33rd Annual International Conference of the IEEE EMBS
Boston, Massachusetts USA, August 30 - September 3, 2011

# Voxel-Based, Parallel Simulation of Light in Skin Tissue for the Reconstruction of Subsurface Skin Lesion Volumes

Brian D'Alessandro, Atam P. Dhawan, *Fellow, IEEE*

***Abstract*—Early detection and diagnosis of skin cancer is essential to treating the malignancy and preventing death. Subsurface features and depth information are critical in evaluating a skin lesion for this early malignancy screening. We present a novel voxel-based Monte Carlo simulation of light propagation in skin tissue which runs in a highly parallel environment on desktop graphics processors, resulting in an extremely fast simulation of millions of photons in less than one second. We then use this model in a genetic algorithm for the inverse 3D volume reconstruction of a skin lesion, given a set of multispectral images obtained using non-invasive transillumination imaging. Our method demonstrates improved accuracy at a superior resolution to existing methods.**

## I. Introduction

MELANOMA is the deadliest form of skin cancer. However, the survival rate can be very good if the cancer is detected early. Deeper subsurface information, such as subcutaneous pigmentation, depth of invasion, and indications of increased blood flow (angiogenesis) are critical factors in early melanoma detection. As a result, noninvasive optical imaging techniques to detect and analyze the depth and morphological changes associated with tumorigenesis could improve patient diagnosis accuracy with minimal need for invasive and time consuming biopsy procedures. The Nevoscope is one such optical imaging device, designed to obtain subsurface information of suspicious skin lesions through multispectral transillumination imaging. However, the reconstruction of depth and chromophore information from transillumination images is not a simple problem.

In this paper, we present an exceptionally fast voxel-based parallel processing Monte Carlo simulation of light propagation in skin tissue. This simulation is then used in an iterative genetic algorithm to find a best estimate solution to the inverse volumetric reconstruction of the melanin and blood components of skin lesions, given only the set of surface transillumination images from the Nevoscope at specific wavelengths [see Fig. 1]. Since the depth of lesion invasion and vascularity are key features in the detection and classification of early skin cancer, this volumetric reconstruction is essential.

Manuscript received April 15, 2011.

Brian D'Alessandro (email: bmd5@njit.edu) and Atam P. Dhawan (email: dhawan@njit.edu) are with the Department of Electrical and Computer Engineering, New Jersey Institute of Technology, Newark, NJ 07102 USA.

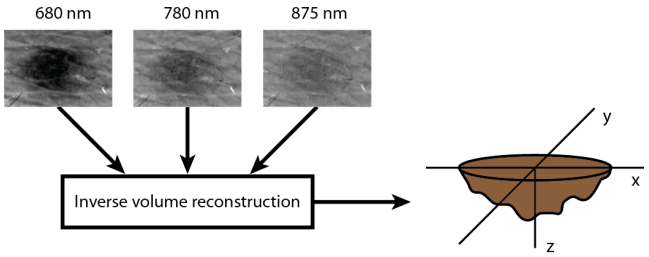


Fig. 1. Objective diagram. A set of multispectral transillumination images are input to the inverse reconstruction algorithm, resulting in a 3D estimate of the skin lesion volume.

## II. Transillumination Imaging of the Skin

With transillumination imaging using the Nevoscope, surface illumination is blocked. Instead, light photons enter directly into the skin at an angle through a fiber-optic illumination ring placed against the skin. These photons undergo multiple internal reflection, scattering, and absorption events depending on the chromophores encountered. The light eventually gets diffused across the layers of the skin and backscattered light photon energy forms a transilluminated image of the skin and skin lesion subsurface structures, which are normally obscured using conventional surface imaging. The backscattered image thus depends on the optical properties of the skin which light encounters on its path from source to destination, such as the chromophore and wavelength dependent absorption coefficient $\mu_a$. Thus the forward model of transillumination imaging can be represented by:

$$I_\lambda = F\left(\mu_a(\lambda)\right) \tag{1}$$

where $F$ is the forward model function for imaging, and $I_\lambda$ is the backscattered image at wavelength $\lambda$.

The most common chromophores encountered in the skin are melanin and the main components of blood: oxygenated hemoglobin and deoxygenated hemoglobin. It is the absorption coefficient at individual locations within the tissue volume which give critical information about the three dimensional distribution and depth of the lesion, as well as the relative amounts of the major chromophores [2]. Thus, our objective is to estimate $\mu_a$ for an unknown tissue volume based on the set of multispectral backscattered transillumination images. In other words, we wish to find the inverse model for volume reconstruction:

$$\mu_a(\lambda) = F^{-1}\left(I_\lambda\right) \tag{2}$$

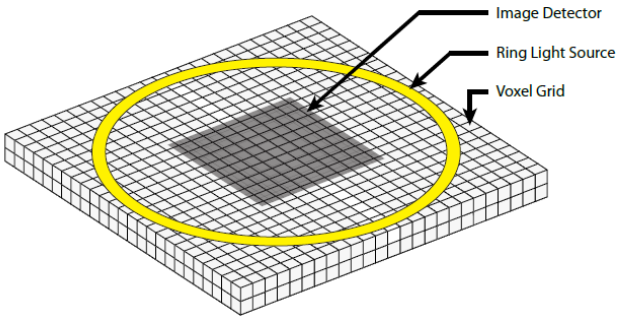


Fig. 2. Elements of the Monte Carlo simulation.

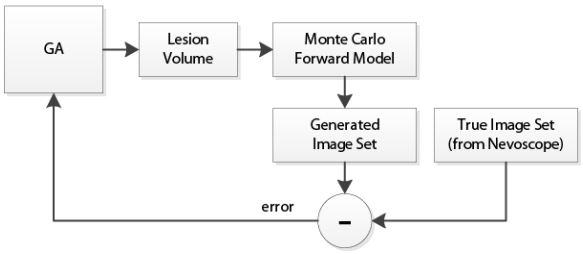


Fig. 3. Flowchart of the inverse reconstruction procedure.

While the inverse model is more difficult to evaluate, the forward imaging model $F$ can be simulated first using Monte Carlo methods.

## III. Parallel Simulation of Light Propagation

Monte Carlo simulation is a statistical technique for simulating random processes. The basic idea of a Monte Carlo simulation for light transport in tissue is to mathematically propagate one photon at a time through a virtual environment. When a large number of these photons are simulated, the stochastic nature of each individual photon's path averages out to a good approximation of light transport and diffusion through the tissue.

The MCML program has been a popular code base for simulating light propagation in tissue [3]. However, this program assumes the tissue volume to be made up only of homogenous layers. A voxel-based approach is more appropriate for simulating irregularly shaped lesions embedded in tissue, such as a skin lesion. With this approach, each voxel may have its own independent set of optical properties. Thus, the resolution of defining irregularly shaped lesions depends on the number of voxels defined in the simulation.

Since each photon is simulated independently, this type of Monte Carlo simulation is highly amenable to parallel execution. Notably, modern desktop graphics cards are designed exactly for parallel processing and can thus run thousands of threads in parallel, resulting in GPU (graphics processing unit) simulation speeds over 1000 times faster than the same program executing serially on a CPU. However, existing parallel implementations of MCML still only use the layer-based approach [4]. In this paper, we have implemented a voxel-based, parallel processing Monte Carlo simulation of transilluminated light transport in skin tissue. This program is capable of simulating the paths of millions of photons through a three dimensional voxel grid where each voxel has its own independent optical properties. Code is executed using the NVIDIA CUDA software development kit on a multi-GPU setup of an NVIDIA GeForce GTX 560 Ti and 9800 GT. Each simulated photon is assigned its own execution thread, and 20,000 threads are executed in parallel.

Simulation of a single photon begins at a random position within the area of the illumination ring. Each photon then proceeds through a number of steps within the voxel grid. At the end of each step, the photon scatters and its intensity weight is decreased. Exact values for step length, scattering direction, and weight reduction are computed based on the optical properties of the voxel in which the photon is currently located; however, each equation integrates some degree of randomness as well.

If a scattered photon finds its way back to the surface of the voxel grid and its position is within the region where the image detector is located, then the photon's intensity weight is counted as part of the diffuse reflectance. The image detector itself is made up of many individual detectors, thus producing a 2D pixelated image similar to what would be detected by the CCD image sensor attached to the Nevoscope during actual usage. Fig. 2 shows the virtual setup and positioning of the illumination ring and image detector over an example voxel grid.

Simulation using a large number of photons is desirable in order to increase the number of photons detected by the image sensor. A large number of detected photons increases the signal to noise ratio of the detected image, at the expense of computation time. However, with our parallel GPU-based implementation, simulation of millions of photons is possible in less than one second in a voxel-based environment.

## IV. Inverse Model for Volume Reconstruction

With the Monte Carlo method, the forward imaging model is well defined for simulation. However, the inverse model is much more difficult to find, and is ill-posed, meaning there are possibly many solutions. Since the subsurface 3D volume of a lesion is unknown, we wish to use the multispectral Nevoscope image set to reconstruct the volume with depth and chromophore information. Thus, the problem is to find the voxel grid with the optimal set of optical coefficients, which when passed through the forward model, produces images which are closest to the actual images obtained from the Nevoscope. This is an optimization problem, which can be solved by numerous methods, one of which is through a genetic algorithm [see Fig. 3].

Previously, this problem has been attempted through the use of Taylor series approximation to the forward model [1]. With this method, the first derivative, or Jacobian, was found using the Monte Carlo method. Thereafter, the detected image could be approximated by multiplying the Jacobian

matrix with a vector containing the absorption coefficients of each voxel. The reason for doing so is that since each proposed solution requires a forward model evaluation to determine its accuracy, a very large number of executions of the forward model are required. Taylor series approximation was used to save computation time.

However, with our voxel-based parallel Monte Carlo algorithm, simulation using enough photons for an acceptable single to noise ratio takes less than one second rather than multiple minutes. Thus, using our simulation directly for the forward model has the benefit of not needing any approximations at all; each solution can be validated by an independent run of the simulation. Taylor series approximation is only valid for small deviations from the point about which higher order derivatives were derived. Thus, large deviations in the absorption coefficient of a voxel from the background absorption value increases the error of reconstruction. Since the problem is already inherently ill-posed, any additional artificial error is desirous to be eliminated. The incredible speed of our GPU based Monte Carlo simulation has the capability to eliminate this approximation error.

A genetic algorithm (GA) is used to iteratively search the solution space for the best solution to the reconstruction problem [5]. Individuals in the population, represented by their chromosomes, are defined to be the set of absorption coefficients for each voxel in the voxel grid. The fitness of each chromosome is evaluated via a comparison between the multispectral images of the true lesion ( $I_\lambda^{\text{real}}$ ) and the images generated by the forward model ( $I_\lambda$ ) on that chromosome:

$$\text{fitness} = \frac{1}{N_\lambda} \sum_\lambda \frac{I_\lambda - I_\lambda^{\text{real}}}{I_\lambda^{\text{real}}} \qquad (3)$$

where $N_\lambda$ is the number of wavelengths used for imaging.

The implemented GA uses steady state reproduction without duplicates, crossover, mutation and a 'nudge' operator which mutates alleles by a small deviation from its current value. Only one operator is chosen for each generation. Initially, the operator probability is 80%, 10%, and 10% for crossover, mutate, and nudge respectively; but these are linearly scaled towards 20%, 40%, and 40% by the final generation. A high crossover rate early in the GA is beneficial for initial exploration through the search space, however, as the population converges, a higher mutation rate is beneficial for improved local search. Furthermore, the number of photons used for simulation is interpolated from $1.5 \times 10^6$ to $3 \times 10^6$ over the course of the GA. The reason for this is to save computation time where possible: faster, but noisier images are acceptable initially, while images with a higher SNR are needed during the final stages of the GA to better differentiate chromosomes which vary only by small mutations. After the conclusion of the GA, a simple hill climbing algorithm is run on the best fit individual to explore small 'nudges' in lesion depth at every point in the x-y plane.

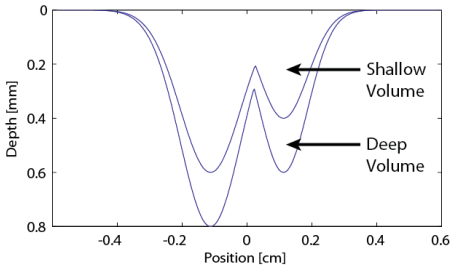


Fig. 4. The two layered skin lesion model. The shallow surface peaks at 0.6 mm and 0.4 mm while the deep surface peaks at 0.8 mm and 0.6 mm.

Nudges which result in a higher fitness are retained to produce a final best estimate solution of the lesion volume.

## V. Results

The speed improvements of our voxel-based parallel Monte Carlo code were substantial. Simulation of 10 million photons using the code from Wang et. al. [1] took 1 hour and 13 minutes on a 3.16 GHz dual core CPU, whereas our GPU enhanced code took only 1.29 seconds. This speedup of over 3000x allows the direct use of the simulation for fitness function evaluation in the GA without any approximations.

To test the accuracy of this method for inverse volume reconstruction, an artificial lesion was simulated. Initially, the same model was used as in the prior work [1] for accurate comparison of results. The model consists of a two-surface phantom lesion set within a 1.2 cm x 1.2 cm x 1 mm volume [see Fig. 4]. The volume bounded by the shallow surface contains only 5% melanin, while the volume bounded by the deep surface contains only 20% blood at 75% oxygen saturation. This suits observations that a skin lesion is made up primarily of melanin closest to the skin surface, with a deeper blood network below the lesion. The specific surfaces defined for simulation were based on a two-peak mixed Gaussian function. A grid of 16x16 control points was defined to be solved by the GA and each point consists of two values: the shallow surface depth and the deeper surface depth. The shallow surface depth was restricted to 0-0.6mm while the deep surface was restricted to 0-0.2mm. Segmentation of the grid to only the region of interest limited the number of control points to 35, significantly higher than the previous work which used only 9. Thus, each chromosome in the GA was of length (35 control points) x (2 depth values) = 70. Population size was set to 200, and the GA was run for 5,000 generations. Upon completion, the best solution was interpolated to produce a lesion volume of resolution 3.75 x 3.75 x 0.1 mm per voxel. The volume bounded by each reconstructed surface was compared with the true volume, and results are presented in Table I. Our proposed method was first evaluated using imaging at single wavelengths separately, 680 nm, 780 nm, and 875 nm. The method was then also evaluated using the multispectral imaging set of all three wavelengths combined, with equal weighting in the combined fitness function. Reconstructed volumes are visualized in Fig. 5.

Table I: Results comparison for a two surface lesion model with an upper volume of 5% melanin and lower volume of 20% blood.

| | Shallow Vol Error | Deep Vol Error |
|---|---|---|
| Wang, et. al. [1] | 2.68% | 16.58% |
| Proposed; 680 nm | 0.62% | 13.84% |
| Proposed; 780 nm | 0.64% | 13.46% |
| Proposed; 875 nm | 0.52% | 9.36% |
| Proposed; 680,780,875 nm | 0.35% | 6.48% |

Table II: Results for a two surface lesion model with an upper volume of 5% melanin and lower volume of a 20% blood, 5% melanin mixture. Imaging was performed using 680, 780, and 875 nm.

| | Shallow Vol Error | Deep Vol Error | Execution Time |
|---|---|---|---|
| 16x16 | 0.45% | 0.88% | 6152 s |
| 32x32 random | 6.10% | 9.20% | 13658 s |
| 32x32 seeded | 0.80% | 1.68% | 8365 s |

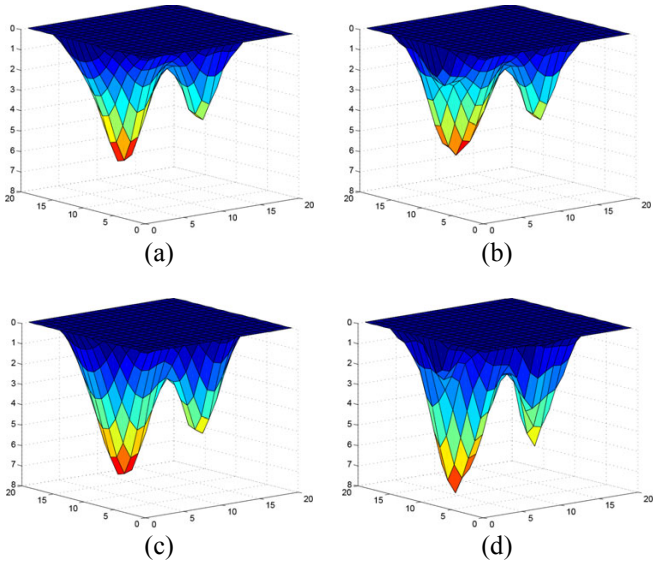


Fig. 5. Reconstruction results using 680, 780, and 875 nm imaging. (a) Shallow surface true model. (b) Shallow surface reconstruction. (c) Deep surface true model. (d) Deep surface reconstruction.

As can be seen, our parallel-processing inverse reconstruction algorithm provides superior results to the previous work in all trials. Furthermore, evaluation using the multispectral imaging set of three wavelengths is superior to using imaging at single wavelengths alone. Multispectral imaging thus provides additional information due to the contributions of different scattering profiles which assist in depth reconstruction. It is also apparent that the higher wavelengths, such as 875 nm, result in a lower error. This may be due to the fact that at lower wavelengths, the absorption due to blood is much lower compared to that of melanin, and hence is more difficult to recover.

Additional experiments were conducted, this time using a mixture of 5% melanin and 20% blood for the deeper volume with three wavelength imaging. Results are indicated in Table II. Reconstruction was very good for the 16x16 grid, with < 1% error for both volumes. Given such good results, reconstruction at a finer resolution, 32x32, was then attempted. Initially, a randomly initialized population of size 400 was iterated through 10,000 generations of the GA. Results were less accurate (volume error: 6.1% shallow; 9.2% deep) due to the much larger solution search space (the number of depth control points was quadrupled). Also, the higher generation count needed to reach convergence translated to a longer GA execution time. We found that rather than starting with a completely random initial population, the 32x32 population can be "seeded" by random variations about an interpolated version of the best 16x16 solution. This efficient initialization reduced the generation requirement to reach convergence back to 5,000 for a population size of 200, and resulted in a quicker and much more accurate solution (volume error: 0.8% shallow; 1.68% deep). Additionally, the shape of these reconstructed volumes matched the true shapes very well, with 2D correlation coefficients of 0.89 for the shallow volume and 0.87 for the deep volume.

## VI. Conclusion

In conclusion, we have presented a novel, voxel-based Monte Carlo simulation of light propagation implemented in a highly parallel manner on graphics processors, capable of simulating over a million photons in less than one second. We have then successfully used this parallel algorithm for an improved inverse reconstruction of skin lesion volumes from multispectral transillumination images with a higher accuracy and resolution compared to what was previously possible. In future work, we will investigate use of the method for the reconstruction of a volume where the concentrations of melanin and blood vary. Furthermore, we will investigate the possible benefit of using multiple populations in the GA to reconstruct skin lesion volumes at an even higher resolution with better accuracy.

## References

[1] S. Wang and A. P. Dhawan, "Shape-based multi-spectral optical image reconstruction through genetic algorithm based optimization," *Computerized Medical Imaging and Graphics,* vol. 32, pp. 429-441, 2008.

[2] B. D'Alessandro and A. P. Dhawan, "Depth-Dependent Hemoglobin Analysis From Multispectral Transillumination Images," *Biomedical Engineering, IEEE Transactions on,* vol. 57, pp. 2568-2571, 2010.

[3] L. Wang, S. L. Jacques, and L. Zheng, "MCML - Monte Carlo modeling of light transport in multi-layered tissues," *Computer Methods and Programs in Biomedicine,* vol. 47, pp. 131-146, 1995.

[4] E. Alerstam, T. Svensson, and S. Andersson-Engels, "Parallel computing with graphics processing units for high-speed Monte Carlo simulation of photon migration," *Journal of biomedical optics,* vol. 13, p. 060504, 2008.

[5] L. Davis, *Handbook of Genetic Algorithms*: Van Nostrand Reinhold, 1991.