

# **T4DT: Processing 4D CT scans of the Lungs**

Robert Fowler   Joe Warren   Yin Zhang  
Rice University

with Thomas Guerrero

M.D. Anderson Cancer Center, Department of Radiation Oncology  
and Bradley Broom

M.D. Anderson Cancer Center, Department of Biostatistics & Applied Mathematics

A proposal submitted to  
**2005 GC4R program for Collaborative Advances in Biomedical Computing (C-ABC)**

**Start and end dates:** - June 1, 2006 - May 21, 2008

**Amount Requested:** - \$220,000

**Project Summary:** We propose to develop information technology for processing time-varying CT scans (4D CT) of the lungs. In particular, we propose to develop tools for estimating the ventilation (air flow) and perfusion (blood flow) from 4D CT and apply these tools to quantitatively assess the effectiveness of current treatments for lung cancer.

# T4DT: Processing 4D CT scans of the Lungs

Robert Fowler   Joe Warren   Yin Zhang  
Rice University

with Thomas Guerrero

M.D. Anderson Cancer Center, Department of Radiation Oncology  
and Bradley Broom

M.D. Anderson Cancer Center, Department of Biostatistics & Applied Mathematics

## 1 Project Description and Objectives

During the last few decades, advances in 3D medical imaging has enabled significant progress in the diagnosis and treatment of disease. One of particular imaging technology, computed tomography (CT), generates 3D images corresponding to a density estimate of the imaged tissue. The 3D images produced by CT are crucial in the diagnosis and treatment of various diseases including many types of cancer.

One obstacle to the computation of accurate CT images is patient motion. In particular, if the tissue being imaged lies in the thorax, the patient's breathing cycle naturally displaces this tissue during image acquisition and reduces the accuracy of the resulting 3D image. In response to this problem, recently developed CT scanners have faster data-acquisition times and are capable of imaging moving tissue 3 to 4 times a second. The resulting *4D CT scan* consists of a time-varying stack of 3D images.

While this new generation of CT scanners can produce 4D CT scans, the information technology necessary to exploit the data contained in these 4D images is mostly lacking. For example, image density variations in the sequence of 3D images of a patient's lungs during a breath cycle are due to changes in the amount of air and blood in the lungs. However, current software used for registering and analyzing these images is very slow and it is not robust enough to automatically analyze these images and determine even basic physiological properties such as ventilation (air flow) and perfusion (blood flow). In current practice, obtaining these quantities is very expensive, both in terms of computational resources and time spent by a human expert.

In response to this challenge, we propose to develop the basic information technology necessary for processing 4D CT scans of the lungs taken during a single breath cycle. The algorithmic aspects of this effort fall in two broad areas: segmentation and registration. The segmentation portion of the project, lead by Professor Warren, will focus on building a 3D deformable model of the lungs and fitting this model to each 3D image in the 4D CT scan. The registration portion of the project, led by Professor Zhang, will focus on extracting the 3D deformation that the lung tissue undergoes between each 3D image of the breath cycle. Dr. Fowler will work with the algorithmic teams and with Dr. Broom on computational performance issues to ensure that the resulting programs can be used productively in both research and clinical settings. This work will be done in close conjunction with Dr. Thomas Guerrero, to ensure that his technology will be applicable to a specific cancer related problem identified by Dr. Guerrero, determining the ventilation and perfusion of the lungs based on 4D CT.

### 1.1 Segmentation of 4D CT scans of the lung (Warren)

One of the central challenges in processing 4D CT scans of the lungs is extracting the 3D shape of the lungs and their various lobes from each 3D image. Such processing allows subsequent computations




Figure 1: Construction of 3D surface models from 2D cross-sections.

(such as registration) to narrow their spatial domain to smaller, segmented portions of the images. Traditionally, construction of a 3D model of the lungs has been done on a high resolution 3D CT scan taken of a patient while holding their breath [1, 2]. However, the resolution of each 3D scan in a 4D CT scan is much lower (around 1mm) due to the need to bound the total radiation exposure to a patient. (Essentially, a 4D CT scan trades spatial resolution for temporal resolution.)

We propose to build a high resolution 3D model of the lungs from a single high-resolution 3D CT scan. This model will be constructed by manually contouring each 2D horizontal image of the 3D scan into a disjoint set of closed regions corresponding to lobes (figure 1.1(a) and then connecting these curves to form a 3D surface model of the lungs. Specifically, we propose to use a contour following method [3] developed by Professor Warren during previous work on building a 3D atlas for the mouse brain (see [geneatlas.org](http://geneatlas.org) for an overview). Figures 1.1(b) shows a partitioning of a 2D sagittal cross-section of the mouse brain into anatomical regions and connection of a sequence of such contours into a partial 3D surface model (figure 1.1).

Given a high-resolution 3D model of the lungs, our main challenge is to deform this 3D model onto each 3D image in the 4D CT scan. Again, we propose to develop a method based on our experience with fitting a 3D deformable model of the mouse brain to 3D images of sectioned mouse brains [4]. This method consists of two steps. First, we fit the outer surface of the lung model to the boundary between the chest wall and the lungs. Note that this boundary is easily distinguished in figure 1.1. The main benefit of this initial fit is to orient and normalize the size of the lung model to a specific 3D image.

Next, we fit the internal surface boundaries of the lung to low-resolution 3D images. In practice, we believe that this step will be the most challenging since the boundaries of the lobes are only faintly visible in low-resolution scans. Again, we plan to exploit our previous experience with the mouse brain project. In that project, we developed a new hybrid image processing techniques based on texture variation that previously has been used to distinguish the boundaries of anatomical regions interior to the mouse brain [5, 6].

## 1.2 Registration of 4D CT scans of the lungs (Zhang)

The result of our segmentation method is a sequence of 3D surface models fit to each image in the 4D CT scan. Our next task is to compute 3D deformation on internal lung tissue through image registration. Among many utilities, the registration results can be used to localize the physiological functions of the lungs, such as perfusion and ventilation [7, 8], and to aid the subsequent treatment. Given the obvious implications to cancer patients, a main challenge to our registration method




Figure 2: Displacement fields generated by the optical flow and the new method for a pair of simple compressible images: artifacts of the optical flow has been avoided by the new method.

is that it must possess a much higher degree of fidelity than what is normally required in other applications such as computer vision. This high-fidelity requirement necessitates the use of temporal information present in 4D CT scans, leading to a 4D image registration problem.

Though image registration has been studied extensively for the past thirty-five years, a vast majority of existing works are based on the "brightness constancy" principle, first proposed by Horn and Schunk in their original optical flow method [9]. Put simply, these methods assume that a pixel will maintain a constant intensity while it moves from image to image. However, in lung CT images, pixel intensities correspond to the density of the material being imaged. Because local density can vary significantly during breathing cycles, the image intensity of that structure also varies. For these "compressible" images, traditional methodologies cannot provide high-fidelity results. Although some "compressible" registration methods [10, 11, 12, 13, 14, 15] have been proposed, they are far less studied and developed than their "incompressible" counterparts.

Dr. Guerrero [8] currently employs an optical flow method based on the incompressibility assumption. As useful as it is, the method produces undesirable artifacts for lung images. Our objective is to collaborate with Guerrero and other PI's in designing and implementing a high-fidelity "compressible" image registration method for breathing lungs that possesses three basic properties: (a) robust with respect to noise, (b) robust with respect to large displacements, and (c) computationally efficient, reliable and parallelizable. Given the conflicting nature of these goals, a careful and thorough collaborative investigation is necessary to find optimal balances.

We have conducted proof-of-concept studies on a new "compressible" image registration method that integrates ideas of Horn and Schunk [9], Lucas and Kanade [16], Wildes *et al* [17], and Weickert *et al* [18]. It is based on the principle of *mass conservation* rather than that of *intensity constancy*, and is robust with respect to noises. The use of mass conservation is justified by Guerrero's observation (private communication) showing that the violation to mass conservation during a breathing cycle is about 5-10% and can be alleviated by a correction procedure. Preliminary numerical results indicate that on compressible images our new method indeed produces registration results of higher fidelity than existing methods such as optical flow (see figure 2).

A straightforward extension of 3D compressible image registration to 4D would incur enormous computational burdens (e.g., a  $512 \times 512 \times 100 \times 10$  resolution leads to around  $0.75 \times 10^9$  variables). We will investigate parametric modeling for dimensionality reduction, effective numerical schemes and parallel computation for the 4D compressible image registration problem. A combined strategy of all these approaches will be necessary. Formulating and implementing such a combined strategy

will require close collaborations between different groups. For instance, based on Guerrero’s observation (private communication), we believe that it is justifiable to model the direction of a voxel velocity as being time-invariant.

### 1.3 High-performance analysis of 4D CT scans (Fowler)

Our initial involvement with Dr. Guerrero was focused on speeding up his existing compressible registration algorithm. Previously, this method required days to register a single pair of high-density 3D CT images. By applying code and data restructuring techniques, we realized a speed improvement of around 6.5, and by moving to faster computer systems, we got another factor of about 2. In the future, we estimate that a simple parallelization scheme could get one more order of magnitude, but with a low (30%) parallel efficiency. (The range of deformation between pairs of 3D images severely constrains performing the spatial decomposition needed to get performance scaling on today’s commodity clusters). While these improvements have greatly improved the productivity of the optical flow code as a research tool, we observe that it is still too slow for use in a clinical setting.

This performance limitation, combined with the the inability of the current method to produce artifact-free results that can be used automatically to compute ventilation and perfusion images, motivates our shift to exploring then new segmentation and registration algorithms discussed above. A key component of the effort, therefore, is to ensure that the software that embodies these methods will perform well enough be used to improve productivity in a clinical setting. In addition to demonstrating high performance in institutions such as Rice and M. D. Anderson, the methods must be usable in regional hospitals with modest computational resources.

We are concerned with two aspects of performance: on node, or single processor behavior; and parallel efficiency.

Broadly speaking, inefficient use of the memory hierarchy is the primary cause of on-node performance problems. Typically, main memory latency is at least 100 nanoseconds, a period during which a processor could have executed many hundreds of instructions. It is therefore vital that programs and data be structured to minimize the effects of memory access as much as possible. The goals of this (re-)structuring are to: (a) minimize the number of cache misses at all levels; (b) use pre-fetching and code scheduling to hide the latency of memory operations; and (c) reduce the total number of memory operations needed.

In some problem domains, it is relatively easy to improve memory hierarchy performance. Scientific programs, including biomedical visualization, however, present more difficulty because they use large data sets iteratively, and most of the data is touched on every iteration. Methods for improving the memory behavior of such programs are well known [19] in the compiler community. When these methods are potentially used in an optimizing compiler, seemingly small changes in a user’s program may inhibit their application. Since these optimizations are expressible in most source languages, it is also possible to apply them by hand as we did in the optical flow code. While this may be useful for speeding up a particular production code, it is tedious, error-prone, and results in a less understandable and maintainable program.

Performance, like correctness, is a global property of a software system. Just as a single error anywhere in the code can make a program incorrect, local performance problems can wreck overall performance and small changes to a program may produce large changes in its potential parallel efficiency. As with correctness, it is a lot easier to design and test for performance from the beginning of a project than to rely on *ex post facto* changes to fix an existing problem. *Performance must be engineered into the code from the beginning.*

Based on these observations, we propose a four-part strategy for ensuring that this project will result in high performance implementations.

- First, we will work closely with all of the participants in the project to ensure that performance instrumentation and analysis is an integral part of the design and development process. In particular, we will use HPCToolkit [20, 21] performance tools developed in HiPerSoft. We will deploy other performance tools as needed.
- Second, we will work closely with the algorithms groups to design data structures and code to work well on today's systems. This means being able to get good data reuse to exploit the system memory hierarchy and to structure the code to be able to exploit the high levels of instruction level parallelism provided by today's systems. A key aspect of this is that we will use the results of performance instrumentation and analysis to guide the decision process and to evaluate our progress in getting vendor compilers to work well on our code.
- Third, in cases in which vendor compilers are not able to generate performant code, we will explore new compiler-based optimization techniques to specifically address the problems we encounter in this project. For example, our preliminary work on the compressible registration code is used to motivate work on optimizations that enable the use of multimedia instruction sets in scientific applications [22]. Such algorithms can be embedded in source-to-source optimization tools, such as LoopTool [23], that can be used to perform complex transformations, while still allowing us to develop and maintain much more readable source code.
- Fourth, we will plan from the outset to make our methods work well in parallel on low-cost computational clusters built from commodity components. It is important that the methods be parallelizable because clinical success will require the ability to complete the computations within a few minutes. By targeting commodity clusters, the methods will be economically viable in a wider range of settings.

Professor Broom will work closely with the Rice computer scientists on algorithm development and implementation. He will also provide the on-site computer science presence for the project at M. D. Anderson.

## 2 Project Innovation and Impact

By most estimates, more than 175,000 Americans will be diagnosed with lung cancer during 2005. During that same year, over 150,000 Americans will die of lung cancer. The societal impact of our proposal lies directly in its clinical applications. Any advance in the diagnosis and treatment of lung cancer has the potential to affect thousands of lives. To assess this potential impact, we first briefly review the state of the art in 3D imaging technology for the lungs. Next, we argue how our proposed method for determining perfusion and ventilation from 4D CT will directly affect treatment for lung cancer.

### 2.1 Current state of the art

The need for 3D imaging methods suitable for use with the lungs has long been recognized. Both SPECT (single photon emission computed tomography) and MR (magnetic resonance) imaging have been used to evaluate ventilation in the lungs. For both imaging techniques, the patient

inhales a tracer gas (typically radioactive) and a 3D image is then computed using specialized equipment.

Unfortunately, these imaging techniques have several drawbacks. First, they require inhalation of a radioactive tracer that may accumulate in the major airways and distort the resulting images. Second, the imaging techniques have only limited spatial resolution (typically around 2cm). Third, the resulting images have no interpretation in terms of direct physiological values (such as ventilation or perfusion). Finally, the use of these images with CT-derived quantities such as radiotherapy radiation dose or tissue-density distributions requires multi-modal image registration, a difficult task and a potential source of error.

In contrast, 4D CT scans do not use a tracer gas, have higher spatial resolution (around 1mm) and gives access to direct physiological values based on tissue density. Moreover, there is no need for multi-modal registration since all data is CT-based.

## 2.2 Innovative aspects of the proposal

Much of the groundwork of this proposal focus on two traditional tasks in medical image processing: segmentation and registration. For both of these tasks, the innovative aspect of this proposal simply lies in adapting existing techniques and developing new techniques for processing 4D CT scans. However, the truly innovative aspect of this proposal lies in its clinical application: determining perfusion and ventilation from 4D CT scans.

We believe that this goal is eminently achievable. Previously, we have developed a method that measures the distribution of air from a single breath (ventilation) from an exhale-inhale pair of breath-hold CT images acquired without contrast for patients [7] and small animals such as rats [24]. In other preliminary work, we have also applied this pairwise technique directly to a sequence of 3D images forming a 4D CT scan [25, 26] and computed rough estimates of ventilation during the breath cycle for a patient's lungs.

Achieving the ability to determine ventilation and perfusion directly from a 4D CT scan would have a significant impact on the treatment on lung cancer. For example, one of the most common complications of radiation treatment of lung tumors is incidental damage to healthy lung tissue. Being able to spatially assess changes in the perfusion and ventilation of a patient's lungs during the course of radiation treatment would allow for quantitative evaluation of the efficacy of radio-protective drugs that reduce this incidental damage [27, 28, 29, 30].

We also believe that our new method would be very beneficial for pre-surgical or pre-radiotherapy pulmonary evaluation. Because our method does not require the use of a radioactive tracer, a 4D CT scan could be computed both before and immediately after a stress test to compute the lung reserve available to patient at rest. This knowledge could then be used to determine a level of lung function loss that would be acceptable as a consequence of resection or radiotherapy.

## 2.3 Interdisciplinary aspects of the proposal

The research described in this proposal would bring together four world-class experts with widely varying skills:

- Thomas Guerrero is a physician at M. D. Anderson and is a leading expert on applying 4D CT to the treatment of lung cancer.
- Professor Yin Zhang is a professor in the CAAM department whose research expertise includes applications of optimization methods to medicine and biology.

- Professor Joe Warren is a professor in Computer Science. He is one of the world’s leading experts on deformable modeling and has published numerous papers in the area.
- Dr. Rob Fowler is a member of the Center for High Performance Computing (HiPerSoft) at Rice and an expert in computer architectures, performance analysis, and high performance compilers.

### 3 Local Impact and Prospects for Funding

This proposal envisions the emergence of a joint research group at Rice that combines the expertise of three researchers from historically separate entities at Rice; CAAM, CS and HiPerSoft. Combining the optimization expertise of Dr. Zhang, the modeling expertise of Dr. Warren and the computing expertise of Dr. Fowler will yield a research group with the skills to tackle a wide range of biological applications.

Previously, Dr. Guerrero and Dr. Zhang have received some limited funding from an NSF VIGRE grant. This money supported their initial research on compressible image registration for pairs of 3D CT images. However, during the course of their research, they realized that the complexity and difficulty of the problem that they were facing required expertise outside their domain. Dr. Guerrero and Dr. Zhang then recruited Dr. Fowler and Dr. Warren to provide those missing skills. This proposal is the first step in our joint plan for proceeding with this project.

**Our rationale for GC4R funding:** Our ultimate goal is to secure an R01 grant from the NIH to provide long-term funding for this project. In our experience, securing such funding requires several important ingredients:

- A strong biological PI
- An important clinical application
- Extensive preliminary work

We believe that that two of these three ingredients are already in place. As our biological PI, Dr. Guerrero has world-class credentials in lung cancer and has successfully obtained R01 funding in several other projects. Our clinical application, lung cancer, is one of the leading killers of Americans. The advent of 4D CT provides an opportunity to provide more quantitative evaluation of current treatments for lung cancer. We believe that the NIH will view developing information technology to exploit this opportunity as being critical.

However, we lack one critical ingredient for success in establishing NIH funding: strong preliminary work. While the previous work between Dr. Guerrero and Dr. Zhang is a start, we have several major obstacles to overcome before we are ready to submit an NIH grant. The current registration method is extremely slow. We believe that localizing the registration problem through segmentation and applying high-performance computing technology will improve the speed of the registration method so that it can be applied in clinical applications. To successfully compete for NIH funding, we must also validate our preliminary work. In particular, we must demonstrate that our registration method can be used to accurately compute perfusion and ventilation from 4D CT scans.

To advance our preliminary work to a state where an NIH proposal has a reasonable chance of success will entail a commitment of at least a year’s worth of effort by Dr. Warren and Dr. Fowler.

Given that most successful NIH proposals go through at least 1 rounds of rejection and revision, we feel the two year scope of the proposed funding is reasonable.

Funding this work via GC4R provides vital support for Dr. Warren and Dr. Fowler to engage in preliminary work during this time period. Dr. Warren's only current research funding expires at the end of this summer. Without this support, he will be unable to participate in this project.<sup>1</sup> While the HiPerSoft group has done some unfunded exploratory work that has already resulted in significant improvement in the performance of the existing registration code, GC4R funding for this project will enable Dr. Fowler (and a student) to devote significant amounts of this time to the collaboration.

## 4 Management Plan

With support of GC4R, we propose to solidify and expand the joint research group described here. At the heart of the current collaboration are the spur of the moment, joint research meetings that Dr. Zhang, Dr. Warren and Dr. Fowler hold on the 3rd floor of Duncan Hall. (Our offices all lie within 100 ft of each other.) We intend to continue holding these impromptu meetings and supplement these meetings with a weekly lunch with Dr. Guerrero and Dr. Broom.

With the student support in this proposal, we intend to recruit two CS Ph.D. students to become involved in this project. One student will work with Dr. Warren on modifying and improving the segmentation methods developed during the course of Dr. Warren's research on the mouse atlas project. Dr. Fowler will use student funding to support parts of several students, each of whom will work on applying a particular technique from high performance computing to methods developed by Dr. Zhang and Dr. Warren. Dr. Zhang will continue to leverage the NSF VIGRE funding for his Ph.D. Student, Edwards Costillo, to work on this project while increasing his own summer commitment.

To foster interactions between PIs and the students, we intend to hold a semi-weekly research seminar in which one participant will give a short presentation concerning their current work or review a related method developed by a group outside Rice. We view these regular meetings as a key method for stimulating future interdisciplinary research since we hope to attract Ph.D. student outside our immediate research group to these meeting.

We intend to follow the rough **work plan** below:

**Summer/Fall 2006** Construct initial 3D model of lungs (Warren). Complete development of compressible image registration method (Zhang).

**Winter/Spring 2007** Develop 3D segmentation method (Warren). Improve performance of registration method (Zhang, Fowler). Prepare NIH proposal (All).

**Summer/Fall 2007** Enhance performance of 3D segmentation method (Warren, Fowler). Extract perfusion and ventilation from registered models (Zhang). Identify other clinical applications to lung cancer (Guerrero).

**Winter/Spring 2008** Validate method for extracting perfusion and ventilation (Guerrero, Zhang). Apply technology to other clinical applications (Warren, Fowler). Prepare revised NIH proposal if necessary (All).

---

<sup>1</sup>Dr. Warren's previous NIH proposals for the mouse atlas project were declined due to the fact that his biological collaborator, Gregor Eichele, moved to Germany and has become the director of an equivalent project at the Max Planck Institute.

## References

- [1] J. M. Kuhnigk, V. Dicken, S. Zidowitz, L. Bornemann, B. Kuemmerlen, S. Krass, H. O. Peitgen, S. Yuval, H. H. Jend, W. S. Rau, and T. Achenbach, “Informatics in radiology (infoRAD): new tools for computer assistance in thoracic CT. Part 1. Functional analysis of lungs, lung lobes, and bronchopulmonary segments,” *Radiographics*, vol. 25, pp. 525 – 536, Mar. 2005.
- [2] D. Aykac, E. A. Hoffman, G. McLennan, and J. M. Reinhardt, “Segmentation and analysis of the human airway tree from three-dimensional X-ray CT images,” *IEEE Transactions on Medical Imaging*, vol. 22, pp. 940 – 950, Apr. 2003.
- [3] T. Ju, I. Kakadiaris, M. Bello., J. Warren, G. Eichele, C. Thaller, W. Chiu, and J. Carson, “Building 3d surface networks from 2d curve networks with application to anatomical modeling,” in *Pacific Graphics*, 2005.
- [4] T. Ju, J. Warren, G. Eichele, C. Thaller, W. Chiu, and J. Carson, “A geometry database for gene expression data,” in *Eurographics Symposium on Geometry Processing*, pp. 166 – 176, 2003.
- [5] I. Kakadiaris, M. Bello., S. Arunchalam, W. Kang, T. Ju, J. Warren, G. Eichele, C. Thaller, W. Chiu, and J. Carson, “Landmark-driven, atlas-based segmentation of mouse brain tissue images containing gene expression data,” in *7th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCIA)*, pp. 192–199, 2004.
- [6] I. Kakadiaris, M. Bello., T. Ju, J. Warren, G. Eichele, C. Thaller, W. Chiu, and J. Carson, “Hybrid segmentation framework for tissue images containing gene expression data,” in *8th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCIA)*, pp. 254–261, 2005.
- [7] T. M. Guerrero, K. Sanders, J. Noyola-Martinez, E. Castillo, Y. Zhang, R. Tapia, R. Guerra, Y. Borghero, and R. Komaki, “Quantification of regional ventilation from treatment planning CT,” *International Journal of Radiation Oncology, Biology, Physics*, vol. 62, pp. 630 – 634, July 2005.
- [8] T. Guerrero, K. Sanders, J. Noyola-Martinez, E. Castillo, Y. Zhang, R. Tapia, R. Guerra, Y. Borghero, and R. Komaki, “Quantification of regional ventilation from treatment planning ct,” *Int J Radiat Oncol Biol Phys*, 2005.
- [9] B. K. Horn and B. G. Schunck, “Determining optical flow,” *Artificial Intelligence*, vol. 17, pp. 185–203, 1981.
- [10] T. Corpetti, E. Memin, A. Santa-Cruz, D. Heitz, and G. Arroyo, “Optical estimation in experimental fluid mechanics,” *Electronic imaging*, vol. 1, pp. 347–352, 1998.
- [11] T. Corpetti, E. Memin, and P. Perez, “Dense estimation of fluid flows,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 3, pp. 365–380, 2002.
- [12] J. Fitzpatrick, “A method for calculating velocity in time dependent images based on the continuity equation,” *Proc. Conf. Computer Vision Pattern Recognition*, pp. 78–81, 1985.
- [13] V. Devlaminck, “A functional for compressible or incompressible elastic deformation estimation,” *IEEE Signal Processing Letters*, vol. 6, no. 7, pp. 162–164, 1999.

- [14] M. Qiu, “Computing optical flow based on the mass-conserving assumption,” *International Conference on Pattern Recognition*, vol. 3, pp. 7041–7045, 2000.
- [15] H. Haussecker and D. Fleet, “Computing optical flow with physical models of brightness variation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 6, pp. 661–673, 2001.
- [16] B. Lucas and T. Kanade, “An iterative image registration technique with an application to stereo,” *Proceedings of DARPA IU Workshop*, pp. 121–130, 1981.
- [17] R. Wildes, M. Amabile, A.-M. Lanzilloto, and T.-S. Leu, “Recovering estimates of fluid flow from image sequence data,” *Computer Vision and Image Understanding*, vol. 80, pp. 246–266, 2000.
- [18] A. Bruhn, J. Weickert, and C. Schnorr, “Lucas/kanade meets horn/schunck: Combining local and global optic flow methods,” *International journal of Computer Vision*, vol. 61, no. 3, pp. 211–231, 2005.
- [19] K. Kennedy and R. Allen, *Optimizing Compilers for Modern Architectures*. San Francisco: Morgan-Kaufmann Publishers, 2002.
- [20] J. Mellor-Crummey, R. Fowler, and G. Marin, “HPCView: a tool for top-down analysis of node performance,” in *Proceedings of the Los Alamos Computer Science Institute Second Annual Symposium*, (Santa Fe, NM), Oct. 2001. Distributed on CD-ROM.
- [21] J. Mellor-Crummey, R. Fowler, and D. Whalley, “Tools for application-oriented performance tuning,” in *Proceedings of the International Conference on Supercomputing (ICS1001)*, (Sorrento, Italy), pp. 154–165, June 2001.
- [22] Y. Zhao and K. Kennedy, “Scalarization on short vector machines,” in *Proceedings of the 2005 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2005.
- [23] A. Qasem, J. Mellor-Crummey, and K. Kennedy, “Automatic tuning of whole applications using direct search and a performance-based transformation system,” in *Proceedings of the LACSI Symposium*, 2004. Published on CD, available at [http://lacsi.rice.edu/symposium/agenda\\_2004](http://lacsi.rice.edu/symposium/agenda_2004), revised version to appear in *Journal of Supercomputing*.
- [24] T. M. Guerrero, R. Castillo, and K. S. *et al.*, “Novel method to calculate pulmonary compliance images in rodents from computed tomography acquired at constant pressures,” *Physics in Medicine and Biology*, 2005. In press.
- [25] T. Pan, T. Y. Lee, E. Rietzel, and G. T. Chen, “4D-CT imaging of a volume influenced by respiratory motion on multi-slice CT,” *Medical Physics*, vol. 31, pp. 333 – 340, Feb. 2004.
- [26] T. M. Guerrero, K. Sanders, and E. C. *et al.*, “Dynamic ventilation imaging from four-dimensional computed tomography,” *Physics in Medicine and Biology*, 2005. In press.
- [27] M. Saunders, S. Dische, A. Barrett, A. Harvey, D. Gibson, and M. Parmar, “Continuous hyperfractionated accelerated radiotherapy (CHART) versus conventional radiotherapy in non-small-cell lung cancer: a randomised multicentre trial. CHART steering committee,” *Lancet*, vol. 350, pp. 161 – 165, July 1997.

- [28] J. G. Rosenman and M. A. Socinski, “Does more aggressive therapy improve outcomes in the treatment of unresectable stage III non-small cell lung cancer?,” *Seminars in Oncology*, vol. 32, pp. ”S13 – S17”, Apr. 2005.
- [29] J. G. Rosenman, J. S. Halle, M. A. Socinski, K. Deschesne, D. T. Moore, H. Johnson, R. Fraser, and D. E. Morris, “High-dose conformal radiotherapy for treatment of stage IIIA/IIIB non-small-cell lung cancer: technical issues and results of a phase I/II trial,” *International Journal of Radiation Oncology, Biology, Physics*, vol. 54, pp. 348 – 356, Oct. 2002.
- [30] C. Schaake-Koning, W. van den Bogaert, O. Dalesio, J. Festen, J. Hoogenhout, P. van Houtte, A. Kirkpatrick, M. Koolen, B. Maat, and A. N. A. *et al.*, “Effects of concomitant cisplatin and radiotherapy on inoperable non-small-cell lung cancer,” *New England Journal of Medicine*, vol. 326, pp. 524 – 530, Feb. 1992.