

Dip Transform for 3D Shape Reconstruction

KFIR ABERMAN*, AICFVE Beijing Film Academy, Tel-Aviv University

OREN KATZIR*, AICFVE Beijing Film Academy, Tel-Aviv University

QIANG ZHOU, Shandong University

ZEGANG LUO, Shandong University

ANDREI SHARF, Ben-Gurion University of the NEGEV, AICFVE Beijing Film Academy

CHEN GREIF, University of British Columbia

BAOQUAN CHEN[†], Shandong University

DANIEL COHEN-OR, Tel-Aviv University




Fig. 1. 3D scanning using a dip scanner. The object is dipped using a robot arm in a bath of water (left), acquiring a dip transform. The quality of the reconstruction is improving as the number of dipping orientations is increased (from left to right).

The paper presents a novel three-dimensional shape acquisition and reconstruction method based on the well-known Archimedes equality between fluid displacement and the submerged volume. By repeatedly dipping a shape in liquid in different orientations and measuring its volume displacement, we generate the *dip transform*: a novel volumetric shape representation that characterizes the object's surface. The key feature of our method is that it employs fluid displacements as the shape sensor. Unlike optical sensors, the liquid has no line-of-sight requirements, it penetrates cavities and hidden parts of the object, as well as transparent and glossy materials, thus bypassing all visibility and optical limitations of conventional scanning devices. Our new scanning approach is implemented using a dipping robot arm and a bath of water, via which it measures the water elevation. We show results of reconstructing complex 3D shapes and evaluate the quality of the reconstruction with respect to the number of dips.

CCS Concepts: •Computing methodologies → Shape modeling;

Additional Key Words and Phrases: shape reconstruction, volume, data acquisition

*Joint first authors.

[†]Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 ACM. 0730-0301/2017/7-ART79 \$15.00
DOI: <http://dx.doi.org/10.1145/3072959.3073693>

ACM Reference format:

Kfir Aberman, Oren Katzir*, Qiang Zhou, Zegang Luo, Andrei Sharf, Chen Greif, Baoquan Chen, and Daniel Cohen-Or. 2017. Dip Transform for 3D Shape Reconstruction. *ACM Trans. Graph.* 36, 4, Article 79 (July 2017), 11 pages.

DOI: <http://dx.doi.org/10.1145/3072959.3073693>

1 INTRODUCTION

3D shape acquisition and reconstruction methods are based on optical devices, most commonly laser scanners, that sample the visible shape surface. These scanners yield point clouds which are often noisy and incomplete. Many techniques have been developed to amend such point clouds, denoise them, and complete their missing parts using various priors [Berger et al. 2014]. All these surface-based reconstruction techniques assume a reasonable sampling rate of the surface. Notably, these techniques fall short in cases where the shapes contain highly occluded parts that are inaccessible to the scanner's line-of-sight. Thus, complex shapes such as the three-elephant statue in Figure 2(a) and Figure 2(b) cannot be properly acquired nor reconstructed based on conventional (optical) scanners. Moreover, some objects, like Figure 2(c), are made of glossy or transparent materials, which pose another challenge that common optics cannot deal with.

In this work, we take a completely different approach to shape reconstruction. The idea is to cast surface reconstruction into a volumetric problem. Our technique is based on the ancient fluid displacement discovery made by Archimedes, stating that: *the volume of fluid displaced is equal to the volume of the part that was submerged*. By dipping an object in liquid along an axis, we can measure the

displacement of the liquid volume displacement and transform it into a series of thin volume slices of the shape along the dipping axis. By repeatedly dipping the object in various orientations (see Figure 1), we generate different volumetric displacements and transform them into what we call a “dip transform”.

3D reconstruction proceeds by measuring the volumes of oblique thin slices of the shape. We refer to the volumes of these slices as samples, and collect such samples along different angles. This, in turn, equips us with the ability to generate enough data to recover the geometry of the input shape. Since our technique is based on using volume samples that are generated by liquid accessing the object, we can acquire occluded and inaccessible parts in a relatively straightforward fashion.

Figure 3 illustrates the dip transform of 2D and 3D shapes. Input shapes (top row) are rotated and dipped along different orientations to form the dip transform (bottom row). In the dip transform 2D images, each column represents measurements along one axis and the rows represent volume displacements (or slices). We show that the dip transform is invertible and thus the shape can be reconstructed back from it.

Our technique of reconstructing 3D shapes with the dip transform is related to computed tomography (see Section 2). Common tomography methods use optical systems with penetrating waves to sample the absorption of the radiation along multiple parallel beams. We show that the dip transform is equivalent to the Radon transform for binary volumes. While tomography acquisition devices based on radiation systems are expensive, bulky and must be used in a safe special environment, our dip-based tomography is inexpensive and safe, while capturing hidden parts and structures. As such, it provides an appealing alternative approach, generating a complete shape at a low computational cost, using a novel data acquisition technique.

The inverse transform requires solving an underdetermined problem. The matrices involved are large and sparse and are nearly orthogonal. As such, they have nonzero patterns and structural properties that can be exploited to accelerate the numerical computations. Given a set of samples of a given object, we use a pre-computed factorized matrix to compute the inverse dip transform in nearly linear time in the number of samples and obtain a stable numerical solution to the problem. When the dimensions of the problem are small to medium size, we solve the problem by (implicitly) computing the pseudo-inverse, giving rise to a solution with minimum norm. For very large problems we apply LSQR or LSMR, which are decomposition-free and minimize the norm of the residual

or related quantities over a Krylov subspace, using matrix-vector products.

The key advantage of the presented approach is that it employs liquid as the sensor. Unlike optical sensors, the liquid has no line-of-sight requirements, it penetrates into cavities and hidden parts of the subject, bypassing all visibility and optical limitations of conventional scanning devices. We demonstrate the merits of our approach by presenting various results of reconstructing complex objects, which include large portions of hidden regions.

2 RELATED WORK

The problem of surface reconstruction from scanned points has been researched extensively in the last couple of decades. Nevertheless scan based acquisition and reconstruction are beyond the scope of this work, see e.g. [Berger et al. 2014] for a comprehensive survey.

Computed tomography (CT) is the process of reconstructing an unknown volume from a collection of two-dimensional projections (shadow images) corresponding to different positions of the electromagnetic radiation source and representing the integral of emissions or absorptions from a given view. A classical way of reconstructing volume densities from line integrals captured by the projection images is via the Fourier Slice Theorem [Bracewell 1956; Kak and Slaney 2001; Radon 1917], which states that the volume densities can be recovered from a number of 1D inverse Fourier transforms of lines in the projection images. Alternatively, Algebraic Reconstruction Techniques (ART) iteratively solve a set of linear equations [Gordon et al. 1970; Kak and Slaney 2001] for the volume densities, using the Kaczmarz iterative scheme.

CT plays a central role in medical imaging as well as in many computer graphics related applications. Trifonov et al. [2006] use visible light tomography to acquire objects made of glass and other transparent media. Specifically, they measure the refraction of light through the transparent object and use tomographic reconstruction to reconstruct the object’s volume.

Still, 3D reconstruction of transparent and specular objects has been a challenging problem in computer vision and graphics. Transparent and specular objects, may have complex interior and exterior structures that can reflect and refract light in complex ways. Since it is beyond the scope of this paper, we refer to the state-of-the-art survey of this field [Ihrke et al. 2010].

Stochastic tomography has been introduced in [Gregson et al. 2012] for the purpose of capturing turbulent fluid mixing behavior. Instead of using the Radon transform (Fourier slice theorem) or algebraic reconstruction techniques (ART), their method uses random walks to reconstruct the volume from an array of cameras. Zhao et al. [2011] use an X-ray CT system to capture the inner structure of fabrics and model their appearance. To accurately render a material, a CT scan is used to capture its inner volume which is then used to compute the material’s scattering properties. Recently, Ijiri et al. [2014] use an X-ray CT system to accurately capture the intricate internal structures of real-world 3D flowers. They segment the volume into flower organs and apply active contours to smoothly reconstruct their surface.




Fig. 2. These three elephants statues are examples for complex shapes that cannot be fully reconstructed by conventional (optical) 3D scanners.




Fig. 3. Examples of 2D and 3D objects (first row) and their corresponding sinograms (second row), $\mathcal{G}(t; \theta)$. The 2D sinograms are described in (a)-(c) for the range of $0 \leq \theta < 180$, $\theta^T = (\cos \theta, \sin \theta)$. A 3D sinogram is described in (d), for the range of $0 \leq \theta < 90$, $0 \leq \phi < 360$, $\theta^T = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$.

3 3D DIP TRANSFORM

According to Archimedes equality, which couples buoyancy and volume, dipping an object in a bath of water allows for measuring the volume of slices along the dipping orientation (perpendicular to the water surface) from the water displacement level. Based on this principle, let us define a *3D dip transform* of an object.

3.1 Definition

Suppose we are given an input object, represented as a volumetric binary function,

$$F(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \in \Omega, \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

where $\Omega \in \mathbb{R}^3$ is the object's interior domain. We define the accumulated volume of $F(\mathbf{x})$ over time t , with orientation θ , as the dip transform of the object, given by:

$$\mathcal{F}(t; \theta) = \int_{-\infty}^{\infty} F(\mathbf{x}) u(\mathbf{x}^T \theta - t) d\mathbf{x}, \quad (2)$$

where $u(t)$ is a step function that equals 1 if $t \geq 0$ and 0 otherwise.

The time derivative of the dip transform, $\mathcal{G}(t; \theta) = \frac{\partial \mathcal{F}}{\partial t}(t; \theta)$, is equivalent to the Radon transform of a binary object [Herman 2009]:

$$\mathcal{G}(t; \theta) = \int_{-\infty}^{\infty} F(\mathbf{x}) \delta(\mathbf{x}^T \theta - t) d\mathbf{x}, \quad (3)$$

where $\delta(t)$ is the Dirac delta function. Relying on the Archimedes equality, we get

$$\int \mathcal{G}(t; \theta) dt = V, \quad \forall \theta,$$

where V is the volume of the object.

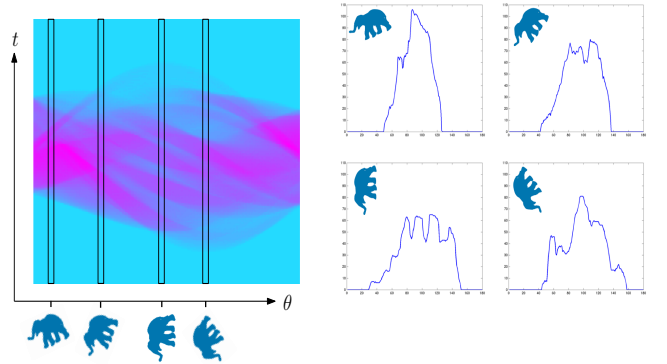


Fig. 4. Dip signals of a 2D elephant shape in various orientations. Each column in the sinogram (left) shows the change in the water level raising for a different angle for $0^\circ \leq \theta < 180^\circ$. The different profiles (right) contribute to the shape reconstruction. The angles from left to right, top to bottom, are $\theta = 20^\circ$, $\theta = 50^\circ$, $\theta = 90^\circ$ and $\theta = 130^\circ$.

Analogously to the computed tomography (CT) field, $\mathcal{G}(t; \theta)$ can be visualized using a sinogram – a visual representation of the object's projections. In our case, each column in the sinogram represents a dipping profile, which is acquired along a specific orientation, as illustrated in Figure 4. The different profiles contribute to the shape reconstruction. Examples of 2D and 3D sinograms of different shapes are shown in Figure 3.

In practice, the object is dipped in different orientations $\theta_i \in \Theta$, the water elevation rate is sampled in a constant rate, T , and a set of sampled measurements, $\mathcal{G}_{\theta_i}[n] = \mathcal{G}(nT; \theta_i)$, is extracted.

3.2 Reconstruction

In order to reconstruct $F(\mathbf{x})$, an inverse transform is applied. We first define a discrete version of the problem. Denoting the sampled measurements by $g_{\theta} \in \mathbb{R}^M$, we recast the reconstruction model as a least-squares optimization problem:

$$\arg \min_{\tilde{F}} \sum_{\theta_i \in \Theta} \|g_{\theta_i} - SR_{\theta_i} \tilde{F}\|_2, \quad (4)$$

where Θ is a set of dip orientations, R_{θ_i} is a rotation operator which rotates the object by an angle of θ_i around its center, and S is a summation operator that sums over the rows of the rotated matrix, counts the number of active voxels, and outputs an M -length vector. We note that g_{θ_i} is equivalent to a vector of volumes of slices, where the height of the slices is uniform. The unknown degrees of freedom that form \tilde{F} are given as a matrix of dimensions $M \times K \times L$. Our eventual goal is to find a voxelized binary solution in $\{0, 1\}^{M \times K \times L}$ that approximates $F(\mathbf{x})$; this solution represents the inversion.

The matrices S and R are defined via

$$S = \begin{pmatrix} \bar{S} & 0 & 0 \cdots & 0 \\ 0 & \bar{S} & 0 \cdots & 0 \\ \vdots & 0 & \ddots & 0 \\ 0 & 0 & \cdots & \bar{S} \end{pmatrix} \quad R = \begin{pmatrix} \bar{R}_{\theta_1} \\ \bar{R}_{\theta_2} \\ \vdots \\ \bar{R}_{\theta_K} \end{pmatrix},$$

where \bar{R}_{θ} is the rotation operator in matrix notation which rotates a vectorized object, generated by using bi-linear interpolation, and \bar{S} is a row summing matrix, which is operated on a vectorized matrix and sums the rows of the original matrix. In mathematical terms, $S = I \otimes v$ where I is an $M \times M$ unit matrix and v is a column vector of ones of size KL .

For notational convenience we express (4) in vectorized form, as follows:

$$\arg \min_f \|g - Bf\|_2, \quad (5)$$

where $f \in \mathbb{R}^N$ is the vectorized version of \tilde{F} (column stack), g is a chained version of all the measured g_{θ} , and $B = SR \in \mathbb{R}^{M \times N}$ is the chained “rotate and sum” matrix after the vectorization procedure. Finally, we use clamping to convert f into a binary solution, namely a vector in $\{0, 1\}^N$.

3.3 Numerical Solution

In this section we discuss several techniques for solving (5), which defines the inverse dip transform.

In practice, acquiring dip signals by measuring water’s elevation rate is a time consuming process. In addition, this mechanical operation may limit the number of dipping angles. Thus, a typical setup of a dip transform reconstruction constitutes an underdetermined problem, where the number of measurements does not uniquely define the desired object.

Conventional inverse Radon transform-based methods assume a full representation of the data within the measurements, and therefore do not fit our case. Instead, we aim at developing a computationally inexpensive method that computes the inverse transform while exploiting the structure of the dipping matrices.

Our underdetermined minimization problem (5) features infinitely many solutions, and one way to express them is by

$$f = f_p + z,$$

where f_p is a particular solution that satisfies $Bf_p = g$, and z is a vector in the null-space of B , namely a vector that satisfies $Bz = 0$. A common way of selecting a specific f , is to compute the solution that minimizes $\|f\|_2$ – this is known as the *min-norm solution*. This method seeks a solution in the active space of the operator, eliminating the null space. Often the null space contains high frequencies, and eliminating them may have a regularizing effect. Computing it is typically done via the pseudo-inverse of B . We solve

$$f = B^\dagger g,$$

where

$$B^\dagger = B(BB^T)^{-1},$$

and we assume that B has full row rank. In cases where the regularizing effect is not accomplished, Tikhonov or other types of regularization may be seamlessly incorporated into the solution procedure. In our case we could add a small diagonal shift to BB^T before inverting it to accomplish the desired effect. Cases of rank deficiency can be dealt with by computing the QR factorization of B and eliminating linearly dependent rows. It is easy to show that the solution in this case has the smallest norm.

In order to convert the solution back to a binary image we determine the value of the binary voxel using a threshold on the resulting image, using Otsu’s method [Otsu 1975]. The steps in the reconstruction algorithm are demonstrated in Figure 6.

Computing the pseudo inverse of B through its singular value decomposition (SVD) is the most numerically stable procedure, but it is unsuitable for the current problem, because the computation of the SVD does not respect sparsity. Instead, we compute the Cholesky factorization of BB^T once and for all: we seek a square lower triangular matrix G , such that $BB^T = GG^T$. We then store the factor G and use it repeatedly as we replace v throughout the computation. To form BB^T , we exploit the block structure of B and construct the matrix product block by block. The solution of g requires backward substitution and forward substitution for G^T and G , respectively. Forming the decomposition takes $O(W^3)$ floating point operations, whereas the backward and forward substitutions are $O(W^2)$. Thus, if we repeat the calculations for different vectors v , the overall computational cost is modest.

Taking a close look at $B = SR$ reveals that it is a sparse matrix with a unique structure. For R_{θ_i} , which rotates the object using a bi-linear interpolation, R typically has 4 nonzero entries per row in 2D reconstruction and 8 nonzero entries per row in 3D reconstruction. The matrix S also has a limited number of nonzero entries. The sparsity pattern of SR for 2D reconstruction is shown in Figure 5. Only 4% of the matrix entries are nonzero.

In a huge-scale setting a direct decompositional approach may consume prohibitive storage resources. In such cases we resort to iterative solvers such as LSQR [Paige and Saunders 1982] or LSMR [Fong and Saunders 2011], with an appropriate preconditioning approach that exploits the special structure of S and R .




Fig. 5. A sample from $B = SR$. The object size was $N = 50 \times 50$, number of orientations $K = 15$ and number of measurements for each angle $T = 50$. Only 4% of the matrix is non zero. Row wise there are approximately N non-zero entries (white), column-wise there are roughly K .

4 DIPPING AND VOLUME SAMPLING

Setting up the 3D dipping device is challenging as it involves physical measurements of water level, whose accuracy can be compromised by many factors such as ripples, tank's cross-section, recorder sensitivity, etc. In addition, rotating and lowering an object in the tank need to be performed in a mechanically-controlled manner, and be precisely calibrated.

We consider three different physical setups to measure the dip transform in practice (see Figure 7):

- (1) **Dipping** – Object is lowered at constant speed into a tank which is partially filled with water. The water elevation is recorded, thereby measuring the slices' volume.
- (2) **Draining** – Object is fixed underwater in a tank. A tap at the bottom of the tank is released, draining the water from the tank at a constant flow rate and the water lowering is recorded. This rate is proportional to the volume of the slices.
- (3) **Flooding** – Similarly to Dipping, the object is lowered at constant speed, however, here water fully filled tank. As the object is dipped, water overflows into a second tank, where the water elevation rate is recorded. The second tank is thinner, enhancing water elevation rate, which potentially leading to higher accuracy.

Each setup has its advantages and disadvantages; we provide a few details below.

The draining process involves an operation of a mechanical tap which introduces noise upon opening and closing it. Furthermore, the draining process may create whirlpools around the tap, resulting in noise in the water level measurement.

The flooding installation involves water transferring from one tank to another. Water level reading in the second tank is error prone as water may not be transferred properly to the second tank. For example, the water surface tension may prevent the water from dripping completely into the second tank even in an overflow scenario. Also, such water transfer takes time, thus it significantly reduces the measuring speed. We have selected the dipping setup to be used during our experiments. Although the object is lowered at a constant step size into the water, the amount of water elevation is changed within each step. When the body is lowered d units in the i th step and the water level rises δ_i units, the volume slice measurement corresponds to $d + \delta_i$ vertically along of the oriented

object. Since in our discrete setting in (4) the object's voxels are accumulated within equally spaced slices, we have to preprocess the measurements so they fit the algorithm input. Thus, we interpolate the acquired non-uniform slices into a continuous function and resample it uniformly. The new samples represent volumes of slices with the same width. The number of slices is determined by the user. That said, the finest level of details is bounded by the original water elevation rate; the lower the rate, the finer the level of detail of the reconstructed geometry. Figure 8 exhibits physical measures of a spherical ball, interpolated and resampled.

It should be noted that our dipping scheme assumes that the object has no vertical *caps* in any orientation. A cap is a vertical cavity that forms a vessel, in which water can be accumulated if the object is elevated vertically and air can be trapped, generating air pockets when the object is dipped in the opposite orientation. Most caps, if they exist, would be small and would have a minor effect akin to noise on the dip transform. Nevertheless, caps can be detected by dipping and then lifting back the object with the liquid trapped in the cap, yielding two different water levels. Flipping the object vertically allows detecting air pockets as they become caps.

5 EXPERIMENTS AND RESULTS

To test the feasibility of our technique we have constructed both 3D dipping simulator, which virtually dips known models, and a physical 3D dipping robot, shown in Figure 9. We have tested our technique both qualitatively and quantitatively.

5.1 Dipping Simulator

We have implemented a 3D dipping simulator to allow experiments using a wide range of 3D objects, as well as different dipping scenarios (varying number of dips, noise, etc.). To simulate our mechanism, a 3D model, represented as mesh surface, is first converted into a volume representation, using standard voxelization method [Patil and Ravi 2005]. Dipping in a specific orientation is simulated by dividing the object by equally spaced surfaces in the requested orientation. The number of voxels in each slab is counted, resulting in a single dip measurement. By collecting a set of dip measurements, we then apply our numerical solution to reconstruct the volume from the dip transform. It should be noted that the number of dipping orientations influences the quality of reconstruction. As the number of dipping profiles increases, more details are visible.

In order to demonstrate the evolution of a shape as a function of the number of dips, we performed the dipping process on several objects, using an increasing number of orientations.

Figure 10 demonstrates the quality of our reconstruction, using the raw voxels representation, as a function of the number of dips. As can be observed, as the number of dips increases more geometric details are reconstructed. Note specifically the fine reconstructed details of the octopus arms.

In Figure 11 we show our reconstruction evolution of different 3D models, with a resolution of 240^3 . While low number of orientations captures the raw shape of the object, using 10% of the required dips (which are theoretically needed for perfect reconstruction) results in high quality reconstructions. Table 1 summarizes several statistics for these models including a reconstruction error using the METRO method [Cignoni et al. 1998].

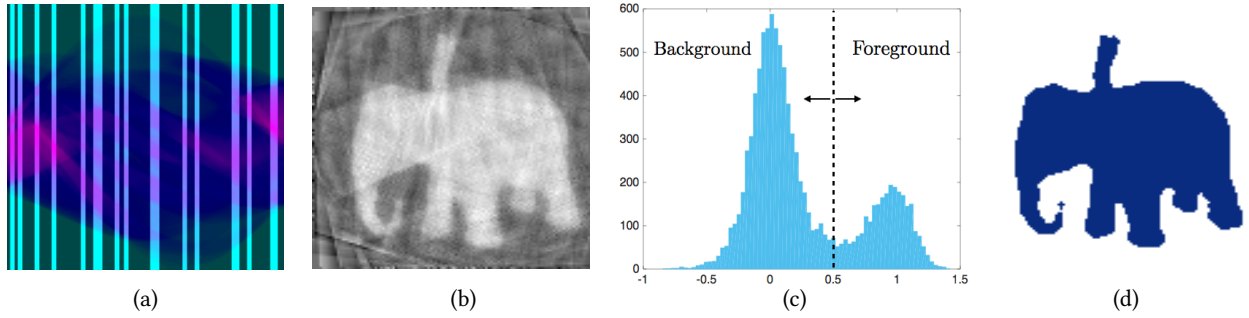


Fig. 6. Underdetermined Reconstruction. (a) A subset of 20 dipping angles (columns) is selected. (b) Reconstructed image using LSMR with 20 dips. (c) Histogram of the reconstructed image. It can be seen the there are two clusters, around 0 and around 1, that divides the image to background and foreground. (d) Resulting shape after thresholding, using Ostu's method.

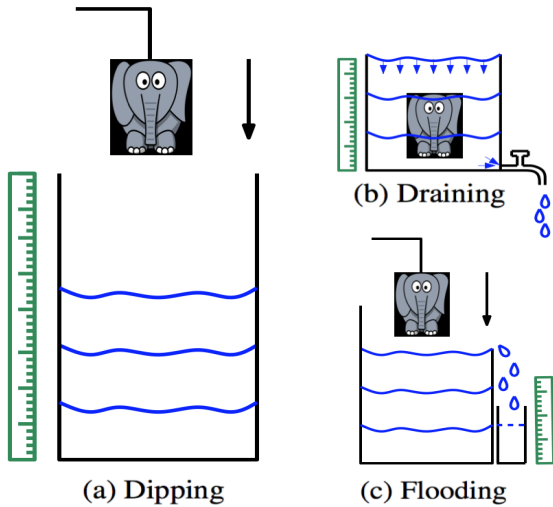


Fig. 7. Water scanning methods. (a) Dipping - the object is lowered into the tank and the water level rise rate is measured. (b) Draining - the water is drained using a tap at the bottom of the tank and the water lowering rate is measured. (c) Flooding - the object is lowered into a fully filled tank, causing an overflow and a transfer of water into a second tank where the rate of the water level elevation is measured.

Our method requires no line-of-sight between the sensor and the captured area, thus, able to reconstruct occluded parts of complex objects such as the hollow cube in Figure 12. It can be seen that the inner parts of the cube, which are completely occluded due to the shape of the body, are reconstructed as can be seen by slicing the objects into slabs.

In addition, we examined the robustness of our method, which can be understood from Figure 13. In 13(a) we have tested the reconstruction error (using METRO [Cignoni et al. 1998]) for increasing number of dips. We averaged over results from different objects and different resolutions, thus, we present the reconstruction error as a function of the percentage of dip orientations (from the number required for perfect reconstruction in each resolution). This

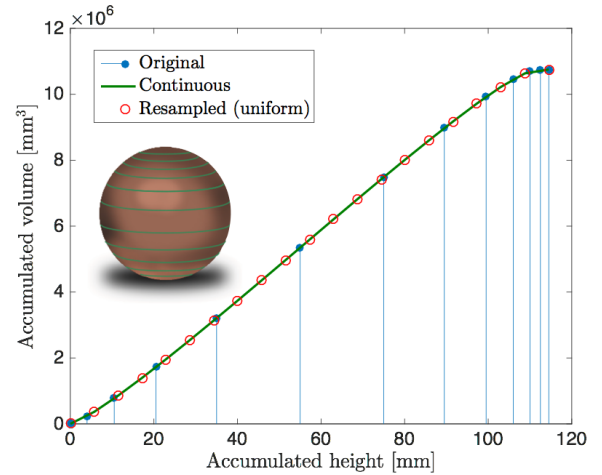


Fig. 8. Dipping method resampling. In order to comply with the inverse dip transform requirement, non-uniform slices of a spherical ball, with a diameter of 260 [mm], are converted into uniform slices. The ball measurements are interpolated and resampled. The number of uniform slices is dictated by the required voxel-resolution.

Model	Volume [%]	Error
Hand	19.41	0.02
Mother-Child	7.71	0.0125
Star	1.76	0.0126
Elephant	12.04	0.017
DNA ball	1.84	0.0433

Table 1. Statistics on the examples shown in Figure 11. The objects in the list are sorted top to bottom, with respect to Figure 11. The distances are calculated using METRO w.r.t to the bounding box diagonal of the object.

observation supports the high quality reconstruction achieved using




Fig. 9. Applying the dip transform in practice using the dipping robot – a mechanical arm able to dip a fixed object in different orientations and track the rate of water elevation.

only 10% of the required angles, as shown in Figure 11. In addition, since our practical system introduces noise to the measurements, we have added a white gaussian noise, to the virtual object's dips, and measured the quality of reconstruction as a function of the signal to noise ratio (SNR). Since our method solves an optimization problem with a matrix structure regularization, it can handle noise, in the sense that the error correlates with the noise in a linear nature. The results in 13(b), which demonstrate false positive and false negative, shows that the reconstruction error is considerable, even for low SNR.

5.2 The Dipping Robot

Following the discussion in Section 4 we constructed a fully automatic system, which incorporates the dipping method. The setup is rather simple, consisting of a robot arm with one linear axis lowering the object (moving perpendicular to the water surface) and two rotation arms responsible for setting the object's orientation. The vertical step was set to 2[mm] and the rotational to 1°. The water sensor used is non-contact absolute position transducer, plus a floating ball, with accuracy of 0.1[mm]. Figure 9, shows our 3D dipping robot in action.

Reconstructing an object from real dip measurements is a challenging task. System noise (due to jitter and water ripples) should be considered and a practical number of dipping orientations may not be sufficiently diverse. However, geometric priors on the reconstructed body, such as smoothness and symmetry, may help to improve the quality of reconstruction.

In order to reconstruct a 3D object we perform the following stages:

- (1) **Noise smoothing** - the measured signals are smoothed, using a normalized 5 samples gaussian kernel, in order to reduce noise which mainly lies in high frequency and does not degrade the object's details due to the smoothness assumption.
- (2) **Symmetry exploitation** - Orientations which yield the same dipping signal are measured only once and the resulting signal is duplicated to all of the appropriate locations.
- (3) **Dipping signals interpolation** - In order to generate dipping signals of missing orientations or to generate more samples between two consecutive measurements, interpolation between neighboring samples is performed.
- (4) **Reconstruction** - using the smoothed duplicated interpolated samples, our reconstruction algorithm is generated.

In the experiments, we dipped different objects as shown in the left column of Figure 15. These objects exhibit symmetry which enables significant reduction in the number of required orientations. For example, the cats object (Figure 15 first row) was dipped in 35 different orientations, followed by noise removal smoothing. Figure 14 compares between a real noisy dip measurement, its smoothed version and the simulated ground truth dip. In order to extend the number of dips, we next exploited the unique geometry of the object which has symmetry of 120° around the z axis and an additional reflection symmetry within each third. Thus, the number of real dips was duplicated into 210 virtual dips. These dips were interpolated, resulting a total number of 420 dips.

This number of dips enables resolution of 40^3 , which can be upsampled and smoothed, due to the smoothness property of the object. The reconstruction of our dipping algorithm on the real measurements is shown in Figure 15(d). In addition, our dipping reconstruction is compared qualitatively to a scan of a 3D structured light scanner. As can be seen in Figure 15(c), the structured light scanner can not reconstruct self occluded parts of the objects. However, our dipping reconstruction results in a complete object.

6 CONCLUSIONS AND FUTURE WORK

The paper presents a novel three-dimensional shape acquisition and reconstruction method based on the well-known Archimedes equality between fluid displacement and the submerged volume. A strong feature of our method is that it relies on fluid displacement as the fundamental shape sensor. The liquid has no line-of-sight requirements, which is advantageous compared to optical sensors. Furthermore, our technique allows for seamless penetration into cavities and hidden parts of the object as well as transparent and glossy materials, thus bypassing all visibility and optical limitations of conventional scanning devices. Our results demonstrate the




Fig. 10. Reconstruction for different number of dip orientations. The octopus resolution is 120^3 . For low number of angles the fine details are lost (legs, eyes, etc.), and as the number of angles increases more of the shape is reconstructed.

capability of acquisition and reconstruction for a wide range of shapes with large occlusions and hidden regions.

One obvious limitation of the proposed method is its speed of acquisition. The mechanical arm dips the object vertically step by step, and it is necessary to read the liquid level in between steps, which introduces a temporal bottleneck. We are currently exploring ways for improving this by developing a continuous dipping and reading process, whereby we measure the ratio between the vertical dipping speed and the liquid elevation. We are also investigating faster computational methods. Future work may include the development of sparse recovery techniques based on compressed sensing, which exploit the redundancy that exists in the binary representation of the voxels, to significantly reduce the number of dipping orientations. This approach may lead to a much smaller optimization problem, which in turn will allow for further acceleration of the solution procedure, and the reduction of the overall computational time required to solve the inverse transform problem.

Another limitation that we have discussed is “no-cap” assumption. The dipping method assumes that the object has no significant cap or vessel. This may introduce errors once the liquid is spilled into the cap and fills it up. We are considering dealing with such caps by measuring the level of liquid also during a vertical lift of the object. The amount of liquid trapped in the cap can be measured to assess the volume of the cap. Analyzing the profiles of the liquid level measured on up and down movements may reveal some additional geometric features of the shape of the cap.

The success of this volumetric technique leads us to speculate that we may be able to consider a multi-modal acquisition setup where the reconstruction combines the information gained by the various modalities. For example, we may be able to utilize the strengths of laser scanners for sampling the shape-visible exterior and the volumetric and occluded information reconstructed from the dip transform.

7 ACKNOWLEDGEMENTS

We thank the anonymous reviewers for their helpful comments. This project was supported in part by the Joint NSFC-ISF Research Program 61561146397, jointly funded by the National Natural Science Foundation of China and the Israel Science Foundation (No. 61561146397), the National Basic Research grant (973) (No. 2015CB352501) and the NSERC of Canada grant 261539.

REFERENCES

- Matthew Berger, Andrea Tagliasacchi, Lee M. Seversky, Pierre Alliez, Joshua A. Levine, Andrei Sharf, and Claudio T. Silva. 2014. State of the Art in Surface Reconstruction from Point Clouds. In *Eurographics 2014 - State of the Art Reports*, Sylvain Lefebvre and Michela Spagnuolo (Eds.). The Eurographics Association. DOI: <https://doi.org/10.2312/egst.20141040>
- R. N. Bracewell. 1956. Strip Integration in Radio Astronomy. *Australian Journal of Physics* 9 (June 1956), 198. DOI: <https://doi.org/10.1071/PH560198>
- Paolo Cignoni, Claudio Rocchini, and Roberto Scopigno. 1998. Metro: measuring error on simplified surfaces. In *Computer Graphics Forum*, Vol. 17. Wiley Online Library, 167–174.
- David Chin-Lung Fong and Michael Saunders. 2011. LSMR: An iterative algorithm for sparse least-squares problems. *SIAM Journal on Scientific Computing* 33, 5 (2011), 2950–2971.
- Richard Gordon, Robert Bender, and Gabor T. Herman. 1970. Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and x-ray photography. *Journal of Theoretical Biology* 29 (1970), 471–481.
- James Gregson, Michael Krimerman, Matthias B. Hullin, and Wolfgang Heidrich. 2012. Stochastic Tomography and Its Applications in 3D Imaging of Mixing Fluids. *ACM Trans. Graph.* 31, 4, Article 52 (July 2012), 10 pages. DOI: <https://doi.org/10.1145/2185520.2185548>
- Gabor T. Herman. 2009. *Fundamentals of Computerized Tomography: Image Reconstruction from Projections* (2nd ed.). Springer Publishing Company, Incorporated.
- Ivo Ihrke, Kiriakos N. Kutulakos, Hendrik P. A. Lensch, Marcus Magnor, and Wolfgang Heidrich. 2010. Transparent and Specular Object Reconstruction. *Computer Graphics Forum* (2010). DOI: <https://doi.org/10.1111/j.1467-8659.2010.01753.x>
- Takashi Ijiri, Shin Yoshizawa, Hideo Yokota, and Takeo Igarashi. 2014. Flower Modeling via X-ray Computed Tomography. *ACM Trans. Graph.* 33, 4, Article 48 (July 2014), 10 pages. DOI: <https://doi.org/10.1145/2601097.2601124>
- Avinash C. Kak and Malcolm Slaney. 2001. *Principles of computerized tomographic imaging*. Society for Industrial and Applied Mathematics, Philadelphia (Pa.). <http://opac.inria.fr/record=b1100961>
- Nobuyuki Otsu. 1975. A threshold selection method from gray-level histograms. *Automatica* 11, 285–296 (1975), 23–27.
- Christopher C Paige and Michael A Saunders. 1982. LSQR: An algorithm for sparse linear equations and sparse least squares. *ACM transactions on mathematical software* 8, 1 (1982), 43–71.
- Sandeep Patil and B Ravi. 2005. Voxel-based representation, display and thickness analysis of intricate shapes. In *Ninth International Conference on Computer Aided Design and Computer Graphics (CAD-CG'05)*. IEEE, 6–pp.
- J. Radon. 1917. Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten. *Akad. Wiss.* 69 (1917), 262–277.
- Borislav Trifonov, Derek Bradley, and Wolfgang Heidrich. 2006. Tomographic Reconstruction of Transparent Objects. In *Proceedings of the 17th Eurographics Conference on Rendering Techniques (EGSR '06)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 51–60. DOI: <https://doi.org/10.2312/EGWR/EGSR06/051-060>
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011. Building Volumetric Appearance Models of Fabric Using Micro CT Imaging. In *ACM SIGGRAPH 2011 Papers (SIGGRAPH '11)*. ACM, New York, NY, USA, Article 44, 10 pages. DOI: <https://doi.org/10.1145/1964921.1964939>




Fig. 11. 3D-dip reconstructions using different number of orientations. An increasing number of orientations (from left to right) allows better reconstruction of the dipped object (shown on the right). The objects resolution is 240^3 . The models are reconstructed using 1%-3% (first column), 5% (second column) and 10% (third column) from the required number of orientations for perfect reconstruction.




Fig. 12. 3D dip reconstruction of a hollow cube (left). The cube inner shape is reconstructed as seen by slicing the reconstruction to 3 slabs (right).

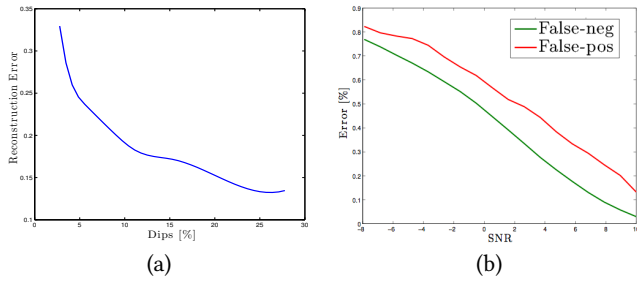


Fig. 13. Reconstruction errors. (a) METRO reconstruction error for different percentage of dip orientations (from the number required for perfect reconstruction). The reconstructed error decreases dramatically with the number of orientations, allowing the use of only a small portion of the required dipping orientations. (b) Normalized reconstruction error as a function of the SNR. Noise is introduced to the samples. False positive error marks the creation of false object's parts in the reconstruction, while false-negative point to missing object's parts. The errors were calculated using different objects with different resolutions.

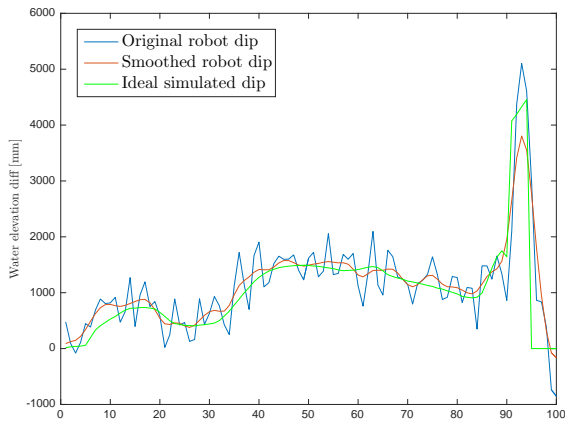


Fig. 14. A measurement of the dipping robot, before (red) and after (blue) smoothing. The profile is compared to an ideal dip in the same orientation (green).



Fig. 15. 3D dip reconstructions comparison. (a) Picture of the objects during the dipping (b) Profile picture of the printed objects (c) Structured light scanner reconstruction (d) Our 3D reconstruction using the dipping robot. Occluded parts of the body have no line-of-sight to the scanner sensor, while the dipping robot, using water, is able to reconstruct these hidden parts.