Conformal light delivery using tailored cylindrical diffusers

A. Rendon[a,b], J. Okawa[c], R. Weersink[b], J.C. Beck[c] and Lothar Lilge[a,b]

[a]Department of Medical Biophysics, University of Toronto, 610 University Ave., Toronto, Canada;

[b]University Health Network, 610 University Ave., Toronto, Canada;

[c]Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Rd., Toronto, Canada

ABSTRACT

Tailored light diffusers offer the flexibility of shaping the delivered light dose (fluence rate) distribution, potentially leading to conformal light delivery. Because of scattering and absorption, tissue acts as a spatial low pass filter of the diffuser's emission profile, and therefore some dose distributions with high spatial frequencies cannot be delivered. We characterize the set of attainable light dose distributions in terms of the spatial frequency of the emission profile and identify regimes where such distributions are less sensitive to changes in optical properties. Furthermore, we contrast two different algorithms to solve the inverse problem: Simulated Annealing (SA) and Non-negative Least Squares (NNLS). SA is plagued by superimposed high frequency components that do not contribute significantly to the cost. We present an iterative low pass filter that smooths the emission profile without considerably increasing the cost. A non-negative least square (NNLS) algorithm is also tested. We conclude that non-negative least squares (NNLS) is superior to simulated annealing (SA) in terms of time performance and cost minimization.

Keywords: tailored diffusers, dosimetry, photodynamic therapy, interstitial, conformal, inverse problem, simulated annealing

1. INTRODUCTION

In principle, the ability to tailor the longitudinal emission profile of fiber based light diffusers opens a great opportunity for conformal interstitial light delivery. Although a great variety of emission profiles can be manufactured, this does not imply that the same flexibility exists for the attainable light dose distributions. This condition also makes the inverse problem, i.e. finding an emission profile that corresponds to a desired dose distribution, difficult to solve. Here we present our efforts of characterizing the attainable light dose distributions and solving the inverse problem, in the context of a single diffuser, as is relevant in the treatment of selected conditions. A single diffuser is also relevant to the multi-diffuser light delivery problem because sharp light attenuation implies that dose is largely dictated by the sources closest to the point of interest. In this case, one can think of a diffuser being responsible for conforming to a particular region of the treatment surface, with little contribution from other diffusers.

To be useful, characterizing the attainable light dose distributions has to be general enough that it is informative, but restricted to a small number of parameters. One would proceed by defining a set of emission profiles and characterizing the resulting light dose distributions, in particular, the resulting isodose lines. Ideally, such emission profiles would constitute an orthogonal basis for expressing any other emission profile. When the light propagation problem is defined as a convolution of a point source function G with an emission profile S, the harmonic functions (ensuring positivity) constitute an adequate basis choice for the emission profile.[1] The fluence rate $\phi(\rho, z)$ takes the form:

$$\phi(\rho,z) = \int_L G(\rho, z-z')S(z')\mathrm{d}z' \quad (1)$$

$$\approx \int_{-\infty}^{+\infty} G(\rho, z-z')(0.5+0.5\cos(kz'))\mathrm{d}z' \quad (2)$$

$$= \mathcal{F}(\rho, k=0) + \cos(kz)\mathcal{F}(\rho,k), \quad (3)$$

Send correspondence to: L.L. `llilge@uhnresearch.ca`

Optical Methods for Tumor Treatment and Detection: Mechanisms and Techniques in Photodynamic Therapy XVI, edited by David Kessel, Proc. of SPIE Vol. 6427, 64270M, (2007) · 1605-7422/07/$18 · doi: 10.1117/12.700996

where the approximation is due to extending the integration interval. The function $\mathcal{F}(\rho, k)$ denotes the Fourier transform in k space of $G(\rho, z)$ with respect to z. Equation (3) is used in Sect. 2 to characterize the attainable fluence rate distributions. Here we use

$$G(\rho, z) = \frac{3\mu_{\text{eff}}^2 a'}{4\pi\mu_{\text{a}}} \frac{\exp(-\mu_{\text{eff}} \cdot r)}{r},$$

with $r = \sqrt{(\rho^2 + z^2)}$, and where a' is the reduced albedo and μ_{eff} and μ_{a} are the effective and absorption coefficients, respectively.[2,3] With this choice,

$$\mathcal{F} = \frac{3\mu_{\text{eff}}^2 a'}{4\pi\mu_{\text{a}}} K_0(\rho\sqrt{\mu_{\text{eff}}^2 + k^2}),$$

where K_0 is the zeroth-order Bessel function of the second kind, which behaves approximately like $\exp(-x)/\sqrt{x}$.

We will describe the isodose lines of Eqn. (3) in Sect. 2 in order to characterize the attainable light dose distributions. In particular we will relate the amplitude d of such contour lines to the distance from their center to the diffuser ρ_0, as a function of k and μ_{eff}. Figure 1 identifies these parameters.

$\mathcal{F}(\rho, k)$ falls off steeply as a function of ρ and k therefore acting as a low pass filter. This makes solving the inverse problem not trivial. While diffuser manufacturing prefers a relatively smooth emission profile, the nature of the inverse problem often leads to solutions with high spatial frequencies or consisting of disjoint collection of point sources. In Sect. 3 we will describe our efforts to solve the inverse problem for a single diffuser.

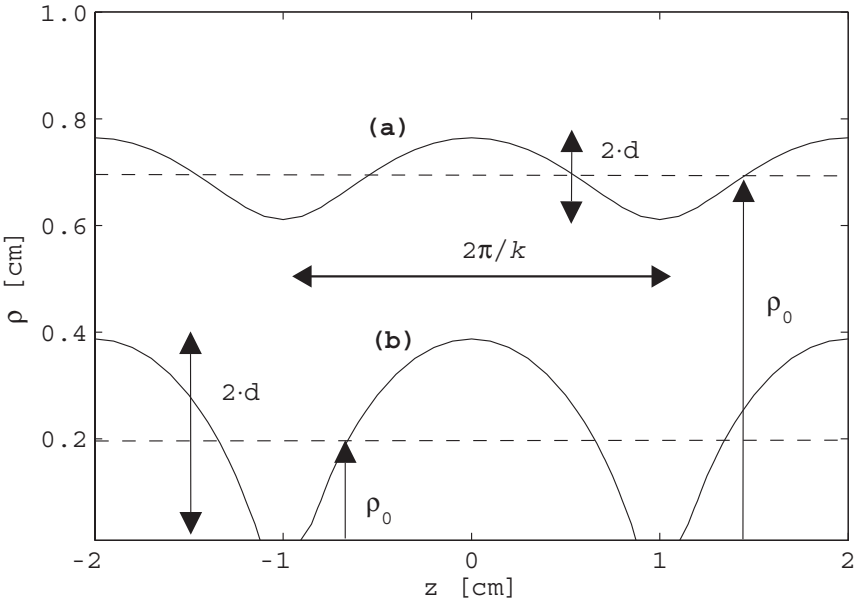

Figure 1. Description of the geometrical parameters used in the text. Two contour lines for the emission profile $S = 0.5 + 0.5\cos(\pi z)$ are depicted. These represent two different regimes: (**a**) contour lines fully detached from diffuser (top line), and (**b**) contour lines continuous with the z axis. The full amplitude of the contour lines is labeled $2d$ and the distance from the center of the contour line to the z axis (the diffuser) is labeled ρ_0. Notice that the principal spatial frequency is maintained but other spatial frequencies that depend on ρ also emerge.

2. FORWARD PROBLEM

Characterizing the forward problem serves several purposes. In particular, it permits the characterization of the attainable fluence rate distributions by examining the shape of the contour lines as a function of spatial

frequency. From this, one can determine regimes where using tailored diffusers is most appropriate. In this respect, there are two aspects of the contour lines that are important: the maximum modulation along ρ (our d parameter) and the robustness of the shape of a contour line with respect to changes in the optical properties. One would like to understand how these properties vary with increasing ρ and with k as a function of μ_{eff}. For this purpose we have numerically found the contour lines of Eqn. (3).

Contour lines are iso-fluence rate curves. Figure 2 shows two representations of Eqn. (3). The two regimes presented in Fig. 1 correspond to the contour lines with fluence rates above or below the intersection of the minima of the fluence rate and the z axis. This fluence rate can be found by taking the limit of Eqn. (3) when ρ approaches 0, for z such that $\cos(kz) = -1$. The value of the fluence rate for the separation of the contour lines from the contour lines is given by $\phi_s = 3\mu_{\text{eff}}^2 a'/(4\pi\mu_{\text{a}}) \log(1 + k^2/\mu_{\text{eff}}^2)$, for the point source function specified above.

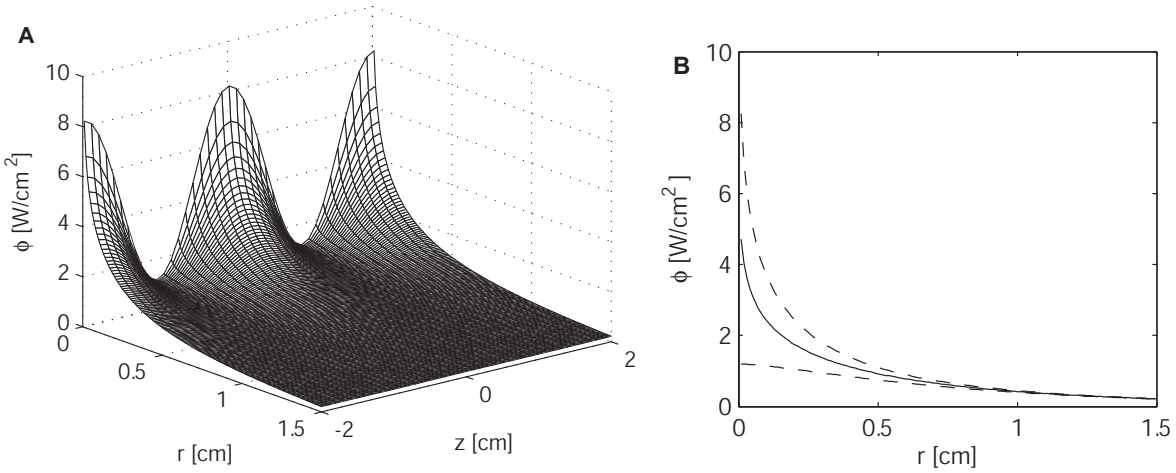

Figure 2. Fluence rate arising from the source profile $S = 0.5 + 0.5\cos(\pi z)$ with $\mu_{\text{eff}} = 1$ cm^{-1}, according to Eqn. (3). Contour lines for this source were shown in Fig. 1. (A) Mesh graph shows the overall shape of the function. (B) Side view along the ρ axis: the z-independent term in Eqn. (3) (solid line) and the values of the maximum (z=0 cm) and minimum (z=1 cm) fluence rates (dashed lines) are shown. Notice that the fluence rate that separates the two contour level regimes depicted in Fig. 1 corresponds to the minimum fluence rate at $\rho = 0$. This is also the fluence rate that achieves maximum modulation as seen in (B).

2.1. Modulation transfer properties

The maximum achievable amplitude d depends on ρ_0, k and the optical properties. Figure 3 depicts this dependence. For contour lines of fluence rates less than ϕ_s the amplitude equals ρ_0, as expected from the definition of these dimensions. This linear behavior is followed by a fast decay of amplitude, the maximum being achieved for ϕ_s. Furthermore, as the spatial frequency decreases (larger periods) the amplitude increases, however its sensitivity to the optical properties increases as well, as is evidenced by comparing d for a given k between the different panels of Fig. 3, and noticing that the change in d is larger for smaller k.

Figure 3 also shows that even for small μ_{eff} and a periodicity of 6 cm^{-1} the maximum value that d takes is approximately 0.7 cm. These dimensions roughly define an ellipsoid with semi-axes 3 and 1.4 cm long, which corresponds to the maximum treatment volume for such a source.

2.2. Robustness to changes in optical properties

In vivo optical properties are difficult to measure and often vary between treatment planning and light delivery, or even throughout light delivery (e.g. due to changes in blood flow and oxygenation). This results in changes in the light dose distribution. The total fluence delivered at a given contour level can be adjusted by increasing the power or the irradiation time. This kind of dosimetry is amenable to online detection and correction. On

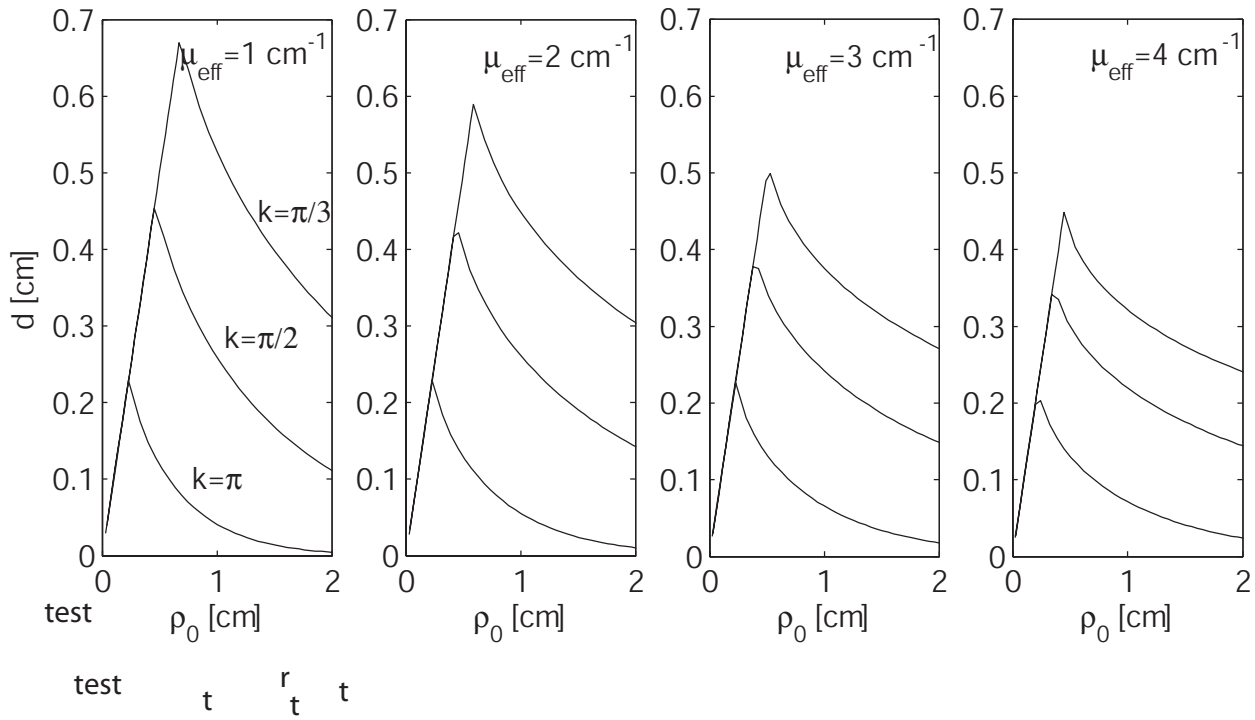

Figure 3. Dependence of d with ρ_0, for different spatial frequencies k and optical properties (as indicated in the panels). The curves in each panel correspond to spatial frequencies $k = \pi, \pi/2, \pi/3$ cm^{-1} from bottom to top, in each panel as labeled.

the other hand, the emission profile is fixed. It then becomes important to identify regimes where the shape of the contour lines is robust to changes in the optical properties.

We have looked at the effect of different optical properties on d for a given ρ_0, shown in Fig. 4. There are two important trends to note. Higher spatial frequencies are less sensitive to changes in the optical properties, however display comparatively smaller d. This is evidenced by vertical stripes with similar d in the first panel of Fig. 4 (compare to the other two panels). Additionally, regime **b** defined in Fig. 1 is more robust to changes in μ_{eff}. Therefore, as a general guideline it is better to position tailored light diffusers within this regime. Solutions in regime **a** are desired because they are further away from the diffuser, thus requiring less diffusers to cover the same volume, but these solutions are less stable and can lead to unexpected regions of over or under-treatment. This effect is more pronounced for sources with small spatial frequencies, i.e. when conforming to large features.

We have only looked at changes in the shape along ρ. These changes are associated with the transfer of modulation. Changes along z also exist and can be relevant to treatment planning, but are not described here. This can be investigated for example by the emergence of new spatial frequencies of the contour lines as a function of ρ. These spatial frequencies are suppressed as ρ_0 increases, until the contour lines are sinusoidal with the same spatial frequency of the source (cf. Fig. 1 regimes **a** and **b**).

3. INVERSE PROBLEM

In practice Eqn. (2) is solved numerically for an arbitrary source $S(z')$. The kernel $G(\rho, z - z')$ can be sampled and given a matrix form following lexicographic ordering of ρ and z along the rows and indexing z' along the columns. Then Eqn. (2) becomes

$$\phi_\alpha = \Delta z' \sum_{\beta}^{N} G_{\alpha\beta} S_\beta, \tag{4}$$

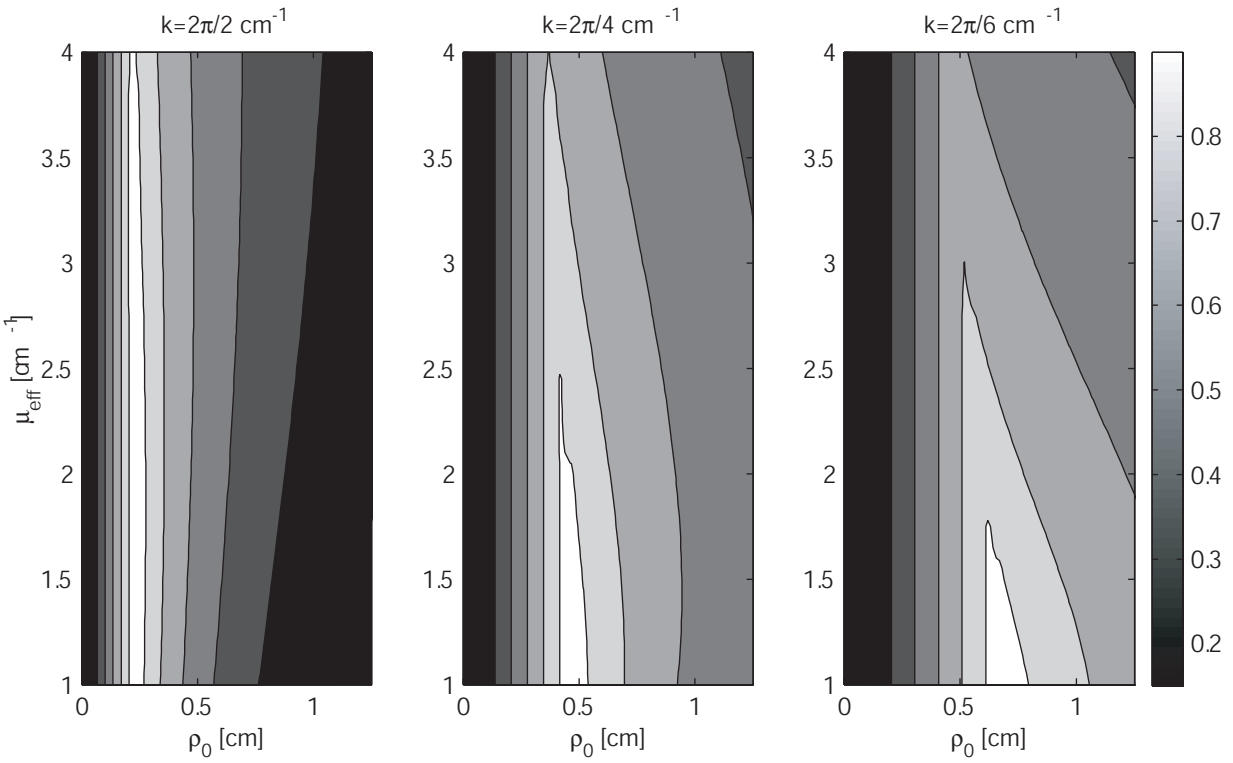

Figure 4. Effects in d due to changes in the optical properties. Three different spatial frequencies of the emission profile are presented as labeled above each panel. Each panel shows the ratio of d to its maximum value in each panel. Contour lines for the ratio are placed at multiples of 0.15. The maximum values of d for each panel from left to right are: 0.23, 0.46, and 0.68 cm.

with $\Delta z' N$ the length of the diffuser. Currently, there is only one manufacturer of tailored diffusers.[4] The manufacturing specifications define $\Delta z'$, the spacing of the emission profile vector S_β, as 1.4 mm. In the remainder, $\Delta z'$ will be conveniently incorporated into $G_{\alpha\beta}$. The spacing of the space variables ρ and z is given by the resolution of the imaging system that acquired the anatomical information. We set the voxel size at a reasonable 1.4 mm on the side.

Assuming a threshold model of tissue response, the clinician's task is to define a set of light doses (fluence values) at locations on the tumor anatomy, e.g. the threshold dose at the treatment volume boundary. Let us call this prescription D_α, where α indexes the space coordinates. The inverse problem becomes finding an emission profile S_β with a resulting dose distribution that best conforms to D_α. To this aim one defines

$$\text{cost} = \left| D_\alpha - \sum_\beta G_{\alpha\beta} S_\beta \right| \tag{5}$$

for the cost function that evaluates how well the dose conforms to the prescription. This cost function is then minimized in terms of S_β. Equation (5) is the simplest of multiple cost functions available. Such cost functions may be subject to sets of constraints

We have investigated two different algorithms to minimize Eqn. (5), non-negative least squares (NNLS) and simulated annealing (SA). NNLS, supplied in MATLAB as `lsqnonneg`, is an iterative algorithm that restricts solutions to non-negative numbers, as required here.[5] SA is a stochastic minimum search method, typically less prone than local search methods to terminating at local minima. We have implemented the algorithm in MATLAB according to the method of Kirkpatrick et al.[6] NNLS is one to two orders of magnitude faster than

SA depending on the data set. For S_β partitioned in 50 intervals and 50^3 voxels, typical execution times on a Pentium 4 (2.8GHz, 512MB) for NNLS range between 0.5 to 1 seconds, and for SA between 30 to 60 seconds. While speed is not critical, it is desirable, particularly when solving multiple diffuser problems.

In addition to speed, solutions must be manufacturable. When a solution exists (see Fig. 5 A), such as when an emission profile is defined and D_α taken from the resulting light distribution, NNLS often converges to the right solution. This convergence is not always achieved because NNLS is somewhat sensitive to initial conditions. For both algorithms, we pre-compute a starting point S_0 by increasing the power at all positions evenly in small intervals, until the first minimum occurs. Starting points for NNLS are then generated randomly from the uniform interval between zero and twice the common power calculated above. In our trials, approximately 45% of these starting points converge to the same result (data not shown). Choosing the initial condition to be S_0, makes NNLS converge to the right solution in all the test cases.

SA is started directly from S_0 and stochastically finds a minimum every time. These minima have very similar cost but correspond to different solutions. The solutions are all characterized by high frequency components superimposed on the original emission profile. This is not surprising when realizing that tissue acts as a low pass filter, therefore high frequencies do not contribute substantially to the cost. This observation is used in Sect. 3.1 to derive and iterative smoothing filter for the SA solutions.

In Sect. 2 we identified regimes where an exact solution does not exist. In these regimes (see Fig. 5 B), NNLS returns spiked solutions constituted by a small collection of spaced point sources with large emission powers. Since diffusers are function by coupling light out of the fiber core via etching of scattering centers, diffusers tailored according to these solutions have particularly high powers remaining at the distal end (since overall little light has been coupled out) and display increased fragility due to fiber etching. In this case SA has the advantage of returning smoother solutions, although with costs higher than the NNLS solutions and still containing high frequency components. The SA algorithm randomly chooses a point along the diffuser and changes its power up or down. The solutions obtained by the NNLS algorithm are difficult to obtain using SA because of the small likelihood that random choices would bring the solution to such shape, given that the cost increase is not significantly higher. This also explains why SA solutions are, notwithstanding the high frequency components, smoother on average. One can alternatively take the NNLS solutions and apply a smoothing filter that recapitulates the positive properties of the SA solutions.

3.1. Iterative low pass filter

There is a simple way to remove the high frequency components from the SA solutions. Considering the stochastic nature of the process, one can average multiple solutions thus reducing uncorrelated noise. To quantify the deviation of a solution from the ideal source profile, one calculates the mean square error (MSE) $\langle |S_\beta - S_\beta^{ideal}| \rangle$ between the two profiles. Averaging multiple solutions decreased both the cost and MSE aproximately as a function of $N^{(0.5)}$, where N is the number of solutions averaged (data not shown). The exact exponent of N depends on the particular ideal solutions. While an effective approach, in practice it becomes prohibitively time consuming to calculate these solutions.

An alternative approach is to directly remove the high frequency components using an iterative low pass filter that determines the cutoff frequency when a significant increase in the cost occurs. We have implemented this by iteratively removing higher frequency components (starting by the highest), calculating the mean and standard deviation of the cost of all preceding iterations and determining whether the current cost is significantly different than the mean. For this comparison, a Student's t-test is used, with the number of iterations corresponding to the number of degrees of freedom. The results of filtering the same data set from Fig. 5 A are shown in figure Fig. 6.

As seen in Fig. 6, a compromise between increasing the cost and reducing the MSE is obtained. Furthermore, increasing the significance level also has the same effect. This simple filter is proof of principle that filtering is possible with a favorable compromise. One can think of filters that can offer better performance in terms of cost and possibly MSE, using a smoother frequency cutoff (here a simple step function was used).

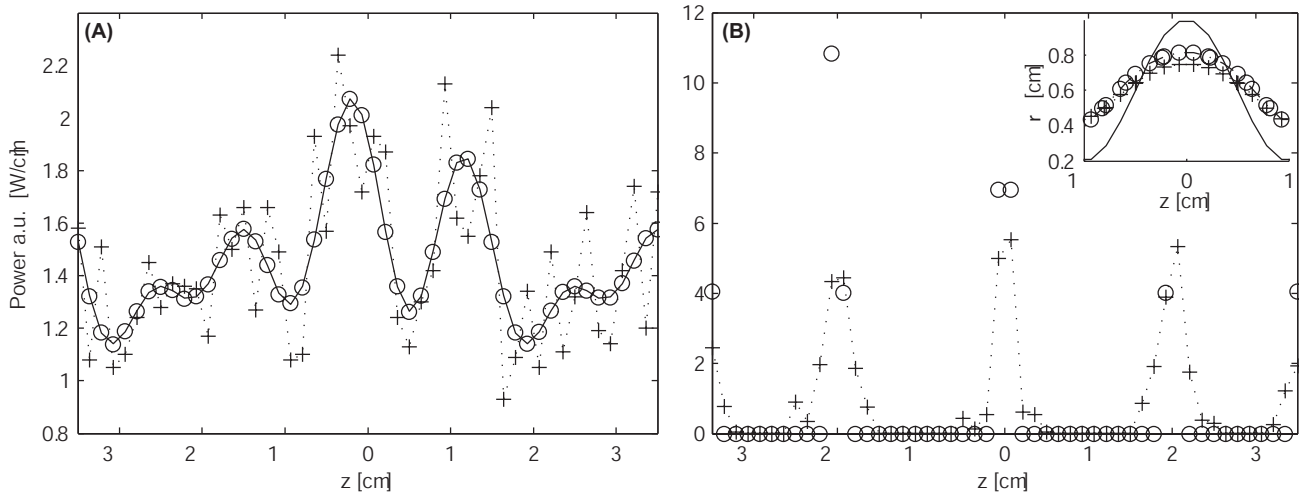

Figure 5. Types of solutions to the inverse problem with the NNLS (circles) and SA (crosses) algorithms. Panel **(A)** shows profiles obtained when a solution exists. Notice that the NNLS solution is indistinguishable from the ideal solution (solid line). Also notice the high spatial frequency components on the SA solutions. Panel **(B)** shows the type of profiles that result when no exact solution exists. The NNLS profile is constituted by a collection of point sources of with large powers, while the SA profile is more continuous and of smaller powers. Solutions in **(A)** were obtained by defining the ideal emission profile shown, forward calculating the fluence rate at $\rho = 0.5$ cm and solving with the indicated algorithms. In **(B)**, a boundary of constant unit fluence rate ($\rho = 0.6 + 0.4\cos(\pi z)$) was used for D_α and solutions calculated accordingly. The ideal light distribution (solid line) and the contour lines for the NNLS and SA solutions are shown in the inset. Notice that the NNLS contour line is closer to the prescribed one.

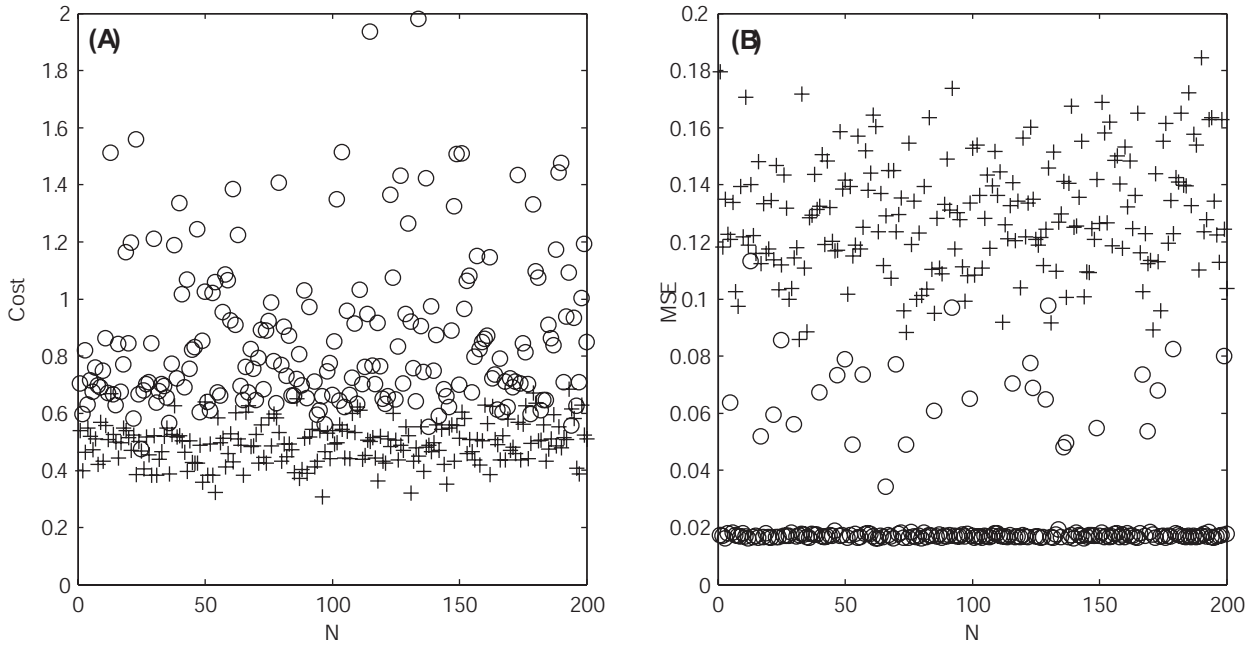

Figure 6. Iterative low pass filter applied to the same data set from Fig. 5 A. 200 SA solutions were analyzed for the statistical properties of the filter. The significance level was placed at 99%. **(A)** shows the cost prior (crosses) and after (circles) filtering. Notice the slight increase in the cost due to filtering. **(B)** depicts the mean square error of the solutions. The filtered solutions that groups at 0.02 all share an equal cutoff frequency lower than that of the remaining points. The ratio average costs is increased 1.8 times while the average MSE is reduced 5 times.

4. CONCLUSIONS

The potential for tailored light diffusers to improve conformal light delivery is limited by the physics of light propagation in tissue. Tissue acts as a low pass filter of the emission profile. As a result, high spatial frequencies rapidly lose modulation leading to a flattening of the light dose distribution. This behavior was characterized in Sect. 2. Additionally, regimes of stability of the light dose distribution with respect to the optical properties were described. Taken together, these two properties help define regimes where tailored diffusers can be useful for conformal light delivery.

We also began to investigate the behavior of the inverse problem. Initially we examined a SA algorithm because it is more likely to find the global minimum of the cost function. It was observed that the SA solutions have high spatial frequency components superimposed on a smooth curve. These high frequency components arise because small deviations from the ideal source have very small incremental cost and are often accepted in the Metropolis step.[6] We implemented an iterative low pass filter that successfully removed the high frequency noise, while marginally increasing the cost. Nevertheless, SA takes a long time to converge and, when compared to NNLS, the second algorithm investigated, it does not converge to a global minimum.

NNLS is a local search iterative algorithm that only results in non-negative solutions. For this reason, it has been classically used in image reconstruction. We found that NNLS executed two orders of magnitude faster than SA and resulted in solutions with overall lower costs. When no emission profile exists that can exactly conform to the prescribed dose, the solutions obtained with NNLS are typically comprised by a collection of point sources. When manufactured, such solutions lead to fragile diffusers with large remaining powers at the distal end. In principle, one can smooth these solutions to achieve better diffuser properties, clearly increasing the cost. Alternatively, one can incorporate the smoothing into the algorithm by solving a regularized problem.

Tailored diffusers are currently manufactured by etching imperfections that couple light out of the fiber core. One can, in principle, solve the problem in the space of the coupling constants (scattering coefficients) plus the input and distal powers. These three sets of parameters are equivalent to the description in terms of emission powers and fully describe the diffuser fabrication. Additionally they can be constrained to obtain solutions with adequate manufacturing properties.

Future work will investigate metrics to determine when a tailored diffuser does not perform significantly better than a standard flat diffuser. Finally, the findings from this work need to be taken into consideration for implementing solutions to the inverse problem with multiple tailored diffusers.

ACKNOWLEDGMENTS

We thank Dr. Leonid Vesselov in Walsh Medical Devices for kindly providing the fabrication specifications for the tailored diffusers. The authors acknowledge funding through the NIH PO1CA43892 and CIHR 68951.

REFERENCES

1. A. Rendon, R. Weersink, and L. Lilge, "Towards conformal light delivery using tailored cylindrical diffusers: attainable light dose distributions.," *Phys Med Biol* **51**, pp. 5967–5975, Dec 2006.
2. W. M. Star, "Light dosimetry in vivo," *Physics in Medicine and Biology* **42**, pp. 763–787, May 1997.
3. R. Graaff and K. Rinzema, "Practical improvements on photon diffusion theory: application to isotropic scattering," *Physics in Medicine and Biology* **46**, pp. 3043–3050, Nov. 2001.
4. L. Vesselov, W. Whittington, and L. Lilge, "Design and performance of thin cylindrical diffusers created in Ge-doped multimode optical fibers," *Applied Optics* **44**, pp. 2754–2758, May 2005.
5. C. Lawson and R. Hanson, *Solving Least Squares Problems*, Prentice-Hall, 1974.
6. S. Kirkpatrick, C. Gelatt, and M. Vecchi, "Optimization by simulated annealing," *Science* **220**(4598), pp. 671–680, 1983.