Some ideas on algorithms for soft-tissue 'tomographic' reconstruction

Berthold K.P. Horn

1995/Feb/10

Background:

The medium is highly scattering. We assume for simplicity that the scattering properties are spatially uniform. There is also some absorption. The absorption density is expected to vary spatially. In fact, the result we are after is the spatial pattern of absorption.

We are allowed to apply light sources and make measurements only on the boundary of the volume of interest.

We expect to gather data by measuring brightness at many points on the surface while illuminating it at a particular point. These brightness measurements are repeated for many possible light source positions. If we assume fixed positions for sensors and emitters, and if there are are n light sensors and m light detectors, we obtain a data vector with $M = m * n$ components.

These measurements may not be all independent. Because of reciprocity, the brightness measure at P when the volume is illuminated at Q is equal to (?) the brightness at Q when the volume is instead illuminated at P.

The task is to recover as much information as possible about the internal distribution of absorbing material. We may expect as a rough rule of thumb that we should not try to recover an image of absorption density with more pixels than the number of independent measurements — $(m * n/2)$ say.

We *can* attempt reconstruction with a larger number of pixels but then we are dealing with an ill-posed problem that will need to be 'regularized'. In this underconstrained case we expect to be able to make changes in some pixel values that can be compensated for by opposite changes in neighboring pixels without affecting the measurements available at the boundary in a noticable way.

'DC' Model:

With intense scattering, the directions of travel of photons are randomized soon after they enter the volume. Any directionally of source or detector optics has very little effect more than a millimeter or so into the volume. At any given point in the volume, there is a given photon flux, with photons going in every which direction. The fraction of these

photons that are absorbed in a given volume depend on the absorption density there.

We can model this system using the diffusion equation. Equivalently we can treat it as a problem in heat conduction, or finally, as a resistive grid.

Consider a uniform resistive grid. The potential at a point in the grid represents the photon flux. We can inject currents (?) on the boundary of the grid and also measure potentials (photon flux) at points on the boundary.

This grid represents the scattering behavior of the material. Without any additional factors (such absorption or boundary conditions), the potential in the grid will rise without bound if we inject a fixed current. This corresponds to a volume of scattering material with no photon loss, and a stream of photons being injected at one point.

Absorption is represented by leakage conductance from nodes of the grid to ground. The leakage current is proportional to the local potential (photon flux) and the conductance of the leakage resistor (absorption density). A grid with some leakage will stabilize with fixed finite potential distribution where the total incoming photon flux is balanced by the total lost through absorption.

An aside: Note that in the 'DC' case, we do not need to know the scattering length per se (which corresponds to the resistivity of the grid). What matters is the ratio of conductance of the leakage path to conductance of the grid. That is, the 'absorption per scattering length'. In reconstruction this means that we would only need to know the scattering length accurately if we wanted to know the absorption density in absolute terms.

Boundary Conditions:

Viewed this way, it becomes clear that we need to be precise about what happens on the boundary. One extreme case would be application of a totally reflective material (either specular/mirror like or diffuse such as a perfectly white powder — it makes no difference). This corresponds to open circuit conditions on the boundary of the resistive grid.

Another extreme case would be total absorption on the boundary (black material applied directly, or the object is placed in a large dark volume). This corresponds to shorting the boundary of the resistive grid to ground.

(Another option might be applying material with the same scattering/absorbing properties on the boundary. This 'matched termination' would correspond to extending the grid indefinitely in space.)

It is not clear at this point whether better signal data can be collected in the 'open circuit' or in the 'short circuit' condition (It is not obvious to me anyway, but maybe to somebody else?)

Forward Problem:

The forward problem is to compute, for each light source position what measurements can be expected at each of the light sensing positions, given the spatial pattern of absorption density. Basically 'all' one need do is solve a spatially varying Poisson's equation:

$$\Delta v(x,y) + \rho\, v(x,y)\, a(x,y)) = 0$$

where $v(x,y)$ is the potential (light flux density), ρ is the resistivity of the grid (light scattering) and $a(x,y)$ is the leakage conductivity to ground (light absorption density). In the forward problem we are given $a(x,y)$ and asked to find $v(x,y)$.

Suitable boundary conditions are defined (open circuit or short-circuit to ground) and current is injected at one point on the boundary. The potential along the boundary is the sought after result. This problem can be solved effectively using iterative methods. (Minor difficulties may arise if $a(x,y)$ has abrupt transitions).

The result of the computation is a set of $M = m * n$ numbers (m light detector outputs for each of n possible light source positions).

Inverse Problem:

Unfortunately we are interested in the *inverse* of the above problem. We can collect the $m * n$ element measurement vector, and want to know the spatial distribution of light absorption density $a(x,y)$.

One way to attack this problem is an iterative approach where we update/correct a current guess for $a'(x,y)$. Given $a'(x,y)$, we can compute the expected measurement vector $E'(i * m + j)$ as above. We can compare this to the actually observed measurements $E(i * m + j)$. Then we can try to make adjustments in $a'(x,y)$ to bring $E'(...)$ into better agreement with $E(...)$. When $E'(...) = E(..)$ we may expect that $a'(x,y) = a(x,y)$.

Several optimization methods come to mind here. One would be a form of gradient descent. Such methods usually require that a derivative/gradient be obtained. That is, we need to know how each of the M components of the measurement vector depend on each of the N absorption density values. Based on this we can make an adjustment that yields the best improvement of fit in a least square sense. That is, it reduces

$$||E' - E||^2$$

This computation could be quite expensive. At each step we need to solve the resistive grid, and then compute the derivatives. There are $N * M$ components to these derivatives! And computing the derivatives may involve numerical estimates that require solving the grid again several times after making small adjustements.

It is impractical to actually store the full gradient. For example, suppose we use 100 photodiode sensors and 100 laser diode emitters. Then we have $M = 10,000$ measurements. Suppose we construct a grid with about the same number of nodes $N = 10,000$. The gradient then would require 10^8 numbers (and 100 sensors and 100 sources is not a lot since one needs a certain number to achieve sufficient spatial resolution — see below).

We also need to be concerned about whether such an iterative method converges, and whether it can get stuck in local minima.

Some ideas on the iterative scheme

The potentials on the boundary do not vary linearly with the conductance of a particular leakage path to ground. One thing that may help is that the grid *is* linear in another respect. If we inject a current at a node, we expect the response to be proportional to the current injected. A small change in a leakage conductance will produce a small change in leakage current. So for small increments, the response is linear.

Another things that may help is good choice of discretization of the resistive grid. For example, if we work with a planar grid (2-d case) and the boundary is circular, then a rotationally symmetric layout has benefits. Consider for example, resistors laid out along radii equally space in angle, with other resistors laid out along concentric circles. While this layout has some problems with uneven sampling and a singularity at the origin, it also has the advantage that we do not need to separately understand the response of the grid (light intensities at n detectors) to an absorber placed at each of the N nodes for *each* of m possible light source positions (that is, 10^8 numbers). Instead, because of rotational symmetry, we need consider only *one* light source position on the boundary (leading to 'only' 10^6 numbers).

Also, methods that do not compute the complete derivative may have an advantage. We may not need to walk in the direction of steepest descent, as long as most of the time we are going 'down hill'. The notion of a 'stochastic derivative' may play a role here. The disadvantage of using a stochastic derivative is that while saving a great deal of computation, it is not guaranteed to take us 'down hill' at every step.

To avoid local minima (either inherent in the problem or because of short-cuts used in the iterative method) one may want to use simulated annealing. Usually this is a very slow method. Recent work on so-called 'fast' simulated annealing (still slow, but not quite glacial) may be relevant here. In 'fast' simulated annealing, the temperature is reduced with the inverse of time rather than logarithmically.

An iterative method can always benefit from a good first guess. It may be possible to get a good first guess by using a fan-beam back-projection approach.

Some heuristics may also come into play. If the signal of detectors on a particular part of the periphery is low for most source positions, then one may expect a concentration of absorbing material near that area. Similarly, if the signal at most detectors is low when for *source* positions in some particular part of the periphery, then one may expect a concentration of absorbing material near *that* area.

Three Dimensions:

In most of what is said above, no mention was made of whether the volume of interest — or the model resistive grid — are two or three dimensional. Many of the same considerations apply in both cases.

However, the computations get more expensive with a full three-d grid. So for initial work and perhaps even later — for computational reasons — one may wish to stay with a 2-d grid. If the volume of interest is really three dimensional, then one is of course distorting the problem in some sense. Reconstructions of the interior will at a minimum be some sort of smeared superimposition of layers. Reconstruction nearer the surface is likely to be less affected by this.

We don't expect to be able to reconstruct well deep inside the object in any case, since the distribution of light on the boundary is most directly affected by (i) absorbing material near a source, and (ii) absorbing material near a sensor.

If computational limitations were not a concern, then distribution of source positions and the sensor positons over the 2-d surface of the 3-d volume would have advantages. We could most likely get better conditioning of the problem. Here is why:

Consider first a circular 2-D case with n sensors and m detectors. Suppose for simplicity that $n = m$. If the radius is R, then the spacing of sources and detectors is $\delta = 2\pi R/n = 2\pi R/m$. If we assume an equispaced grid of nodes covering the area we would use a grid interval of perhaps a little more than $\epsilon = \sqrt{\pi R^2/n * m}$. If $n = m$, we find that the

pixel size is about

$$\epsilon \approx 1/2\sqrt{\pi}\delta,$$

where δ is the spacing of sources and sensors on the peripheri. The problem becomes underconstrained if we try to reconstruct on a grid finer than ϵ. (This is not to say we can neccessarily achieve this kind of resolution). So the pixel spacing is (within some small factor) the same as the sensor/detector spacing.

Consider instead a spherical volume. Here the spacing of sensors and detectors on the surface would be roughly $\delta = \sqrt{4\pi R^2/n} = \sqrt{4\pi R^2/m}$. A reasonable voxel size would be $\epsilon = \sqrt[3]{(4/3)\pi R^3/(n * m)}$. So

$$\epsilon \approx 1/(12\pi)^{1/3}\delta\,(\delta/R)^{1/3}$$

Note the $(\delta/R)^{1/3}$ term. Here the voxel size drops *faster* than linearly with the sensor spacing. So we may expect to get resolution substantially smaller than the sensor and/or source spacing.

This is not too surprising. In the two-d case we have a match of dimensionality. We gather a '2-d' data set (for all source positions times all sensor positions) and we want to reconstruct a 2-d absorption density. In the three-d case we gather a '4-d' data set (for all source positions times all sensor positions) and we want to reconstruct a 3-d absorption density. More constraint. But also likely to be a lot more work. Still, this is the 'real' problem one ultimately wants to solve. The two-d version is an approximation.

AC Model:

If we use a modulated source, then we need to take into account something other than the simple static case described above. This situation can be modelled by adding uniformly distributed capacitance to ground from the resistive grid. The forward problem now involves solving something like

$$\Delta v(x,y) + \rho\, v(x,y)\, a(x,y) = \rho C \frac{\partial v(x,y)}{\partial t}$$

Which takes considerably more computing power. On the other hand, we now measure both phase and amplitude on the boundary. Also, results from simpler geometries suggest that the 'AC' signal is more useful for reconstruction than the 'DC' signal. Its just a little scary to look at iterative methods that would involve solving the capatively loaded grid many times.

One additional problem is that here we need to know ρ/C (related to the scattering length). In the 'DC' case, the actual resistivity of the resitive grid was of little concern, only the ratio of resistivity to leakage resistance was of importance (related to the absorption over the scattering length).

Using an analog network as a fast Poisson solver:

Nothing solves the resistive grid equations faster than the resistive grid itself! Since solving the resitive grid is an intimate step of the iterative process, one might consider actually building such a grid and using it to solve the forward part of the problem.

One problem would be the need for a fairly high bandwidth to/from the grid, since N conductances to ground need to be loaded up and then the grid solved for each of n choices for the position of the point where current is injected. Each solution involves reading out m boundary potentials.

Maybe more serious is that we expect the signal to have a very wide dynamic range. The photon flux density on the side opposite to the source is likely to be several orders of magnitude lower than that on a part of the boundary tyhat is closer to the source. And analog resistive grids may be difficult to build so that wide dynamic range is not a problem.

Other thought:

As in most 'inverse' problems, we need to be aware that reconstructed image quality may be poor in terms of resolution and/or signal to noise ratio. This is seen clearly in 'image restoration' problems, which are near ill posed. Still, when there are enough measurements, and the measurement accuracy is high enough, good results can be obtained, as seen in CAT reconstructions. So one thing we need to strive for is a large enough volume of measurement data to over-constrain the problem and high enough quality in the measurements to compensate for amplification of noise in the reconstruction process.

The 'imaging' process we use produces serious 'smearing' if we want to look at it that way, which suggests one will be faced with the kind of problems that are seen in image enhancement and image restoration. On the other hand, we are getting many 'projections' so we may find that tomographic reconstruction provides a better analogy. It is hard to predict, since this problem really seems to fall somewhere in between.

References

1. Horn, B.K.P., "Density Reconstruction Using Arbitrary Ray Sampling Schemes," *Proceedings of the IEEE*, Vol. 66, No. 5, May 1978, pp. 551–562.
2. Horn, B.K.P., "Fan-Beam Reconstruction Methods," *Proceedings of the IEEE*, Vol. 67, No. 12, December 1979, pp. 1616–1623.
3. Horn, B.K.P. *Robot Vision*, MIT Press, Cambridge, Massachusetts and McGraw-Hill, New York, 1986.
4. Horn, B.K.P. & M.J. Brooks (eds.) *Shape from Shading*, MIT Press, Cambridge Massachusetts, 1989.
5. Fogarty, M.A. "Light transmission and diffusion in breast tissue," B.S. in Physics, M.I.T., June 1993.
6. M. Lee "Optical Diffusion Tomography Data Acquisition System," S. M. Thesis MIT, January, 1995.
7. M. Lee, F. H. Bowman, A. B. Brill, B. K. P. Horn, R. C. Lanza, J. Y. Park, and C. G. Sodini, "Optical Diffusion Data Acquisition System," presented at *Future Direction of Lasers in Surgery and Medicine*, Salt Lake City, Utah, July, 1995.
8. Ware, K.M., H.-S. Lee and C.G. Sodini, "A 200 MHz CMOS phase-locked loop with dual phase detectors," *Int. Solid-State Circuits Conference Digest of Technical Papers*, pp. 192-193, 1989.
9. Barbour, R.L., et al. "Imaging of diffusing media by a progressive iterative backprojection method using time-domain data," *SPIE*, Vol. 1641, pp. 21-34, 1992.
10. Chance, B., et al. "Comparison of time-resolved and unresolved measurements of deoxyhemoglobin in brain," *Proc. Nat. Acad. Sci.*, Vol. 85, pp. 4971-4975, 1988.
11. Wang, L.M., P.P. Ho, C. Liu, G. Zhang and R.R. Alfano, "Ballistic 2-D imaging through scattering walls using an ultrafast optical Kerr gate," *Science*, Vol. 253, 769, 1991.
12. Hebden, G. and R.A. Kruger, "A time-of-flight breast imaging system: spatial resolution performance," *Proc. Time Resolved Spectroscopy and Imaging of Tissues, SPIE*, Vol. 1431, pp. 225-231, 1991.
13. Arridge, S.R., M. Cope, P. van der Zee, P.J. Hillson and D.T. Delpy, "Visualization of the oxygenation state of brain and muscle in newborn infants by near infrared transillumination," *Information Processing*

in Medical Imaging, S.L. Bacharach, ed., (Martinus Nijhoff: Holland), pp. 155-176, 1985.

14. Aronson, R., R.L. Barbour, J. Lubowsky and H. Graber, "Application of transport theory to NIR medical imaging," *Modern Mathematical Models in Transport Theory: Advances and Applications*, Vol. 51, Birkhouser Press, pp. 64-75, 1991.
15. Chance, B., et al. "Highly sensitive object location in tissue models with linear in-phase and anti-phase multi-element optical arrays in one and two dimensions," *Proc. Nat. Acad. Sci.*, Vol. 90, pp. 3423-3427, April 1993.
16. Tromberg, B.J., L.O. Svaasand, T.-T. Tsay, R.C. Haskell and M.W. Berns, "Optical property measurements in turbid media using frequency domain photon migration," *SPIE*, Vol. 1525, pp. 52-58, 1991.
17. Firbank, M. and D.T. Delpy, "A design for a stable and reproduciple phantom for use in near infra-red imaging and spectroscopy," *Phys. Med. Biol.*, Vol. 38, pp. 847-853, 1993.
18. Grünbaum, F.A., P. Kohn, G.A. Latham, J.R. Singer and J.P. Zubelli, "Diffuse tomography," *Proc. Time-Resolved Spectroscopy and Imaging of Tissues, SPIE*, Vol. 1431, pp. 232-238, 1991.
19. Arridge, S.R., P. van der Zee, M. Cope and D.T. Delpy, "Reconstruction methods for infrared absorption imaging," *Proc. SPIE*, Vol. 1431, pp. 204-215, 1991.
20. Graber, H.L., R.L. Barbour, J. Lubowsky, R. Aronson, B.B. Das, K.M. Yoo and R.R. Alfano, "Evaluation of steady-state, time- and frequency-domain data for the problem of optical diffusion tomography," *SPIE*, Vol. 1641, 1992.
21. Muller, G. et al., "Medical Optical Tomography: Functional Imaging and Monitoring. SPIE Optical Engineering Press, Bellingham, WA, 1993, p. 5.
22. Webster, J. G., ed., *Electrical Impedance Tomography*, Adam Hilger, New York, 1990.
23. Ishimaru, A., *Wave Propagation and Scattering in Random Media, Volume 1: Single Scattering and Transport Theory*. Academic Press, New York, 1978.
24. Graber, Chang, Aronson, and Barbour, "A Perturbation Model for Imaging in Dense Scattering Media: Derivation and Evaluation of Imaging Operators," *Medical Optical Tomography: Functional Imaging and*

Monitoring. SPIE Optical Engineering Press, Bellingham, WA, 1993,p. 121-143.

25. Berg, Andersson-Engels, and Svanberg, "Time-resolved Transillumination Imaging," *Medical Optical Tomography: Functional Imaging and Monitoring*. SPIE Optical Engineering Press, Bellingham, WA, 1993, pp. 397-424.
26. Yoo, Das, Liu, and Alfano, "Ultrashort laser pulse propagation and imaging in biological tissue and model random media - steps towards optical mammography," *Medical Optical Tomography: Functional Imaging and Monitoring*. SPIE Optical Engineering Press, Bellingham, WA, 1993, pp. 425-449.
27. Izatt, Hee, Huang, Fujimoto, Swanson, Lin, Schuman, and Puliafito, "Optical Coherence Tomography for Medical Diagnostics," *Medical Optical Tomography: Functional Imaging and Monitoring*. SPIE Optical Engineering Press, Bellingham, WA, 1993, pp. 450-472.
28. Hebden, "Time Resolved Imaging for the Detection of Breast Disease," *Medical Optical Tomography: Functional Imaging and Monitoring*. SPIE Optical Engineering Press, Bellingham, WA, 1993, pp. 473-480.
29. Sevick, Burch, Frisoli, Hohnson, Nowaczyk, Szmacinski and Lakowicz, "The physical basis of biomedical optical imaging using time-dependent measurements of photon migration in the frequency-domain," *Medical Optical Tomography: Functional Imaging and Monitoring*. SPIE Optical Engineering Press, Bellingham, WA, 1993, pp. 485-512.
30. Arridge, S., "The Forward and Inverse Problems in Time Resolved Infra-Red Imaging," *Medical Optical Tomography: Functional Imaging and Monitoring*. SPIE Optical Engineering Press, Bellingham, WA, 1993, pp. 35-64.
31. Barbour, R. et al., "A Perturbation Approach for Optical Diffusion Tomography Using Continuous-Wave and Time-Resolved Data," *Medical Optical Tomography: Functional Imaging and Monitoring*. SPIE Optical Engineering Press, Bellingham, WA, 1993,p.87-120.
32. Arridge, R., "Photon-measurement density functions. Part I: Analytical Forms," *Applied Optics*, Vol. 34, No. 31, Nov. 1995, pp. 7395–7409
33. Arridge, R., "Photon-measurement density functions. Part II: Finite-Element-Method Calculations," *Applied Optics*, Vol. 34, No. 34, Dec. 1995, pp. 8026–8047.