

6-2008

Optimal Treatments for Photodynamic Therapy

Allen G. Holder

Trinity University, aholder@trinity.edu

D L Lagostera

Follow this and additional works at: https://digitalcommons.trinity.edu/math_faculty



Part of the [Mathematics Commons](#)

Repository Citation

Holder, A., & L Lagostera, D. (2008). Optimal treatments for photodynamic therapy. *4OR*, 6(2), 167-182. doi:10.1007/s10288-007-0046-4

This Post-Print is brought to you for free and open access by the Mathematics Department at Digital Commons @ Trinity. It has been accepted for inclusion in Mathematics Faculty Research by an authorized administrator of Digital Commons @ Trinity. For more information, please contact jcostanz@trinity.edu.

Optimal Treatments for Photodynamic Therapy

A. Holder ^a and D. Llagostera ^a

August 2, 2004

Abstract

Photodynamic therapy is a complex treatment for neoplastic diseases that uses the light-harvesting properties of a photosensitizer. The treatment depends on the amount of photosensitizer in the tissue and on the amount of light that is focused on the targeted area. We use a pharmacokinetic model to represent a photosensitizer's movement through the anatomy and design treatments with a linear program. This technique allows us to investigate how a treatment's success varies over time.

Keywords: Optimization, Photodynamic Therapy

^a Trinity University Mathematics, San Antonio, Texas, USA. Corresponding author's email is aholder@trinity.edu.

1 Introduction

Photodynamic therapy (PDT) is an emerging treatment that destroys undesirable tissues and is based on the idea that chemical compounds, called *photosensitizers*, cause cellular damage when activated by light. While the underlying chemical processes of PDT have been studied over the past century [6, 7, 14], the technique only recently received US approval for clinical treatments of some cancers, like esophageal and lung cancer. Canada recently approved trials for prostate cancer. The technique is widely used to treat age-related macular degeneration [3], but because of the growing potential for PDT to treat cancer [4], our focus is on modeling oncological treatments.

Photosensitizers exist in a variety of molecular compositions and structures, and different compositions react to different wavelengths of light [12, 14]. When a photosensitizer is activated by light, an electron is excited from a ground state to a more reactive state, converting it to a more energetic species. The excited photosensitizer converts oxygen to singlet oxygen that in turn destroys cellular material. The excited electron also causes the formation of super oxide that directly destroys cells. Though the photochemical reaction results in two destructive processes, the formation of singlet oxygen is regarded as the primary source of cell destruction.

A photosensitizer is designed to aid photodynamic therapy and is synthesized so that it has an affinity for targeted cells [5]. Cancerous cells possess a high content of low-density lipoprotein (LDL) receptors, and as a result, the photosensitizer is designed to be lipophilic in character. A photosensitizer is also constructed to be water soluble so that it easily flows through tissue. Moreover, a photosensitizer should quickly degrade, as this reduces the threat of a cytotoxic response (an extensive exposure of tissue to a foreign species that results in cell death) and the possibility of future destruction from non-clinical light sources.

In the clinical setting, a patient is injected with a photosensitizer, and after the drug accumulates in tumorous cells, a light source is focused on the cancerous region. This causes the photochemical reaction to occur where the photosensitizer is present. After treatment, the patient avoids light exposure for a period that depends on the rate at which the photosensitizer is eliminated from the body, a process that can take up to three months. If the patient does not avoid strong light sources, like the sun, it is possible for unwanted damage to occur.

The tissue destruction caused by PDT depends on the amount of light, photosensitizer, and oxygen in the tissue. Since these amounts vary over time, modeling cell destruction is complicated. Only a fraction of the photosensitizer is activated, and an even smaller fraction of these activated molecules form toxic substances. These

toxic substances accumulate until they surpass a threshold after which cell death occurs. As discussed in [7], the goal of the physician is to destroy as much tumor as possible while limiting the damage to surrounding tissues. Our goal is to accurately model a photosensitizer's flow through the anatomy and then use this model to design treatments that address a physician's goals. Blood and tissue concentrations of a Federal Drug Administration (FDA) approved photosensitizer are measured via standard pharmacokinetic techniques (pharmacokinetics is the mathematical analysis of drug absorption, drug distribution, and drug elimination [2]). Once the distribution of the photosensitizer is understood, an optimization model is designed that answers several questions related to PDT. Our optimization model shows how a treatment's success depends on the photosensitizer's affinity for cancerous cells and how light should be focused during treatment.

2 A Pharmacokinetic Model of Photofrin

The photosensitizer we consider is Photofrin, which is the only FDA approved photosensitizer for cancer treatment. While other photosensitizers are being clinically tested, accurate pharmacokinetic data is not available for these photosensitizers. Photofrin is intravenously injected at 2 mg/kg over a period of five minutes, and it attains a high saturation level from 48 to 72 hours after injection. The 24 hour period between the 48th and 72nd hour is called the treatment window, and at some time within this period a physician focuses a 630nm light on the targeted area. After treatment, a patient is asked to avoid light exposure for six weeks.

We measure Photofrin's concentration by placing a grid over the anatomy and calculating the concentration in each rectangular region. Different tissues absorb and expel substances from the blood at different rates, and hence the concentration of Photofrin varies from tissue to tissue. We make the simplifying assumption that each region is comprised of arterial and non-arterial tissue, where the non-arterial tissue is homogeneous. Our pharmacokinetic model estimates a region's concentration by first calculating the arterial concentration and second using this concentration to gauge the amount of Photofrin in the remaining tissue.

The arterial concentration is modeled via a differential equation that describes how Photofrin accumulates and degrades in blood-plasma. We let C_I be the plasma concentration during injection, k be the rate of elimination, and k_0 be the rate of infusion. The volume of distribution, V , for a particular drug is the volume of fluid that the drug would occupy if the total amount of drug in the body were at the same concentration as is present in the plasma. We model C_I with the following first order,

linear differential equation,

$$\frac{dC_I}{dt} = \frac{k_0}{V} - kC_I,$$

whose solution is,

$$C_I = \frac{k_0}{Vk} [1 - e^{-kt}].$$

The plasma concentration after infusion is C_A , which exponentially decays. So, assuming that the infusion ends at time τ , we have that

$$\frac{dC_A}{dt} = -kC_A, \quad t \geq \tau,$$

from which we conclude that

$$C_A = C_I(\tau)e^{-kt} = \frac{k_0}{Vk} [1 - e^{-k\tau}]e^{-k(t-\tau)}, \quad t \geq \tau.$$

The infusion rate k_0 is based on the fact that the drug is injected at 2mg/kg over a period of 5 minutes. Using an average 70 kg adult, we have that 140 mg of Photofrin are delivered in 5 minutes, and thus k_0 is 1680 mg/h. The mean volume of distribution for Photofrin is 0.49 L/kg, and in the case of the 70 kg person, V is 34.3 L. The rate of elimination, k , is calculated by solving,

$$t_{\frac{1}{2}} = \ln(2)/k,$$

where the half-life for Photofrin is $t_{\frac{1}{2}} = 516$ hours. Hence, $k = 0.00134 \text{ h}^{-1}$. Using these pharmacokinetic parameters for Photofrin¹ and an infusion time of $\tau = .083$ h (5 min), we have for $t \geq .083$ that the plasma concentration is

$$C_A = 4.08e^{-.00134(t-.083)}.$$

We convert C_A to Molarity by dividing the concentrations by the molar mass of Photofrin (596798.4 mg/mol), which means we redefine C_A to be

$$C_A = \frac{1}{596798.4} (4.08e^{-.00134(t-.083)}).$$

We use the arterial concentration C_A to model Photofrin's concentration within a region that contains both arterial and non-arterial tissues. We adapt the blood and

¹All parameters are FDA approved and are based on clinical data located at http://www.fda.gov/cder/foi/label/2003/20451s012_photofrin_1b1.pdf

tissue flow equations in [13], where a two-stage blood flow model sufficiently estimates tissue concentrations of radioactively labeled water. We let $C_p(t)$ be the Photofrin concentration at time t in the non-arterial tissue and $C_{porph}(t)$ be the Photofrin concentration in the entire region. We let k_1 and k_2 be the tracer rate constants that describe the flow from tissue to blood and from blood to tissue, respectively. Allowing V_A to be the percentage of arterial tissue in a region, we use the following model to calculate a region's concentration at time T ,

$$\left. \begin{aligned} C_p(T) &= k_1 \int_0^T C_A(t) dt - k_2 \int_0^T C_p(t) dt \\ C_{Porph}(T) &= V_A C_A(T) + C_p(T). \end{aligned} \right\} \quad (1)$$

We assume that the rates k_1 and k_2 are the same ($k_1 = k_2 = 1$). The time interval is divided into equal time steps of length Δt , and we have from [13] that Model (1) leads to the following approximate Photofrin concentrations,

$$C_{Porph}(T) \approx \frac{V_A C_A(T) + (k_1 + V_A k_2) \int_0^T C_A(t) dt - k_2 \left[\int_0^{T-\Delta t} C_{Porph}(t) dt + \Delta t C_{Porph}(T - \Delta t) / 2 \right]}{1 + k_2 \Delta t / 2}.$$

This approximation is exact as $\Delta t \downarrow 0$.

3 Photofrin Activation and Tissue Destruction

The goal of this section is to calculate the rate at which tissue is destroyed during treatment. This rate of destruction depends on the Photofrin concentrations of the previous section and on the path of the light source, and we begin this section with an investigation into how light travels through the anatomy. Consider Figure 1, where a light source is directed at the patient along angle a . We let $d_{(p,a)}$ be the depth of region p along angle a and $\delta_{(p,a)}$ be the distance between the primary ray and region p . Light attenuates as it passes through the anatomy (see [10]), meaning that it loses energy due to absorption and scatter. The attenuation is experimentally close to exponential decay, with region p receiving

$$\Delta = e^{-\mu_1 d_{(p,a)} - (\mu_2 \delta_{(p,a)} / d_{(p,a)}^{1/2})}$$

units of light from a single unit of light entering the anatomy along angle a . The parameters μ_1 and μ_2 are decided by the type of light and the width of the beam, with

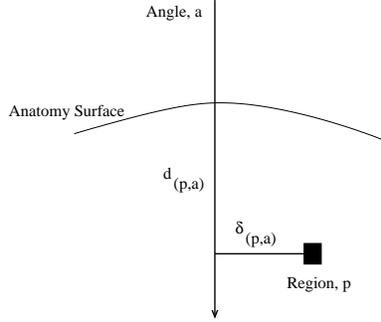


Figure 1: The depth measurements of a simple light model.

current technology producing a 5 mm cylindrical beam of 630 nm infrared light. We set $\mu_1 = 0.02$ and $\mu_2 = 0.93$ so that the light model in Figures 2 and 3 matched the experimental data in [10].

We have from [10] that a typical cell perishes once the concentration of activated Photofrin reaches 17 mM. Allowing $\alpha_{(p,a,t)}$ to be the rate at which Photofrin is activated in region p by a light source focused along angle a at time t , we have that the activated Photofrin in region p is

$$\alpha_{(p,a,t)}x_{(a,t)},$$

where $x_{(a,t)}$ is the length of time (in seconds) that the light is focused along angle a at time t . The rate at which Photofrin is activated depends on the laser irradiance Ψ , the extinction coefficient of the photosensitizer ε , and the concentration of photosensitizer C_{porph} at time t . Including the off-axis penetration of light, we have that

$$\alpha_{(p,a,t)} = \varepsilon\Psi C_{Porph}(t)\Delta. \quad (2)$$

The irradiance of our light source is $\Psi = 0.2 \text{ J/cm}^2\text{s}$, and the extinction coefficient is $\varepsilon = 10^4 \text{ L/(mole cm)}$. In order for $\alpha_{(p,a,t)}$ to be in units of Molarity, the equation is multiplied by the conversion constant $\lambda/(N_Ahc) \text{ mol/J}$, where λ is the wavelength of light, N_A is Avogadro's number, h is Planck's constant, and c is the speed of light.

4 An Optimization Model

In this section we design an optimization model that is based on the rates $\alpha_{(p,a,t)}$. The goals of this model are to 1) calculate the best time to treat within the standard 24 hour treatment window, 2) optimally design a treatment, and 3) investigate

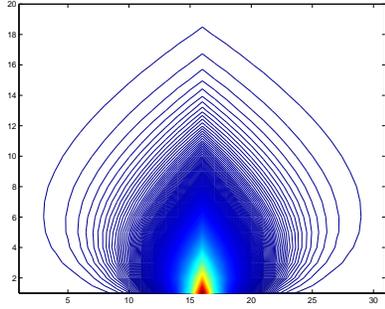


Figure 2: A contour plot of light attenuation and scatter.

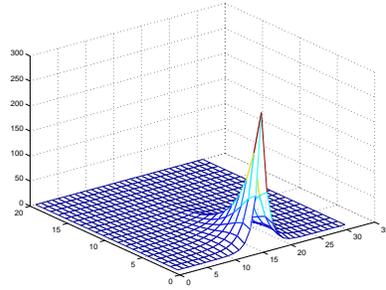


Figure 3: Our model of light deposition.

how a photosensitizer’s affinity effects a treatment’s success. We point out that our problem is similar to that of optimally designing radiotherapy treatments, for which a substantial literature already exists (see [1, 8, 15, 16]). From an optimization perspective, the problems are nearly identical, and we modify the model in [9] to meet our needs. However, radiotherapy uses high-energy particle physics to model how ionizing radiation damages cells, whereas PDT uses biochemical models to understand how an activated photosensitizer destroys tissue. The physics models estimate how the anatomy absorbs radiation, and radiotherapy is grounded in the fact that cancerous tissues are slightly more susceptible to radiation than normal tissues, a treatment property called *therapeutic advantage*. In these models, physicians use their experience to decide the amount of radiation that is likely to kill a cancerous cell but at the same time spare healthy cells. Unlike PDT, where we know that cell death occurs once the concentration of activated Photofrin surpasses a threshold, the physics models do not have a value after which we are guaranteed cell death. Whether or not there is an innate therapeutic advantage for PDT is unknown, and hence a healthy cell’s survival depends directly on treatment design and the photosensitizer’s affinity for cancerous cells.

As in the previous section, we let $x_{(a,t)}$ be the time (in seconds) that a light source is directed along angle a in time period t . Since $\alpha_{(p,a,t)}$ is the rate at which the concentration of activated Photofrin accumulates in region p , we have that the concentration of activated Photofrin in region p is $\sum_{(a,t)} \alpha_{(p,a,t)} x_{(a,t)}$. Notice that we are assuming an additive accumulation of activated Photofrin, which means that we are ignoring cell regeneration during treatment. Cell growth and division have

been investigated in the radiotherapy literature, where it is understood that cellular repair is negligible compared to the damage observed during treatment. While PDT and radiotherapy differ at the molecular level, it is reasonable to assume that the regenerative effects of the cell cycle are minute in both treatments during the relatively short treatment times (under an hour in both).

We form a *dose matrix*, A , from the rates $\alpha_{(p,a,t)}$ by allowing the rows of A to be indexed by p and the columns by (a,t) . Allowing x to be the vector of $x_{(a,t)}$'s, where the indices correspond to the columns of A , we have that $x \mapsto Ax$ is the linear operator that takes exposure times, x , and maps them to concentrations of activated Photofrin, Ax . We partition the rows of A in the following way,

$$A = \begin{bmatrix} A_T \\ A_C \\ A_N \end{bmatrix} \begin{array}{l} \leftarrow \text{Tumor dose points} \\ \leftarrow \text{Critical dose points} \\ \leftarrow \text{Normal dose points.} \end{array}$$

With this notation, we have that A_Tx , A_Cx , and A_Nx are the concentrations of activated Photofrin in the tumorous, critical, and normal tissues, respectively.

One of our goals is to decide exposure times that best treat the patient. We assume that the patient image is divided into a 64×64 grid (so a patient image has 4096 regions). A physician delineates the three types of tissue on this grid, and our task is to use this information to decide when the photosensitizer is best distributed and how the treatment should be administered. For each of the hours within the standard treatment window, we generate a dose matrix that represents the hour in 15 minute intervals. For example, for the 50th hour we fix time at 50, 50.25, 50.50, 50.75 and 60 hours, and for each of these times we calculate a dose matrix A_i , $i = 0, 1, 2, 3, 4$. The dose matrix for the entire hour is $A = [A_0|A_1|A_2|A_3|A_4]$. For each hour in the treatment window we solve the following linear optimization problem,

$$\begin{array}{l} \min \quad \omega_1 e^T \rho + \omega_2 e^T \beta + \omega_3 e^T \gamma \\ \text{s.t.} \quad \left. \begin{array}{l} A_T x \geq (0.017 + \varepsilon) - \rho \\ A_C x \leq 0.017 + \beta \\ A_N x \leq 0.017 + \gamma \\ x, \rho, \beta, \gamma \geq 0 \\ \rho \leq (0.017 + \varepsilon), \end{array} \right\} \end{array} \quad (3)$$

where e is the vector of ones (length is decided by the context of its use). The first three constraints are *elastic*, meaning that their right-hand sides adjusts with ρ , β , and γ . If these variables are zero, we have achieved the destruction of the tumor and spared other tissues (assuming the 2nd and 3rd inequalities hold strictly, more

on this below). The objective function attempts to make ρ , β , and γ as small as possible, where the weights ω_1 , ω_2 , and ω_3 decide the relative importance of each tissue type. The parameter ε is used so that ρ need not be zero to guarantee the destruction of a cancerous region. If ω_2 and ω_3 are held constant, we have from [9] that $e^T \rho^*(\omega_1) = O(1/\omega_1)$, where $\rho^*(\omega_1)$ is optimal for ω_1 . This result shows that for an appropriately large ω_1 we are guaranteed that $\rho^*(\omega_1) < \varepsilon$. We use $\omega_1 = 10$, $\omega_2 = 1$, and $\omega_3 = 1/10$. The value of ε is 10^{-3} .

The choice of solver is important to a treatment's success. There are two (competing) perspectives on solving linear optimization problems. The classical simplex approach is combinatorial in nature and terminates with a basic-optimal solution. Such solutions guarantee that many constraints hold with equality, which is not desirable as this guarantees that several non-tumorous regions attain the threshold of 0.017M, causing unwanted and often unnecessary damage. Instead of altering the model, we use an interior-point method that converges to a solution with clinical advantages. Many interior-point algorithms follow a structure, called the central path, toward optimality, and this structure terminates in the strict interior of the optimal face. Unless there is a unique optimal solution, these algorithms return a solution that is characteristically different than a basic-optimal solution. In particular, interior solutions strictly satisfy as many inequalities as possible. Hence, using a path-following interior-point algorithm ensures us that we spare as much non-tumorous tissue as possible.

Our experiments are conducted on the 4cm \times 4cm geometries in Figures 4, 5, and 6. We only consider angles that are orthogonal to the image, and we assume that the light source can be focused on the patient from any of the 64 regions of an image's face (so there are $4 \cdot 64 = 256$ angles). As discussed in Section 2, each region is assumed to contain arterial and non-arterial tissues. Unfortunately, we do not know of a technique that estimates the arterial proportion of a region from a patient image, and we randomly assign arterial concentrations from a uniform distribution.

We solved each example four times, two with Photofrin having no affinity for cancerous cells and two with Photofrin having an increased affinity for cancer cells. This separation allows us to investigate how a treatment's success is affected by higher absorption rates in cancerous tissues. This is important because new photosensitizers are designed so that they accrue in cancer cells up to 6 times more than in muscle [11]. Unfortunately, the chemical properties of these new photosensitizers are not available, and as a surrogate, we use the known properties of Photofrin and assume that it has a higher affinity for cancerous cells. From [11] we know that the level of increased accumulation depends on time and peaks near the middle of the treatment window. For our experiments we consider the situation where Photofrin concentrations in

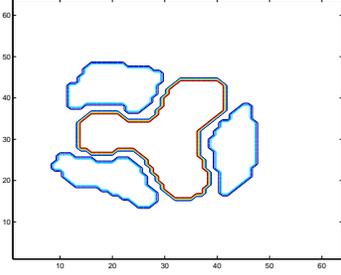


Figure 4: A tumor growing between three critical structures.

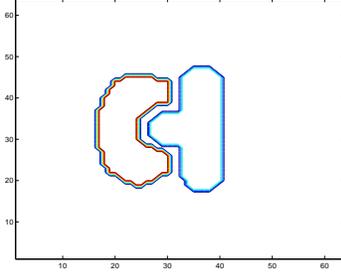


Figure 5: A tumor growing around the left side of a critical structure.

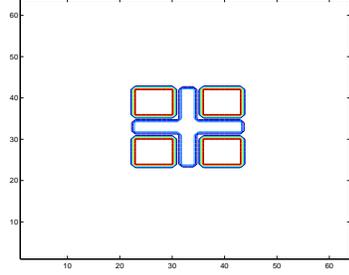


Figure 6: Four cancerous regions divided by a sensitive structure.

cancerous regions are multiplied by

$$6 - \frac{5}{12}|t - 60|, \text{ for } 48 \leq t \leq 72.$$

Destruction of the entire tumor is not guaranteed by solving model (3) with our choices of ω_1 , ω_2 and ω_3 . If ρ and ε are removed, the model ensures that the entire tumor is destroyed and minimizes the damage to the critical and normal tissues. Since the primary focus of the treatment is to remove the threat of cancer, we designed treatments for each example with and without ρ and ε . So, each example was solved 4 times:

- Photofrin without an increased cancer affinity and ρ and ε removed,
- Photofrin without an increased cancer affinity and ρ and ε included,
- Photofrin with an increased cancer affinity and ρ and ε removed, and
- Photofrin with an increased cancer affinity and ρ and ε included.

Figures 7, 8, and 9 show the percent of tissue that is destroyed when there is no increased Photofrin accumulation in cancer cells. The curves labeled ‘Tumor’ and ‘Critical’ show how the tumor and critical structures are destroyed by a treatment that is designed with ρ and ε included. The curve labeled ‘Critical with Tumor Destroyed’ indicates the critical structure damage with ρ and ε removed -i.e. when the entire cancerous region is destroyed. The fact that these curves are nearly flat indicates (surprisingly) that without heightened Photofrin concentrations in cancerous tissues, a treatment’s characteristics are nearly invariant over the treatment window. This

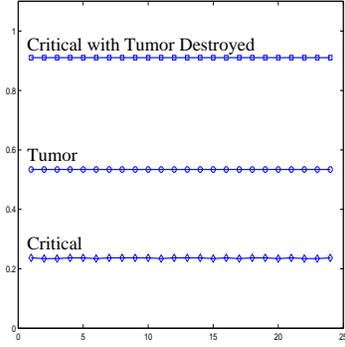


Figure 7: Percent tissue destruction for the example in Figure 4 with increased tumor affinity ignored.

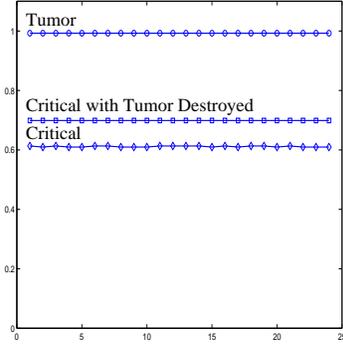


Figure 8: Percent tissue destruction for the example in Figure 5 with increased tumor affinity ignored.

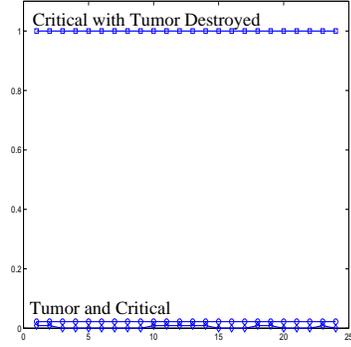


Figure 9: Percent tissue destruction for the example in Figure 6 with increased tumor affinity ignored.

does not mean that the decision variables are the same for each hour, but instead that the treatments adjust as the Photofrin concentrations change so that the same amount of tissue is destroyed.

The three examples illustrate that a treatment's success can vary significantly. From Figure 7 we see that it is possible to destroy 53% of the tumor in Figure 4 while destroying 24% of the critical structures. However, guaranteeing the destruction of the entire tumor requires the destruction of 91% of the critical structures. The problem in Figure 5 fairs better, as Figure 8 shows that 99% of the tumor is destroyed while less than 62% of the critical structures is destroyed. If we guarantee tumor destruction, the critical structures lose 70% of their functional mass. The situation in Figure 9 is significantly worse, and these curves indicate that destroying the tumor is the same as destroying the critical structures.

Figures 10, 11, and 12 depict what happens if Photofrin has an increased affinity for cancerous tissues. The tissue destruction varies over the treatment window, and each example has an improved treatment. This is especially true for the treatments in Figure 12, where we see that it is possible to destroy 90% of the tumor and only 36% of the critical structures. Compared with Figure 9, we see that for this example the increased concentrations of Photofrin in the cancerous regions are crucial to a successful treatment. The statistics in Table 1 are for the hour having the largest separation in destruction between the tumor and critical structures. We point out that the graphs in Figures 10, 11, and 12 exhibit irregular behavior in the 23rd and 24th

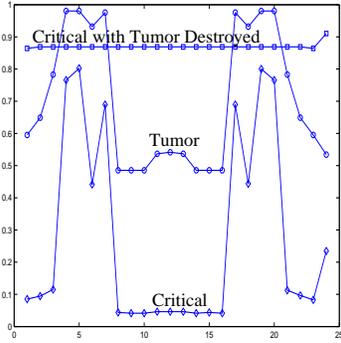


Figure 10: Percent tissue destruction for the example in Figure 4 with increased tumor affinity.

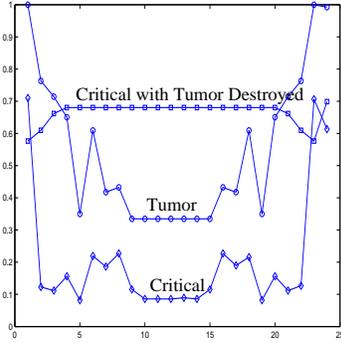


Figure 11: Percent tissue destruction for the example in Figure 5 with increased tumor affinity.

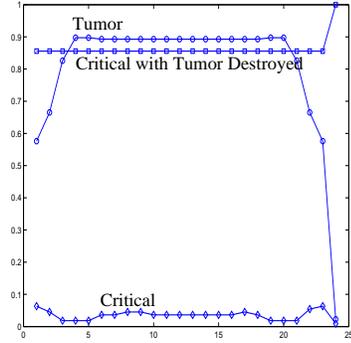


Figure 12: Percent tissue destruction for the example in Figure 6 with increased tumor affinity.

	Hour	Tumor	Critical	Separation
Prob. 1	69	78.29%	11.26%	67.03%
Prob. 2	50	76.32%	12.27%	64.05%
Prob. 3	52	89.73%	1.8%	87.93%

Table 1: The treatment hour with the largest separation between destruction of the cancerous and critical regions.

hour of the treatment window. After a lengthy investigation, we found that the system used by the interior-point algorithm to decide a step direction becomes increasingly ill-conditioned as time approaches the 24th treatment hour. This behavior was observed on numerous examples, and appears to be a product of the differential equation that models Photofrin concentration.

5 Conclusion

Photodynamic therapy has significant potential in the fight against cancer, especially as chemists improve a photosensitizer’s capability to localize in cancerous cells. Our goals were to 1) model a photosensitizer’s blood and tissue concentrations, 2) show that optimization can aid in the design and delivery of treatments, and 3) investigate how increased localization in cancerous cells effects a treatment’s success. As

one would expect, magnifying a photosensitizer's concentration in tumorous tissue improves the optimal treatment. This change is substantial for the example in Figure 6, and it is clear that improving a photosensitizer's affinity for cancerous cells will dramatically improve PDT's viability.

References

- [1] F. Bartolozzi et al. Operational research techniques in medical treatment and diagnosis. a review. *European Journal of Operations Research*, 121(3):435–466, 2000.
- [2] D.W.A. Bourne, E.J. Triggs, and M. J. Eadie. *Pharmacokinetics for the Non-Mathematical*. MTP Press Limited, 1986.
- [3] Neil M. Bressler and Susan B. Bressler. Photodynamic therapy with verteporfin(visudyne): Impact on ophthalmology and visual sciences. *Investigative Ophthalmology and Visual Science*, 41:624–628, 2000.
- [4] Hopper C. Photodynamic therapy: a clinical reality in the treatment of cancer. *Lancet Oncol.*, 1:212–9, 2000.
- [5] A. Chwilkowska, J. Saczko, T. Modrzycka, A. Marcinkowska, A. Malarska, J. Bielewicz, D. Patalas, and T. Banas. Uptake of photofrin ii, a photosensitizer used in photodynamic therapy, by tumor cells in vivo. *Acta Biochimica Polonica*, 50(2):509–513, 2003.
- [6] T. Hasan, B. Ortel, and A. Moor. Mechanisms of photodynamic therapy. In T. Patrice, D. Haeder, and G. Jori, editors, *Reviews in Photosciences: Photodynamic Therapy*. Amsterdam: Elsevier, 2000.
- [7] T. Hasan, B. Ortel, A. Moor, and B. Pogue. Photodynamic therapy of cancer. In R.E. Pollock, R.R.Weichselbaum, T.S. Gansler, J.F. Holland, E. Frei III, R.C. Bast Jr, and D.W. Kufe, editors, *Cancer Medicine 5th Edition*, pages 489–502. B.C. Decker, 2003.
- [8] A. Holder. Radiotherapy treatment design and linear programming. Technical Report 70, Trinity University Mathematics, San Antonio, TX, 2002. to appear in the Handbook of Operations Research/Management Science Applications in Health Care.

- [9] A. Holder. Designing radiotherapy plans with elastic constraints and interior point methods. *Health Care and Management Science*, 6(1):5–16, 2003.
- [10] Steven L. Jacques. Laser-tissue interactions: Photochemical, photothermal, and photomechanical. *SURGICAL CLINICS OF NORTH AMERICA*, 72(3):531–558, 1992.
- [11] Giulio Jori. Factors controlling the selectivity and efficiency of tumour damage in photodynamic therapy. *Lasers in Medical Science*, 5(115), 1990.
- [12] J. McMurry. *Organic Chemistry*. Brooks/Cole, 5th edition edition, 2000.
- [13] V. Oikonen. Model equations for [¹⁵O]h₂O pet perfusion (blood flow) studies. Technical Report report TPCMOD0004, Turku PET Centre Modeling, Turku University Central Hospital, 2002.
- [14] Ravindra K. Pandey and Gang Zheng. Porphyrins as photosensitizers in photodynamic therapy. In K.M. Kadish, K.M. Smith, and R.Guilard, editors, *The Porphyrin Handbook*, volume 6. Academic Press, 2000.
- [15] I. Rosen, R. Lane, S. Morrill, and J. Belli. Treatment plan optimization using linear programming. *Medical Physics*, 18(2):141–152, 1991.
- [16] D. Shepard, M. Ferris, G. Olivera, and T. Mackie. Optimizing the delivery of radiation therapy to cancer patients. *SIAM Review*, 41(4):721–744, 1999.