected by the inverse square effect and because it has the ability to describe spatial dose distribution. Unlike surface area or volume separately, BI is indicative of the shape of topological features that define an isodose surface rather than the size alone. A higher BI correlates with isodose surfaces that are broken into or are about to be broken into multiple surfaces, thus indicating a plan in which hotspots are restricted. Because the restriction of hotspots is desirable, a high BI on average suggests a successful implant.

The C++ algorithm created for this study uses a three-dimensional 1 mm spaced dose grid produced from a brachytherapy plan after dwell positions and times are assigned. The dose grid

¹An earlier poster presentation defined bubble index as the maximum rate of change of the surface area to volume ratio of all isodose surfaces [9]. However, this concept did not prove useful; therefore, bubble index was redefined as the surface area to volume ratio of isodose surfaces as a function dose contour value, which is given in this thesis.




Figure 2.1 The Marching Cubes method requires the division of the dose grid into voxels such as the one depicted here [10]. The red and green numbers show how each vertex and edge is referenced in the algorithm code. Each red number is a point in a 3D dose grid associated with a dose value. If the voxel is intersected by an isosurface, the isosurface will cross at least three of the indexed edges.

contains only information about the location of dose in space and does not explicitly include anatomical structures or implant placement. In order to calculate the bubble index for a given dose, the surface area and volume must first be calculated individually.

2.1.1 Surface Area Calculation

The algorithm made for this study employs the Marching Cubes method to calculate surface area [10]. To use the Marching Cubes method, the dose grid is first divided into cubic volume voxels, each with eight vertices each corresponding to a specified point in the dose grid. Neighboring voxels share a face and four dose grid values. Each voxel is programmed as an object with eight attributed vertices each given an index 0-7. Given a voxel in the dose grid and a vertex index, the associated dose value is known. Each edge of a voxel is also referenced by an index 0-11. A voxel with indices labeled is shown in Fig. 2.1.

The surface area is calculated for a specified dose value that defines one or more 3D isodose

surfaces. The isodose surface or surfaces intersect some set of voxels in the dose grid. To find the intersected voxels, each vertex in the dose grid must be investigated. A vertex on a given voxel can be either less than the dose value that defines the isodose surfaces or equal to/greater than this value. Given this binary option for each of eight vertices, there are 256 possible states for any voxel, each state being a set of voxel indices below and a set of voxel indices equal to or greater than the given dose value. Each state corresponds to exactly one set of voxel edges through which the isodose surface intersects. In the most common case, the voxel is in a state in which every vertex is above or every vertex is below the given dose value, meaning that there is no isodose surface intersection in the voxel. In this case, the set of intersecting edges is empty, and the surface area contributed by that voxel is zero. Otherwise, these edges are arranged in groups of three to define triangles that make up the intersecting plane. All 256 edge sets corresponding to each possible voxel state are tabulated for easy reference [10].

Once the set of edges that are intersected by the isodose surface is known, the location P along the edge where the intersect occurs can be calculated using a linear interpolation:

$$P = P_1 + \frac{(D - d_1)(P_2 - P_1)}{d_2 - d_1}, \quad (2.1)$$

where D is the dose level of the isodose surface(s), P_1 and P_2 are the vertex indices at either end of the edge, and d_1 and d_2 are the dose values at P_1 and P_2 , respectively. An example of an intersection is shown in Fig. 2.2, with the intersecting piece of the isodose surface shown in blue. Performing the linear interpolation on each intersected edge results in a set of points along the edges of the voxel that define triangles in an intersecting plane. Each point can be given a vector to describe its location in respect to an origin arbitrarily defined as vertex 3. The area of each triangle can then be calculated as the surface area contribution for one voxel using the equation

$$\text{Area} = \frac{1}{2} |AB| |AC| \sin \theta, \quad (2.2)$$




Figure 2.2 A possible intersection of an isodose surface with a single voxel in the Marching Cubes method [10]. As in Fig. 2.1, the red numbers represent voxel indices and the green numbers represent edge indices. While this example demonstrates a triangular intersection, also possible are rectangular, trapezoidal, pentagonal, and hexagonal intersections.

where $|AB|$ and $|AC|$ represent the distances between vector A and vectors B and C , which point from the origin to the location of the isodose surface intersection, and θ is the angle between AB and AC .

In order to calculate the surface area of all isosurfaces for a given dose, the surface area contribution from each individual voxel must be summed over the entire dose grid. The process is repeated for a range of dose values to obtain the isodose surface areas for any given dose contour value for the plan.

2.1.2 Volume Calculation

Similar to the surface area calculation, volume was calculated by evaluating the volume contributed by each voxel and then summing every individual voxel in the dose grid.

As described above, there are 256 possible states for a given voxel, where each state is a set of vertex indices with doses below and a set of vertex indices at or above the specified dose value. In




Figure 2.3 Example of how the volume approximation might look for a particular voxel. The origin is defined as the near lower left corner, and x , y , and z directions defined as shown. In this particular example, three intersecting points are represented. However, if no intersecting point is given in a particular coordinate direction, and the associated prism has no volume. Also note that in this illustration, the origin is assumed to be inside the isosurface; in the case that it is not, the volume can be taken as the volume of the voxel minus the shaded volumes.

the case that all vertices have doses above the specified dose value, the voxel is completely inside the isodose surface and contributes its entire volume to the isodose surface volume. Similarly, in the case that all vertices have doses below the specified dose value, the voxel is completely outside the isodose surface and contributes no volume. In all other cases, the voxel is partially intersected by an isodose surface and contributes some fraction of its volume to the total isodose surface volume.

Unlike surface area intersections, which can be split into triangular sections for easy calculation, the volume contribution of each voxel is a unique three-dimensional polyhedron and is, therefore, more difficult to calculate generally. Thus, a method to approximate the volume contribution of a voxel was created for this work. First, the edges intersected by the isodose surface were found as described above. After defining the origin, three points that describe the intersecting point with the lowest value in each three-dimensional coordinate were selected. These points each define a

rectangular prism that must at minimum contribute to the total isodose surface volume. Care must be taken not to double-count regions of overlap. See Fig. 2.3, which shows how the minimum volume contribution is found. While this process is not as precise as the process for calculating the surface area contribution of each voxel, a minimum volume approximation is achieved.

2.1.3 Bubble Index

After surface area and volume are computed, BI is calculated as the surface area to volume ratio for a given dose level. Surface area, volume, and BI computations are repeated for a range of doses. The plot of BI as a function of dose contour value is then numerically integrated to produce the “bubble index area-under-curve” (BI-AUC). BI-AUC is easily used for comparison between different implants and has the advantage of summarizing isodose surface transitions for an entire dose range in a single number. BI-AUC in this study are given from 100%-150% of the prescription dose (Rx), because isodose surface transitions in this interval are critical to the formation of hotspots in the dose grid. A treatment plan in which isodose surfaces break apart at a low dose contour level would have higher BI throughout this dose range. Thus, such a plan would accumulate a higher BI-AUC than a plan in which isodose surface breaking occurs only at a greater dose contour level.

2.2 Direct Structure Count

In addition to BI and BI-AUC, the integer number of separate isodose surfaces at each dose level is of interest because it would allow for the direct observation of the point at which one large isosurface becomes multiple, smaller, disconnected isosurfaces. The isodose surface count is performed via a recursive function that investigates each voxel in the dose grid.

At a given dose level, all voxels can be marked as being either inside or outside of an isodose surface, depending on whether at least some of their vertices are above the defining dose value.

Thus any voxel that is at least partially intersected by an isodose surface is considered inside it for the purpose of counting isodose surfaces. The structure-counting function iterates through each voxel in the dose grid. When a voxel inside an isodose surface is found, a counter increments by one. The function then "marks" that voxel as well as neighboring voxels if they are inside an isosurface, and the same recursive function applies to these neighboring voxels as well. When another voxel is found that is inside an isosurface and unmarked, it must be inside another distinct structure—so, the counter again increments by one.

The number of distinct isosurfaces can be displayed as a function of dose contour value. A plot of this function is used to define a "burst dose" (BD) in which a single isodose surface in the entire dose grid splits into two or more isodose surfaces.

2.3 Validation

A dose grid simulating a single brachytherapy source at the origin was created to validate the algorithm for accuracy in calculating the surface area, volume, and BI of isodose surfaces. The dose grid consists of 49 x 49 x 49 dose points spaced by 1 mm, with the central point assigned to a source strength S of 1000 Gy m² and all other points given values calculated by Eq. (1.1) according to their distance r from the central point in order to simulate a single central brachytherapy dwell position. Since ideal single sources create perfectly spherical isodose surfaces, an ideal surface area (SA), volume (V), and BI can be calculated as a function of dose contour value d by using equations of surface area and volume for a sphere and manipulating Eq. (1.1) to substitute for r :

$$SA = 4\pi \frac{S}{d}, \quad (2.3)$$

$$V = \frac{4}{3}\pi \left(\frac{S}{d}\right)^{\frac{3}{2}}, \quad (2.4)$$

$$\text{BI} = \frac{\text{SA}}{V} = 3\sqrt{\frac{d}{S}}. \quad (2.5)$$

Values for surface area, volume, and BI returned by the algorithm described in this paper are compared to the ideal output for a single source in Fig. 2.4. The ideal values are represented by a blue line. The average difference between the output of the algorithm and the ideal output for BI over the dose range shown is 4%. This error is consistently due to overestimation, and is not random.

Several factors contribute to the inaccuracies in surface area and volume calculation. The interpolation implemented between dose points is linear, when ideally it would follow an inverse square form. As a result, the surface area calculation always yields a higher result than the actual value because isosurface intersections in individual voxels are calculated as slightly farther from sources than they would be in an inverse square interpolation. Since the dose grid does not contain information about the location of brachytherapy sources, this inaccuracy was not corrected for this report. Approximations utilized in volume calculation are described above, and contribute to a consistent underestimation in volume calculation. Since volume is in the denominator of BI, this underestimation leads to an overestimation of BI. Additionally, while the theoretical calculations in Eqs. (2.3), (2.4), and (2.5) assume perfectly spherical isosurfaces, the simulated dose grid has limited resolution of 1 mm and therefore could not produce ideal surface area, volume, or BI even if the algorithm was perfectly accurate.

The higher degree of accuracy obtained in volume calculation is likely due to the large "buffer" created by the voxels inside isodose surfaces that contribute their entire volume to the total calculation, whereas the surface area calculation used only voxels along the border of isodose surfaces. Since BI was calculated simply by dividing surface area by volume at each dose level, the error in BI is accounted for by the propagation of errors in surface area and volume each.

The bubble counting functionality was validated by generating 3D renders of isodose surfaces for

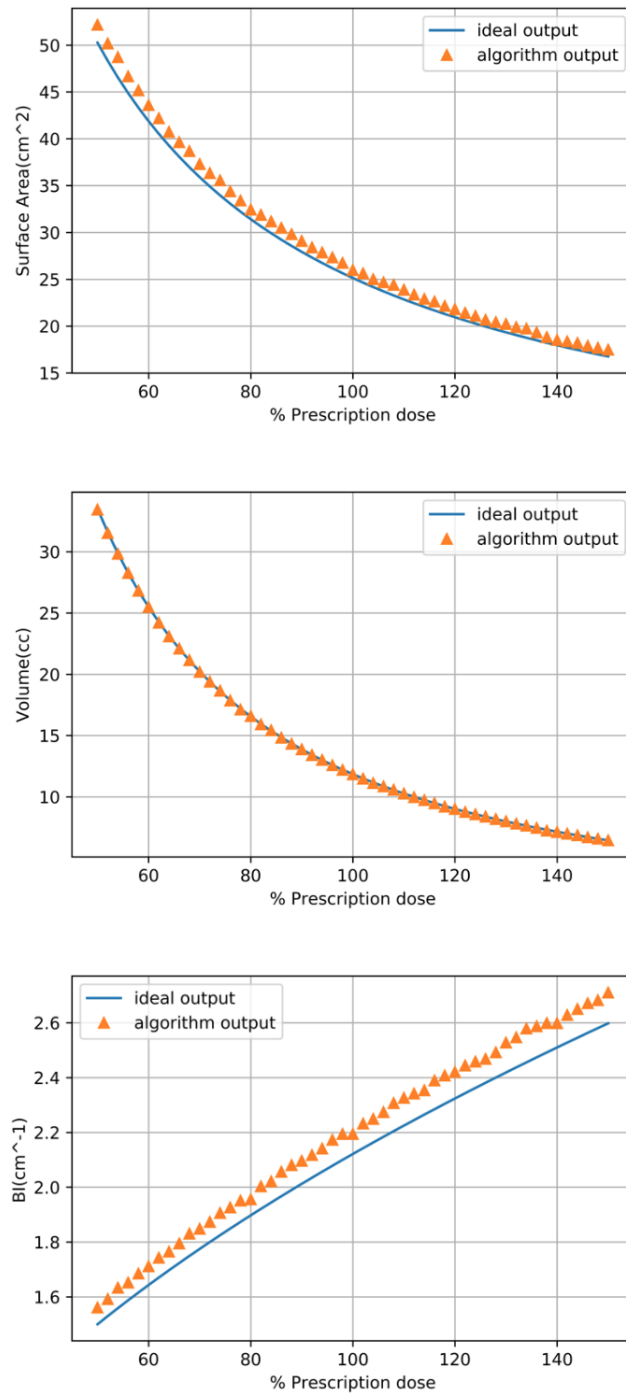


Figure 2.4 Surface area, volume, and bubble index of the single source simulation as a function of the percentage of prescription dose that defines the isodose surface contours. The source strength was set at the arbitrary value of 1000 Gy m², and the prescription dose was 5 Gy.

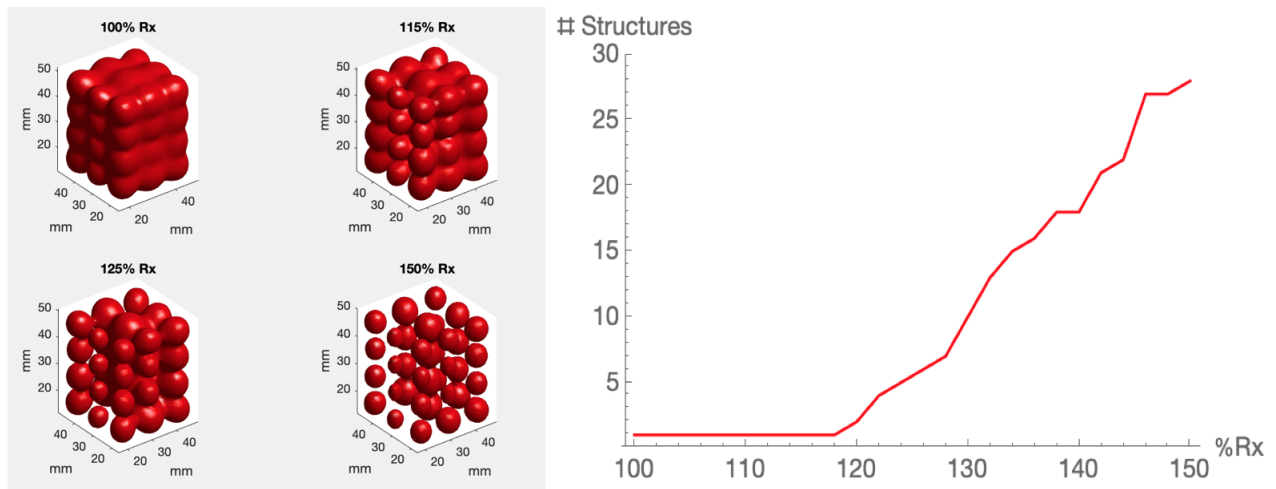


Figure 2.5 On the left, a 3D isodose surface rendering is shown for one of the test cases described below. The single isodose surface breaks apart between 115% and 125% of the prescription dose. On the right, the number of separate isodose structures found by the algorithm is plotted as a function of dose contour value.

situations with multiple brachytherapy sources and comparing these to the output of the algorithm. Figure 2.5 shows several 3D renders of isodose surfaces for the $3\text{ cm} \times 3\text{ cm} \times 4\text{ cm}$ target plan with 36 sources described in Section 2.4 compared to the bubble count computed by the algorithm. The plot designates the burst dose at 120% Rx, which agrees with the isodose surface renders. Additionally, the bubble count for the simulated single source yielded the expected constant one isodose structure across all dose levels.

2.4 Testing

In the clinic, a target is a designated region where a certain amount of dose is prescribed to treat a disease—typically, a tumor. However, clinical targets were not used for initial testing because the purpose of this project was to create a metric for implant efficiency that is otherwise difficult to determine; thus, simpler benchmark targets are needed to know whether the outputs of the algorithm were truly indicating quality implants. To test the algorithm’s ability to distinguish

between implants of varying efficiency, five simple, geometric targets were created. Although the most efficient implant may not be known beforehand for each geometric target studied, the simplicity of these tests allowed for insight into the implant configurations that correlate with high BI-AUC or low BD. These targets and the corresponding implants were simulated in Oncentra Treatment Planning (Elekta, Veenendaal, The Netherlands) and are shown in Table 2.1. In each image, the value of s is the only parameter that changes between different implant configurations made for that target.

The first target was a $1\text{ cm} \times 1\text{ cm} \times 1\text{ cm}$ cube. Three different implant configurations were created: one source at the center of the target, one source 2.5 mm laterally removed from the center, and one source 4 mm laterally removed from the center. These three implant configurations were compared to one another using the output from the algorithm. Because of the simplicity of this structure, the most symmetrical implant is most efficient.

The second target was a $2\text{ cm} \times 1\text{ cm} \times 1\text{ cm}$ rectangular prism. Two sources were placed along the long axis either 5 mm, 10 mm, or 15 mm apart. This target can be thought of as two of the first target adjacent to one another. The implant separated by 10 mm places each source in the center of its cube. This configuration should be more efficient than the implant separated by 5 mm, which is too central. However, the implant separated by 15 mm further provides peripheral coverage and may be the most efficient.

The third target was an equilateral triangular prism with 2 cm sides and a thickness of 1 cm. Lines were drawn connecting each corner to the center of the triangular face, and one source was placed in the center of each line. Three additional implant configurations were created by moving these sources 2 mm closer to one another, 2 mm farther from one another, or 4 mm farther from one another than in the original configuration.

The fourth target was a $5\text{ cm} \times 5\text{ cm} \times 5\text{ cm}$ cube. To increase in size and complexity, this implant was composed of four parallel catheters each containing five sources. The catheters were

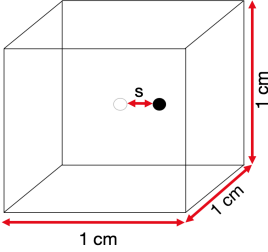
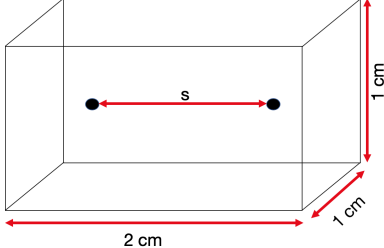
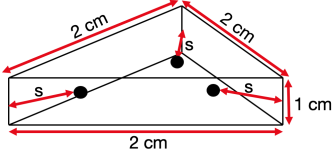
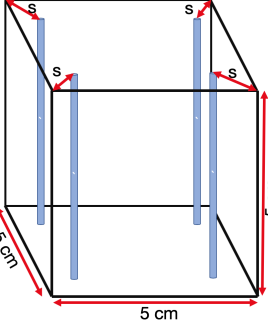
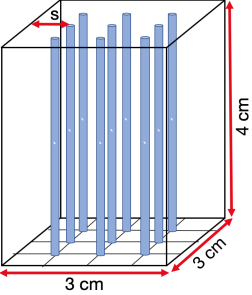
Test	Target
1	
2	
3	
4	
5	

Table 2.1 Diagrams for each target created and tested in this project. The black outline frames the geometric target, while black dots represent brachytherapy sources and blue cylinders represent catheters that may contain multiple sources. The parameter s in each image was varied to create several different catheter configurations for each target.

placed perpendicular to one of the square faces. Implant configurations varied in distance of each catheter from the nearest corner of the square face.

The fifth target was a $3\text{ cm} \times 3\text{ cm} \times 4\text{ cm}$ rectangular prism. Nine parallel catheters each containing four sources were placed along the long axis. One catheter was placed in the center of the square face, and the other eight were placed symmetrically surrounding it. The outer catheters were placed either 3 mm or 7 mm from the nearest target edge. The more central plan is expected to be less efficient because it should have difficulty achieving peripheral coverage.

Once each target and implant configuration was simulated, Inverse Planning Simulated Annealing (IPSA), a dose optimization tool, was used to calculate dwell timing. For each target simulation, V_{100} was kept constant within 1% target volume by adjusting dwell timing across changing catheter configurations to test whether implant efficiency could be compared in plans with similar DVHs. 3D dose grids were then extracted for input into the algorithm so that BI, BI-AUC, and BD could be calculated and compared.

Chapter 3

Results and Conclusions

3.1 Outcome of Simulations

The results presented in the following tables and figures indicate that BI-AUC is capable of choosing an advantageous implant configuration. Tables 3.1-3.5 display the BI-AUC from 100% to 150% Rx and BD in % Rx of each target simulation created for this project with its corresponding implant variations. The "Description" column designates a value for s , a distance parameter that specifies implant placement in the context of each individual target simulation. The target images in Table

Implant Description	V_{100}	BI-AUC	BD
$s = 0$ mm	70.2%	302.9	none
$s = 2.5$ mm	71.1%	279.9	none
$s = 4$ mm	70.6%	241.1	none

Table 3.1 Results for the simple cubic target volume with one brachytherapy source placed in the center. Different implants are made by moving the source laterally from the center. The implant with the source has the best BI-AUC. There is no BD for any implant, as expected.

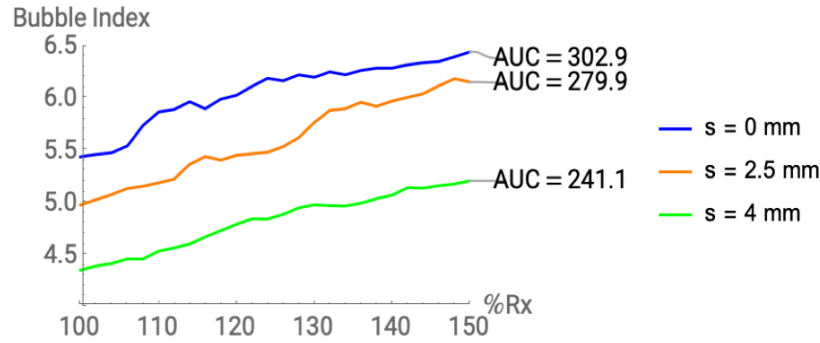


Figure 3.1 BI versus prescription dose for implants in the 1 cm cubic target. The most centralized implant has the highest BI over the entire dose range, and thus has the highest BI-AUC.

Implant Description	V_{100}	BI-AUC	BD
s = 5 mm	89.8%	183.0	none
s = 10 mm	89.8%	197.2	380%
s = 15 mm	89.7%	250.7	148%

Table 3.2 Results for the 2 cm × 1 cm × 1 cm rectangular prism target with two brachytherapy sources. Different implants are made by changing the lateral separation distance of the sources.

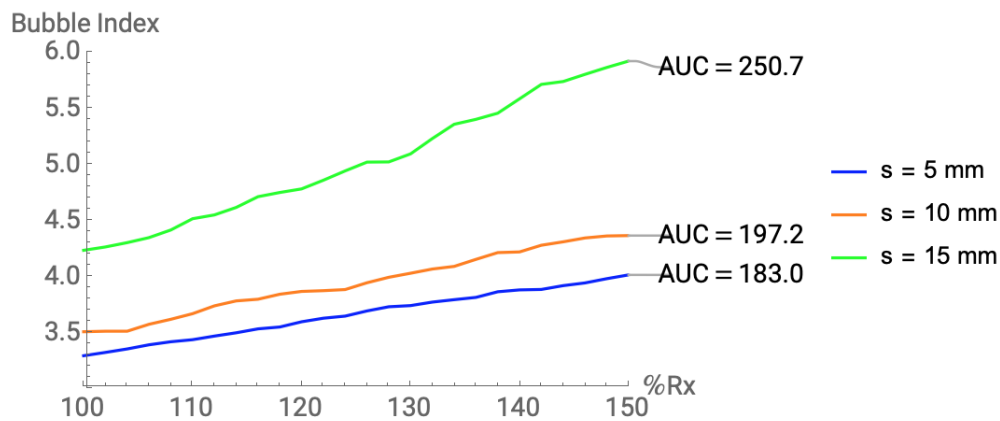


Figure 3.2 BI versus prescription dose for implants in the 2 cm × 1 cm × 1 cm rectangular prism target. The implant with the most separation between dwell positions achieved the highest BI over the entire dose range by a decent margin.

Implant Description	V_{100}	BI-AUC	BD
$s = 7.5$ mm	90.8%	177.6	185%
$s = 9.5$ mm	90.8%	200.4	206%
$s = 11.5$ mm	90.8%	193.9	480%
$s = 13.5$ mm	90.7%	176.4	none

Table 3.3 Results for the triangular prism target with three brachytherapy sources. Different implants are created by changing the distance of each source from the nearest corner of the triangular face. While the least centralized ($s = 7.5$ mm) implant has the lowest BD, the second-least centralized implant ($s = 9.5$ mm) has the greatest BI-AUC.




Figure 3.3 BI versus prescription dose for implants in the triangular prism target. While the $s = 9.5$ mm target has the highest BI-AUC, at the lowest doses it does not have the highest BI of all the implants. No implant has the highest or the lowest BI over the entire dose range.

Implant Description	V_{100}	BI-AUC	BD
$s = \frac{\sqrt{2}}{2}$ mm	89.5%	83.1	127%
$s = \sqrt{2}$ mm	89.1%	85.2	141%
$s = \frac{3\sqrt{2}}{2}$ mm	89.1%	56.5	250%
$s = 2\sqrt{2}$ mm	89.9%	53.5	748%

Table 3.4 Results for the 5 cm cubic target volume with four catheters each containing five brachytherapy sources. Different implants were created by changing the distance of each catheter from the nearest corner. While the least centralized ($s = \frac{\sqrt{2}}{2}$ mm) implant has the lowest BD, the second-least centralized implant ($s = \sqrt{2}$ mm) has the greatest BI-AUC.

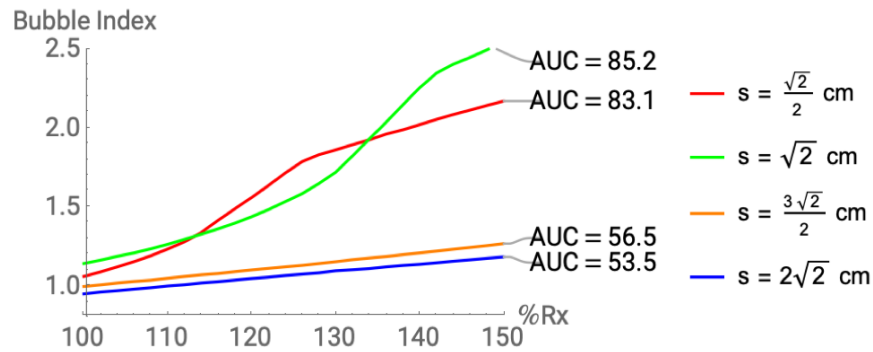


Figure 3.4 BI versus prescription dose for implants in the 5 cm cubic target. While the two most centralized implants consistently result in the lowest BI, the BI curves of the two other implants cross over one another over the given dose range.

Implant Description	V_{100}	BI-AUC	BD
$s = 3$ mm	91.9%	228.2	120%
$s = 7$ mm	91.0%	108.0	none

Table 3.5 Results for the 3 mm \times 3 mm \times 4 mm rectangular prism target with nine catheters each containing four brachytherapy sources. Different implants were created by changing the distance of each catheter from the nearest edge, excepting the central catheter, which was kept in the middle of the square face.




Figure 3.5 BI versus prescription dose for implants in the 3 cm × 3 cm × 4 cm rectangular prism target. The implant with the most separation between dwell positions achieved the highest BI over the entire dose range by a decent margin, although for doses close to the prescription the BI values for both implants were similar.

2.1 show the appropriate context for s in each case.

For the 1 cm cubic target shown in Table 3.1, the highest BI-AUC correlated with the implant that was most radially symmetric and therefore best able to provide dose to the entire target without waste. For both rectangular prism targets, shown in Tables 3.2 and 3.5 the highest BI-AUC correlated with the implant with the highest degree of peripheral loading. However, the 5 cm cubic target in Table 3.4 and the equilateral triangle in Table 3.3 show that there may be a limit to how much peripheral loading is advantageous: in these cases, the highest BI-AUC correlated not with the outermost implant configurations, but those that were peripherally loaded without nearly approaching the perimeter of the target.

Plots of BI versus dose give a more holistic view of some of the implants. Each plot shows the BI curve for every implant variation in one of the simple geometric targets. The BI-AUC of each curve is also given. Figure 3.1 shows the BI curves for the implants in the 1 cm cube. The most radially symmetric implant has the highest BI over the entire dose range. For the 2 cm × 1 cm × 1 cm target shown in Figure 3.2, the implant with the most separation between dwell positions has the highest BI over the entire dose range. Similarly, the BI curves for the implants in the 3 cm ×

3 cm × 4 cm target plotted in Figure 3.5 show that the implant with the most separation between catheters has the highest BI for the entire dose range. However, this is not always the case, as is seen in Figs. 3.3 and 3.4. The BI curves for the implants in the triangular target are shown in in Fig. 3.3. While the $s = 9.5$ mm implant has the highest BI-AUC overall, the BI lines cross over one another at different points in the dose range. In Fig. 3.4, which shows the BI curves for implants in the 5 cm cubic target, the BI lines also cross one another at several points. Thus the most advantageous implant according to BI-AUC is affected by the chosen dose range.

The results for BD are less informative. Many test cases did not result in a burst dose within a dose range typically calculated by the treatment planning software (<800% Rx). In targets where a burst dose did consistently occur, the lowest burst dose correlated with the implant configuration with the most separation between catheters. This indicates that BD is more a measure of implant distance than of implant quality. Furthermore, BD is fundamentally limited as a metric for implant evaluation because it reports only one "event" in isodose transitions, whereas BI-AUC summarizes the entire transition over a specified dose range. Thus BD could be sensitive to minor anomalies in an implant, such as one catheter positioned farther away from others that causes an early isodose surface break.

3.2 Directions for Future work

Before functionality is further developed, attention should be given to decreasing the computation time of the algorithm presented in this project. While the simple test cases presented here produced relatively small dose grids (about 20 x 20 x 20 dose points), real anatomical cases will produce much larger dose grids (closer to 90 x 90 x 90 dose points). For these larger dose grids, the algorithm can spend as much as three minutes calculating BI and structure count for each dose level. Because this algorithm is only in its preliminary stages, code efficiency was not a priority. Therefore it is

expected that many improvements can be made.

While simple geometric targets provide insight into how these new metrics evaluate implants, it is of more relevant interest to test them on patient anatomy. For instance, the BI-AUC and BD could be calculated using catheters already programmed in a previously treated patient. Then, more catheters could be added and BI-AUC and BD could be calculated again and compared to the values for the original case. Because the second case would retain the original catheters but also include additional catheters, the new implant must be able to deliver dose to the target at least as efficiently as the original case. Such a test would apply these parameters to clinical work, and further show whether BI-AUC and BD can consistently choose the implant that is known to be more efficient.

Eventually, an improved version of the algorithm created for this study could be joined with optimization programs like IPSA so that implant creation and evaluation could be completed simultaneously. Rather than using these parameters to compare multiple human-made implants, an optimal implant could be then be created in treatment planning software.

3.3 Conclusions

Our results suggest that BI-AUC could eventually be used to evaluate brachytherapy implants before physical placement. If the radiation oncologist was working with an abnormal or unfamiliar target, she could simulate several possible implant configurations, and then use BI-AUC along with the DVH to decide which implant configuration would be most advantageous. If not clinically applied, BI-AUC could be similarly used in resident training. Furthermore, we aim to eventually merge BI-AUC with dose optimization software, so that the optimal implant can be created without human input.

The new metrics presented in this thesis have built upon past efforts to quantitatively evaluate implant efficiency and make up for inadequacies in the standard DVH. This project has shown that

in simple geometric targets, BI-AUC has the ability to differentiate between implants of varying efficiency.

Bibliography

- [1] F. M. Khan and J. P. Gibbons, *The physics of radiation therapy*, 5th ed. (Wolters Kluwer, Philadelphia, 2014), pp. 97, 309, 423–424.
- [2] Mayo Clinic, “Brachytherapy,” <https://www.mayoclinic.org/tests-procedures/brachytherapy/about/pac-20385159> (Accessed April 07, 2020).
- [3] A. Wu, K. Ulin, and E. S. Sternick, “A dose homogeneity index for evaluating ^{192}Ir interstitial breast implants,” *Medical Physics* **15**, 104–107 (1988).
- [4] D. Baltas, C. Kolotas, K. Geramani, R. F. Mould, G. Ioanndis, M. Kekchidi, and N. Zambiglou, “A conformal index (COIN) to evaluate implant quality and dose specification in brachytherapy,” *Int. J. Radiation Oncology Biol. Phys.* **40**, 515–524 (1998).
- [5] A. Viswanathan and B. Thomadsen, “American Brachytherapy consensus guidelines for locally advanced carcinoma of the cervix. Part I: General principles,” *Brachytherapy* **11**, 33–46 (2012).
- [6] R. E. Drzymala, R. Mohan, L. Brewster, J. Chu, M. Goitein, W. Harms, and M. Urie, “Dose-volume Histograms,” *International Journal of Radiation Oncology* **21**, 71 (1991).
- [7] L. L. Anderson, “A natural volume-dose histogram for brachytherapy,” *Medical Physics* **13**, 898–903 (1986).

-
- [8] N. Fournier-Bidoz, T. Inoue, T. Inoue, and T. Nose, “Use of the Anderson dose-volume histogram to evaluate interstitial implants,” *J. Brachytherapy Int.* **13**, 235–242 (1997).
- [9] D. Kim, J. A. Cunha, L. Beaulieu, and I. Hsu, “A topological analysis of brachytherapy dose distributions: the bubble index,” *Brachytherapy* **15**, S158 (2016).
- [10] P. Bourke, “Polygonising a scalar field,” <http://paulbourke.net/geometry/polygonise/> (Accessed August 08, 2019).

Index

brachytherapy procedure, 4, 5

dose grid, 4, 8, 10–13, 16, 17, 22, 28

dose-volume histogram (DVH)
 definition, 5, 6
 shortcomings, 6

efficiency, 4, 7, 8, 10, 20, 29

gray, 3

inverse square effect, 6, 16

linear interpolation, 12, 17

Marching Cubes, 11, 13

natural dose-volume histogram, 7

oncologist, 4, 5, 29

optimization, 4, 22, 29

peripheral loading, 4, 20, 22, 27

voxel, 5, 7, 11, 13