

Volumetric Reconstruction from Multi-Energy Single-View Radiography

Sang N. Le[†]
linsang@nus.edu.sg

Mei Kay Lee[‡]
lee.meikay@gmail.com

Shamima Banu[†]
shamima@comp.nus.edu.sg

Anthony C. Fang[†]
afang@comp.nus.edu.sg

Department of Computer Science
School of Computing
National University of Singapore[†]

Physical Education and Sports Sciences
National Institute of Education
Nanyang Technological University[‡]

Abstract

We address the volumetric reconstruction problem that takes as input a series of orthographic multi-energy x-ray images, producing as output a reconstructed model space consisting of uniform-size mass density voxels. Our approach solves the non-linear constrained optimization formulation problem by constructing a compliant estimate of volumetric distribution, subject to projective and domain constraints, and minimizes variational irregularities. To resolve the inherent ambiguities of single-view formulation, an optional shape model may be introduced to aid the reconstruction process. We demonstrate our method's practical usage as a new in-vivo method for estimating three-dimensional body segmental compositions, and compare its results with those of existing methods.

1. Introduction

A wide range of applications require measurements of properties internal to a physical shape in a convenient and non-invasive way: In airport control, security devices for inspecting luggage on conveyor belts; in manufacturing or food industries, devices for evaluating products without interrupting flow; in biomedical practices, devices for obtaining body compositions as inputs to analysis, rehabilitation or prosthetic reconstruction systems; and in biologically-inspired learning systems, where models incorporate physical properties of the subjects. While magnetic resonance imaging (MRI) and computed tomography (CT) scanners in the medical domain offer possible solutions, they are typically considered impractical in terms of cost and logistics for the above-mentioned applications.

Multi-energy radiographic imaging has seen uses in chemical identification for threat detection systems [7, 19]. Dual-energy imaging is also widely used for estimating body compositions of fat, muscular or bone mineral content. Although they are safe and cost effective, these

methods suffer the limitations of projection-based imaging modalities—the acquired pixels are measures of a complex physical integral and segmentation along the projection direction is generally not possible. Tomographic techniques can recover three-dimensional densities from multi-view images, but incur considerable cost in radiation exposure, scan duration and a more restrained setting.

In this paper, we examine the problem of three-dimensional reconstruction from single-view, multiple-energy radiographic images—insofar as it is possible to do so. The problem is inherently ill-posed, resembling ambiguities in single-view scene interpretation. Thus, a *complete* reconstruction of volumetric content—even in the case of homogeneous substances—is neither possible nor to be expected. However, in applications where partial reconstruction suffices, the methods described in this paper readily provide measurements that surpass the results of present acquisition methods in accuracy and dimensionality.

Our approach estimates the mass distribution internal to an object by discretizing its domain into volumetric elements. The objective is to compute a distribution that is consistent with: (1) the absolute values and (2) the differential values observed in x-ray images, among (3) other constraints that are reasonable to the application domain. We cast the problem as a non-linear constrained minimization problem, to which we implement a least-squares solution comprising steps of iterative back-projection and constraints resolution. A convergent solution is obtained through interleaved stages of Gauss-Seidel iterations. The general idea is related to voxel-based scene representation and object reconstruction techniques, but with fundamental differences in imaging modalities, application domains and methodologies.

We further apply our algorithm to the estimation of the physical compositions of human subjects, and demonstrate that the quantities obtained through the method are empirically superior to the existing alternatives.




Figure 1. Examples of dual-energy x-ray scans and image processing steps. From left, subject with 18% body fat: (a) High-energy image (intensity proportional to mass). (b) Low-energy image. (c) Low-to-high ratio image. (d) Image isolating a single material (bone mineral); (e) Segmentation by material composition: bone (white), lean tissues (red), and fat (yellow). (f) Image enhanced with dual-energy information. Another subject with 33% body fat: (g)-(h) Segmented and enhanced images. The two leftmost images are the raw images. Note that simple thresholds on either input image will not reveal bone mineral content, for instance, at the thorax or the hands.

1.1. Multi-Energy Imaging

Consider a single photon beam of energy level E . As the beam traverses through objects in its path, physical interactions diminish its intensity. With an intervening object of homogeneous material, the source and attenuated intensities are related via the classic attenuation formula:

$$I = I_0 e^{-\mu_m \times \rho \times L} \quad (1)$$

where L is path length (cm), and ρ and μ_m are the object's density (g/cm^3) and mass attenuation coefficient (cm^2/g) respectively. For objects with heterogeneous compositions—the ubiquitous case, the attenuation contribution by each differential element can be expressed as:

$$d(\ln(\frac{I}{I_0})) = -\mu_m \times \rho \times dL \quad (2)$$

where the material properties are now specific to the volumetric element.

Discretizing the object domain by the imaging resolution ($X - Y$) across a proportional depth (Z), and assuming constant material in each volume element, the log-attenuation ratio is a sum of elemental attenuation over the pixel:

$$\ln\left(\frac{I(E, x, y)}{I_0(E)}\right) = \sum_z -\mu_m(E, K) \times \rho(x, y, z) \times dz \quad (3)$$

As it turns out, the mass attenuation coefficient for a substance K (c.f. effective atomic number) is constant at a given beam energy E and can be obtained through experimentation. If the number of substance types is finite and known, an alternate form of the attenuation equation is:

$$\ln\left(\frac{I(E, x, y)}{I_0(E)}\right) = \sum_j -\mu_m(E, K_j) \times \rho(K_j) \times L_j \quad (4)$$

where L_j is the effective thickness of material K_j within the mass column.

It is clear from the equation above that by introducing beams of distinct energies, it is possible (at least in theory) to solve for the effective thickness of each material in a heterogeneous object, provided that we have prior knowledge of the object's composition. This is possible for organic lifeforms and also in manufacturing and food industries, where material constituents are known *a priori*, but more difficult for applications whose object contents cannot be anticipated, such as in luggage inspection. *Henceforth, in this paper, we assume the former circumstances.*

To illustrate, suppose we obtain the log-attenuation values, p_1 and p_2 , corresponding to the attenuation of two beams of energies E_1 and E_2 . From Eq. 4, we have:

$$\begin{pmatrix} p_1 \\ p_2 \end{pmatrix} = \begin{pmatrix} \rho(K_1) & 0 \\ 0 & \rho(K_2) \end{pmatrix} \times \begin{pmatrix} \mu_m(E_1, K_1) & \mu_m(E_1, K_2) \\ \mu_m(E_2, K_1) & \mu_m(E_2, K_2) \end{pmatrix} \begin{pmatrix} L_1 \\ L_2 \end{pmatrix} \quad (5)$$

which is solvable for any two-substance object so long as the constant matrices on the right are non-singular.

Where a mixture of component materials is present, e.g., soft tissues, the number of distinct beam energies can be fewer than the component materials. A classification algorithm is required to identify each pixel by material groups, such as fat+lean or soft tissue+bone, each of which is solvable as in the case of non-mixture materials. Examples of raw and processed images are shown in Figure 1. We refer the interested reader to Pietrobelli et al.'s review article [18] for further details. For the remainder of this paper, it suffices to understand that with multi-energy imaging, it is possible to infer compositional mass estimates for a finite number of pre-determined material types from a single view.

2. Background

Physically-based methods that involve human subjects are common in clinical studies, biomechanical studies [12], biologically-inspired learning [23], simulation systems [20] and advanced animation [15], to name a few. The physical quantities that these systems require are known as *body segmental parameters* (BSP), comprising measures of mass, center of mass, and principle moments of inertia for each body segment [22]. Despite the existing research, obtaining accurate estimates of three-dimensional segmental parameters in living subjects remains a fundamental challenge.

Typically, BSP estimates are obtained from cadaver-based regression models (e.g., [9, 25]). While applying these regression equations is relatively straightforward, the models are typically constructed from small, demographically limited or physically-biased samples of a population. It is widely acknowledged that indiscriminate use of these predictive models can lead to a propagation of unsubstantiable errors [10]. Other methods based on stereophotogrammetry (e.g., [13]) relied on the flawed premise that the mass within a body is uniformly distributed. Some experimental studies compute BSP from MRI and CT (e.g., [8]) despite practicality and safety concerns.

More recently, methods to estimate BSP from X-ray projections have been developed [11]. Besides providing empirical validation that accurate *in-vivo* estimates of BSP are possible, these studies point out substantial errors in the unqualified use of regression equations [10]. Their method is practical and safe, relying on dual-energy scanners that are readily available in sports and clinical facilities. Although the authors acquired dual-energy images, only one of the radiographic image was used, thereby limiting their methods to planar uses. To the best of our knowledge, the use of multi-energy image properties to recover three-dimensional mass, geometry or segmental properties, as described in this paper, has not been discussed previously.

Our methods are not limited to biomechanical purposes. In vision, knowing the physical properties of an object offers an important advantage in predicting object movement dynamics [2]. This advantage can be exploited in physically-based tracking [4] and model-based learning systems [3]. Although computing the physical properties of a rigid object is relatively easy, the problem becomes significantly more complex for non-rigid motion and deformable structures [16, 17] such as human subjects.

3. Volumetric Reconstruction I (with known exterior geometry)

We first consider a volumetric reconstruction algorithm where the object’s geometry is known, and extend it to the unknown case in the next section.

Let the three-dimensional bounding manifold of the ob-

ject \mathbb{G} be aligned with the plane of the acquired images $I \in \mathbb{R}^{2 \times N}$. For dual-energy scans, $N = 2$. We discretize the space such that \mathbb{G} is enclosed within a volumetric domain \mathbb{D} comprising $m \times n \times d$ cells. Each cell, commonly termed a *voxel*, $v = (x, y, z)^T \in \mathbb{D}$, is assumed to have uniform material property.

Let the scalar function $f : \mathbb{D} \rightarrow \mathbb{R}$ denote the unknown mass density field. We constrain the values of f such that the implicit surface $f = 0$ delineates \mathbb{G} :

$$f(v) = \begin{cases} +ve & v \text{ is interior to } \mathbb{G} \\ 0 & v \text{ lies on } \mathbb{G} \\ -ve & v \text{ is exterior to } \mathbb{G} \end{cases} \quad (6)$$

Note that geometric closure and orientability are required of \mathbb{G} —for otherwise, the signed distance function is ill-defined [21].

From the properties of radiographic transfer, a photon beam of sufficiently high energy passes through all matter and is attenuated linearly by the amount of mass in its path [14]. The associated pixel is thus proportional to the mass integral along the depth. Assuming unit depth (cm) for each cell, the integral of each column is:

$$I(x, y) = \sum_{z=0}^d \{ f(x, y, z) \mid f(x, y, z) > 0 \} \quad (7)$$

which confines the summation to voxels interior to the object. Our objective is to estimate f within \mathbb{D} subject to the $m \times n$ constraints of Eq. 7. Referring to Eq. 2, we note that $f(x, y, z) = -\mu_m \times \rho$ is the differential attenuation of the photon beam through a path of unit length.

Since the formulation is under-constrained, a solution to f is not viable without further assumptions or constraints. Further, direct matrix solvers are ill-advised as the volumetric domain is typically large but locally structured.

Drawing from discrete level-set approaches, we proceed to formulate an iterative solution by first initializing f with a signed function f^0 whose value equals the shortest linear distance to \mathbb{G} . The distances are signed such that points interior to \mathbb{G} are positive while those exterior are negative (Eq. 6). Efficient methods of computation exist, e.g., [1], whose details shall be omitted here.

Although a signed distance function is largely expected to be different from a mass distribution function, the boundary condition coincides by design:

$$\mathbb{G} = f^0(p) = f(p) = 0 \quad (8)$$

for all $p \in \mathbb{G}$. To preserve the implicit geometry, we introduce a positive field $r^+ : \mathbb{D} \rightarrow \mathbb{R}^+$ such that:

$$f(x, y, z) = r^+(x, y, z) f^0(x, y, z) \quad (9)$$

where we initialize $r^+(x, y, z) = 1$ so that Eq. 8 and Eq. 9 are initially identical. Since the field r^+ is uniformly positive, the boundary (Eq. 8) and sign constraints (Eq. 6) are unconditionally maintained.

To unify the solution, we introduce a regularization scheme which minimizes total variational energy modeled as the anisotropic diffusion of the multiplier field r^+ :

$$\operatorname{argmin}_{r^+} \int \nabla \cdot r^+(x, y, z) \quad (10)$$

As smooth divergence is not necessarily true of real body masses, the condition only serves to ensure generally non-sporadic variations in the reconstituted volume.

The volumetric reconstruction problem can thus be cast as a non-linear constrained minimization problem: Solve the unknown positive field $r^+ \in \mathbb{D}$ that minimizes Eq. 10 subject to the constraints in Eq. 7, and given the inputs \mathbb{G} , f^0 and I .

We formulate an iterative solution by refining r^+ through three steps. The first step locally resolves the projected mass constraints (Eq. 7) for each pixel by normalizing the column of mass densities through the following pair of computations:

$$I'(x, y) = \sum_{z \in \mathbb{G}} r^+(x, y, z) f^0(x, y, z) \quad (11)$$

and

$$r^+(x, y, z)' = r^+(x, y, z) \frac{I(x, y)}{I'(x, y)} \quad (12)$$

The next step refines the current estimate (Eq. 12) through its first-order differential. Expressing Eq. 10 in terms of finite differences, we then obtain a strictly diagonally dominant linear system whose solution can be efficiently approximated through a Gauss-Seidel relaxation scheme:

$$\begin{aligned} \delta r^+(x, y, z)' &= (r^+(x, y, z)' - t(r^+(x-1, y, z)' \\ &\quad + r^+(x, y-1, z)' + r^+(x, y+1, z)' \\ &\quad + r^+(x, y, z-1)' + r^+(x, y, z+1)' \\ &\quad + r^+(x+1, y, z)')) / (1 + 6t) \end{aligned} \quad (13)$$

where t is the relaxation parameter. We typically set t as 0.01 and perform a fixed number of relaxation iterations (8–15) at each update cycle.

Finally, we update the current iterate, projecting it to the positive domain when necessary:

$$r^+(x, y, z) \leftarrow \{r^+(x, y, z)' + \delta r^+(x, y, z)'\}_{\geq 0} \quad (14)$$

4. Volumetric Reconstruction II (without exterior geometry)

We extend the previous framework to the case where the object geometry is unknown. Recall from §1.1 that the effective thickness of materials in heterogeneous objects may




Figure 2. Image pixels plotted as ratios of low-to-high attenuation (Y-axis) against mass (X-axis) for a single subject. The two calibration curves partition the pixels into three zones: Above the upper calibration curve are the pixels that contain bone. Pixels plotted near the bottom curve contain only fat tissues. The remaining bulk lying in between the curves are pixels of lean/fat mixtures.

be obtained by exploiting differences in interactions with beam energies. The idea then is to identify the fractional composition of each material and to make use of estimated material thickness in place of object geometry.

We first segment the image pixels by material composition, making use of the fact that the attenuation coefficients of distinct materials are known constants (see §1.1). Illustrating the process through an example, Figure 2 shows the classification of pixels in a full-body dual-energy image (Figure 1, Subject 1) into compositions of bone, lean and fat. The plot shows pixel ratios (Y-axis) of dual-energy images against projected pixel mass (X-axis). The material calibration curves, experimentally determined using the method described in [14], partition pixels by known compositions. The segmented image is shown in Figure 1(e).

From each pixel’s classification, we estimate each material’s effective thickness using densities that have been experimentally obtained [18]. The thickness of the column of mass at each pixel, $L(x, y)$, can therefore be estimated as the sum of each material’s effective thickness.

It thus remains to position the mass column along Z , denoted $H(x, y)$. We introduce a heuristic region-based operator that first seeks out local symmetries in the image. A method we found relatively resilient to scale and noise is a split-operator scheme based on edge extraction and weighted Principle Component Analysis [24]. High-intensity pixels lying along the symmetry axes are “anchored” to the ground plane: $H(x, y) = 0$. Conversely, low-intensity pixels in directions orthogonal to the symmetry axes and border both background and object (and hence have zero thickness) are assigned a height equal to half the



Stack	Measured (cm)	Computed (cm)	Difference (%)
1	5.5	5.130	-6.727%
2	11.0	10.732	-2.436%
3	16.5	16.643	0.866%
4	22.0	21.333	-3.031%
5	27.5	27.013	-1.771%

Figure 3. Experimental setup for the validation of single-material reconstruction. The scene consists of stacked reams of standard A4-size photocopy paper. The top-left shows the low-energy x-ray image. The volume interior comprises a uniform density material (carbon). The reconstructed scene and the measured versus computed differences are shown. The RMS% difference is 3.582%.



average thickness of their associated symmetry axes pixels. The remaining unprocessed pixels are assigned heights that are relative to the normalized spread of their neighbors' heights.

With an estimate of $\mathbb{G} = (H, L)$, algorithm §3 can be invoked as it is. It is worth noting that the heuristics presuppose errors in finding local symmetries and borders. Further, the constructed geometry is consistent with the class of plane-symmetric shapes. In practice, the scheme also works favorably for generally regular shapes. For scans specific to human subjects in a supine pose, a heuristic that works reasonably well is to ground pixels that contain high bone mineral content, the rationale being such columns of masses are more likely to be in contact with the underlying support.

5. Experimental Results

5.1. Validation

Experiments were carried out to evaluate the viability of estimating the effective thickness of materials using multi-energy imaging. Figure 3 illustrates an example of a single-material experimental setup. Multiple stacks of photocopy



Step	Measured (cm)	Computed (cm)	Difference (%)
Aluminum	0.15	0.148	-1.333%
	0.35	0.354	1.143%
	0.80	0.802	0.250%
Acrylic	3.70	3.737	1.000%
	8.80	8.908	1.227%
	19.80	19.643	-0.793%

Figure 4. Experimental setup for the validation of two-material reconstruction. The scene consists of a six-step device supplied by the scanner's manufacturer. Each step consists of either acrylic or an aluminum/acrylic mixture. The reconstructed scene and the measured versus computed differences are shown. The RMS% difference is 1.023%.

paper reams were arranged in columns of different heights. The paper reams were assumed to compose of a homogeneous material (carbon) with an experimentally determined density of $0.73g/cm^3$. The stacked paper reams were scanned with a dual-energy scanner (Hologic Discovery Model A). Knowledge of scene geometry was not assumed in baseline validations. The image obtained from the low-energy x-ray beam is shown in the figure. The volume was computed with the algorithms described in §4. A similar two-material experiment was carried out using a manufacturer-calibrated device (Figure 4).

In both tests, larger discrepancies were observed (as were expected) near scene edges which can be accrued to the image resolution (pixel size is $1.303\text{ cm} \times 0.618\text{ cm}$) relative to the scene size (approximately $195.5\text{ cm} \times 65.5\text{ cm}$). Increased errors are thus typical of pixels that border both object and air, and relatively more pronounced with smaller-size objects. Scanners with higher image resolution would partly alleviate these errors. On the other hand, differences in reconstructed object heights for the manufacturer-calibrated device (RMS%=1.023) are lower than our stacks of paper reams (RMS%=3.582), possibly indicating the importance of material purity and calibration.

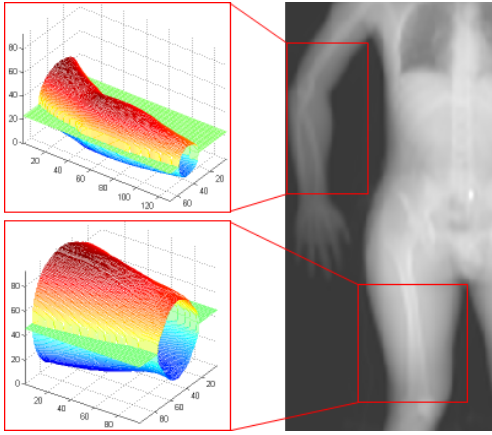


Figure 5. Reconstructed thigh and arm from projected multi-energy images using the algorithm described in §4. They reflect closely the thickness and shape of the actual body segments.

5.2. Human Body Geometry

We applied the algorithm in §4 to reconstruct the geometry for human body segments (Figure 5). Each segment was assumed to consist of bone, lean tissues and fat with densities of $1.6g/cm^3$, $1.06g/cm^3$ and $0.9g/cm^3$ respectively. Bone regions are distinguished from non-bone regions in the images since the attenuation coefficient of bone is distinctly higher. In non-bone regions, the mass fractions of constituent lean tissues and fat were estimated by making use of their attenuation coefficients and low-to-high ratio image. Their thicknesses were then computed using known mass densities and summed up to obtain the total thickness. Bone pixels were resolved in a similar way with two components to be bone and lean/fat mixture. The average density of lean/fat mixture above and below the bone is derived from that in surrounding non-bone neighborhood. The thickness of each bone pixel was the aggregate of the computed bone and lean/fat mixture at that pixel.

Figure 5 shows the three-dimensional geometry of the thigh and arm reconstructed from the projected multi-energy images. Other complex body segments might also be constructed by taking into account additional materials, such as brain tissue in the head and air in the lung.

5.3. Human Body Composition

The mass distribution volumes were computed for estimating body segment parameters. To provide the inputs required for the algorithms, we obtained the dual-energy x-ray images for a group of 18 healthy male subjects; 4 of which also underwent body geometry scanning. Additionally, physical measurements of each subject were taken for purposes of comparisons with the alternative methods. Informed consent was collected from all subjects prior to their participation. The acquired 3D geometry can be registered




Figure 6. Cross-sectional views of the volumetric results at the head, lung, thighs and calves. Each row, from left to right, illustrates the mass distribution estimates after initialization (Eq. 9), normalization (Eq. 12) and Gauss-Seidel relaxation (Eq. 13). The darkness in the head and lung cross-sections corresponds to brain tissue and air, respectively. The bright regions in the thighs and calves indicate the bones in each segment.

with the image using a non-rigid shape registration [6].

Cross-sectional views of the reconstructed volumetric mass density are shown in Figure 6. Different distributions can be seen at the head, lung, thighs and calves. Although the obtained results may not perfectly resemble the real distribution in the human body, they clearly reflect the specific compositions of each body segment.

5.4. Body Segment Parameters

To analyze the proposed method, we also compared our results with available BSP data from alternative methods. Contrary to our approach which is subject-specific and measures segmental mass directly from high-energy x-ray image (Figure 1(a)), the compared methods estimate BSP by assuming known uniform densities, based on statistical studies of cadaver [5, 9], gamma-ray [25] and MRI [8]. Table 1 shows the mean segmental masses of all subjects and the corresponding root mean square deviation against masses estimated from other regression-based methods. Although the variations in segments definition might affect the

Segment	Mass (kg)	RMSD (kg)			
	Vo	DEM [9]	CHA [5]	ZAT [25]	CHE [8]
Head+neck	5.77	1.08	N.A.	0.85	1.08
Trunk	31.06	2.52	N.A.	1.4	2.25
Upper arm	2.3	0.47	0.45	0.42	0.53
Forearm	1.07	0.12	0.16	0.12	0.11
Hand	0.38	0.12	0.13	0.11	0.17
Thigh	9.03	2.4	1.78	N.A.	0.64
Calf	3.08	0.34	0.33	0.21	0.31
Foot	0.9	0.25	0.14	0.17	0.61

Table 1. Comparison of mean segmental mass obtained by our proposed method (Vo) and alternative methods

Segment	Center of mass (%)			Moments of inertia ($kg \times m^2$)						
	Vo	DEM [9]	CHE [8]	Vo			DEM [9] CHE [8]			
				I_{xx}	I_{yy}	I_{zz}	I_{xx}	I_{xx}	I_{yy}	I_{zz}
Head+neck	48.9	43.3	49.6	0.0308	0.0285	0.0172	0.031	0.036	0.037	0.016
Trunk	45.9	N.A.	60.2	0.9832	1.0785	0.2236	N.A.	0.94	0.99	0.18
Upper arm	45.2	43.6	43.4	0.0152	0.015	0.0026	0.022	0.027	0.026	0.003
Forearm	42	43	47.3	0.0071	0.0077	0.003	0.0072	0.0067	0.0066	0.0007
Hand	39	49.4	42	0.0005	0.0006	0.0003	0.0011	0.001	0.001	0.0002
Thigh	42.2	43.3	44.7	0.1268	0.1256	0.0231	N.A.	0.14	0.14	0.017
Calf	39.3	43.3	44.2	0.036	0.0356	0.0041	0.047	0.025	0.025	0.003
Foot	47.4	42.9	54	0.003	0.0007	0.0028	0.004	0.009	0.007	0.009

Table 2. Comparison of centers of mass and moments of inertia obtained by our proposed method (Vo) and alternative methods

computed mass, major differences can be seen across all segments, which is reasonable since our method does not assume any density for the mass distribution computation.

In Table 2, mean values for three principle moments of inertia and center of mass (CM) of each body segment were computed for the 4 subjects with known body geometry and compared between our method and [8, 9]. The segmental CM positions were computed along the main axis of each segment. The moments of inertia were calculated about three anatomical axes defined in [8]. Note that at the foot segment, our method computes a significantly small I_{yy} compared to the corresponding I_{xx} and I_{zz} . This is to be expected because y is the main axis of the foot. However, this main-axis effect was not reflected clearly in [8].

6. Discussion

Typically, the model of an object comprises directly observable aspects such as size, shape, position, motion and texture. A more complete description, however, would include properties such as mass distribution. The importance in modeling these physical attributes is growing, along with increasingly sophisticated methods for processing real-life data, especially where motion is dominated by dynamics. Already, physically-based methods that pertain to processing or understanding human movements are reported in vision, graphics and robotics. Yet despite these advances, the fundamental challenge of estimating the input parameters to those physics algorithms has received little attention.

Our method recovers partial depth structure from x-ray images by exploiting a fundamental principle of materials—that differential photonic attenuation observed at a point reveals the object’s composition along the direction of projection. However, a complete three-dimensional reconstruction of an object’s internal distribution from a single view—for that matter from any limited number of views—is unlikely. Consequently, the majority of work that concerns three-dimensional construction in x-ray imaging solves large systems of algebraic equations formulated from a large number of images, each obtained from different angles. The application of such systems is highly specialized

and has mostly been confined to medical diagnostics.

Multi-energy, single-view x-ray imaging has a distinct advantage that linear “sandwich-type” projectors (without major rotary mechanisms) are possible. Commonly seen ones are those at airport security checkpoints. Among these systems, those that exploit multi-energy image properties are used mainly to detect specific chemicals such as explosives. The ability of these systems to detect structural forms, such as banned carry-on items, remains to be seen. While image-based object detection methods may work for head-on images of a known object, they may not work equally well if the recognizable 2D structures are concealed, by malice or chance, along the direction of projection.

The applicability of our algorithms largely depends on the domain’s tolerance to scale. While it is mostly acceptable to say “the human spine lies central to its trunk”, the statement is meaningless in a close-up view of a spinal section. This is the main limitation of our method. Nevertheless, we have demonstrated the ability to estimate whole-body physical properties in living subjects, with estimates of segmental parameters that surpass the accuracies of conventional, non-subject-specific methods (i.e., regression).

Our method resolves ambiguities of scene interpretation through spatial coherence and a model-based algorithm. To minimize complications to the presentation, we have avoided a more extensive use of domain-specific knowledge, and have focused on information that can be recovered through the nature of multi-energy images. From a computer vision perspective, the ability to reconstruct object models from a single orthographic view may be noteworthy, despite assumptions of shape and material composition. More importantly, the geometry of a human body comprising mixtures of substances can be estimated from a pair of x-ray images.

Extensions incorporating domain-specific knowledge are likely to further improve our results or widen the range of applications. One simple extension is to resolve the spatial sequencing of lean, bone and fat voxels through the anatomical knowledge that fat is always subcutaneous. A more sophisticated approach is to apply the stochastic distribution (e.g., [26]) of the acquired images on variations along the

projected direction. It will be interesting to see these developments, and more generally, the extent to which structure and form can be recovered from multi-energy images.

Acknowledgments

Research supported by grant R-252-000-256-305 from the Science and Engineering Research Council (A*STAR) in collaboration with National Institute of Education (PESS), Republic Polytechnic (SHL) and Singapore Sports School. Further thanks to the support of students and staff of the Computer Graphics Research Lab, National University of Singapore.

References

- [1] J. Baerentzen and H. Aanaes. Signed distance computation using the angle weighted pseudo-normal. *IEEE Trans. Visualization and Computer Graphics*, 11(3):243–253, 2005. **3**
- [2] K. S. Bhat, S. M. Seitz, J. Popovic, and P. K. Khosla. Computing the physical parameters of rigid-body motion from video. In *ECCV, 7th European Conference on Computer Vision*, page 551, 2002. **3**
- [3] C. Bregler. Learning and recognizing human dynamics in video sequences. In *IEEE Conf. Computer Vision and Pattern Recognition*, pages 8–15, June 1997. **3**
- [4] M. A. Brubaker, D. J. Fleet, and A. Hertzmann. Physics-based person tracking using simplified lower-body dynamics. In *IEEE Conf. Computer Vision and Pattern Recognition*, pages 1–8, 2007. **3**
- [5] R. F. Chandler, C. E. Clauser, J. T. McConville, H. M. Reynolds, and J. W. Young. Investigation of inertial properties of the human body. Technical Report 74-137, Aerospace Medical Research Laboratory, Wright-Patterson Air Force Base, Ohio, 1975. **6**
- [6] H. Chen and B. Bhanu. Global-to-Local Non-Rigid Shape Registration. In *18th Intl. Conf. Pattern Recognition*, volume 4, pages 57–60. IEEE Computer Society Washington, DC, USA, 2006. **6**
- [7] Z. Chen, Y. Zheng, B. Abidi, D. Page, and M. Abidi. A combinational approach to the fusion, de-noising and enhancement of dual-energy x-ray luggage images. In *IEEE Conf. Computer Vision and Pattern Recognition*, 2005. **1**
- [8] C. K. Cheng, H. H. Chen, C. S. Chen, C. L. Chen, and C. Y. Chen. Segment inertial properties of Chinese adults determined from magnetic resonance imaging. *Clinical Biomechanics*, 15:559–566, 2000. **3, 6, 7**
- [9] W. Dempster. Space requirements of the seated operator. Technical Report TR-55-159, Wright Air Development Center, Wright-Patterson Air Force Base, Ohio, 1955. **3, 6, 7**
- [10] J. L. Durkin and J. J. Dowling. Body segment parameter estimation of the human lower leg using an elliptical model with validation from dxa. *Annals of Biomedical Engineering*, 34(9):1483–1493, 2006. **3**
- [11] J. L. Durkin, J. J. Dowling, and D. M. Andrews. The measurement of body segment inertial parameters using dual energy x-ray absorptiometry. *Journal of Biomechanics*, 35:1575–1580, 2002. **3**
- [12] K. J. Ganley and C. M. Powers. Determination of lower extremity anthropometric parameters using dual energy x-ray absorptiometry: the influence on net joint moments during gait. *Clin Biomech (Bristol, Avon)*, 19(1):50–6, 2004. **3**
- [13] R. K. Jensen. Estimation of the biomechanical properties of three body types using a photogrammetric method. *Journal of Biomechanics*, 11(8-9):349–358, 1978. **3**
- [14] T. L. Kelly, N. Berger, and T. L. Richardson. DXA body composition: Theory and practice. *Appl. Radiat. Isot.*, 49(5/6):511–513, 1998. **3, 4**
- [15] C. K. Liu, A. Hertzmann, and Z. Popović. Learning physics-based motion style with nonlinear inverse optimization. *ACM Trans. Graph.*, 24(3):1071–1081, 2005. **3**
- [16] D. Metaxas and D. Terzopoulos. Shape and nonrigid motion estimation through physics-based synthesis. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 15(6):580–591, 1993. **3**
- [17] A. Pentland and B. Horowitz. Recovery of nonrigid motion and structure. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 13(7):730–742, July 1991. **3**
- [18] A. Pietrobelli, C. Formica, Z. Wang, and S. B. Heymsfield. Dual-energy X-ray absorptiometry body composition model: review of physical concepts. *Endocrinology and Metabolism*, 271(6):941–951, 1996. **2, 4**
- [19] V. Rebuffel and J. M. Dinten. Dual-energy X-ray imaging: benefits and limits. *Insight - Non-Destructive Testing and Condition Monitoring*, 49(49):589–594, October 2007. **1**
- [20] V. D. Sapio, J. Warren, O. Khatib, and S. Delp. Simulating the task-level control of human motion: a methodology and framework for implementation. *Visual Computer*, 21:289–302, 2005. **3**
- [21] W. P. Thurston. *Three-Dimensional Geometry and Topology*. Princeton University Press, 1997. **3**
- [22] D. Winter. *Biomechanics and motor control of human movement*. Wiley New York, 1990. **3**
- [23] K. Yamane, J. Hodgins, and H. Brown. Controlling a marionette with human motion capture data. In *Proceedings of IEEE International Conference on Robotics and Automation*, volume 3, pages 3834–3841, 2003. **3**
- [24] T. Yuan and X. Tang. Efficient local reflectional symmetries detection. In *IEEE International Conference on Image Processing*, volume 3, pages 1180–3, 2005. **4**
- [25] V. M. Zatsiorsky, V. N. Seluyanov, and L. Chugunova. In vivo body segment inertial parameters determination using a gamma-scanner method. In N. Berme and A. Cappozzo, editors, *Biomechanics of human movement: Applications in rehabilitation, sports and ergonomics*, Worthington, OH: Bertec Corporation, pages 187–202, 1990. **3, 6**
- [26] L. Zollei, E. Grimson, A. Norbash, and W. Wells. 2D-3D rigid registration of X-ray fluoroscopy and ct images using mutual information and sparsely sampled histogram estimators. In *IEEE Conf. Computer Vision and Pattern Recognition*, volume 2, pages 696–703, 2001. **7**