A Bucket Grid Structure to Speed Up Table Lookup in Gauge-Based Photometric Stereo

Helena Cristina G. Leitão
Inst. of Computing, Fed. Fluminense Univ. – R. Passo da Pátria, 156, Niterói, RJ, Brazil
hcgl@ic.uff.br

Rafael Felipe V. Saracchini
Jorge Stolfi
Inst. of Computing, State Univ. of Campinas – Cx. P. 6176, 13083-970, Campinas, SP, Brazil
ra069320,stolfi@ic.unicamp.br

Abstract

In this paper, we show how to speed up the table lookup step in gauge-based multi-image photometric stereo. In that step, one must find a pixel of a gauge object*, of known shape and color, whose appearance under m different illumination fields is similar to that of a given scene pixel. This search reduces to finding the closest match to a given m-vector in a table with a thousand or more m-vectors. Our speed-up method exploits the fact that the table is in fact a fairly flat two-dimensional manifold in m-dimensional space, so that the search can be efficiently solved with a two-dimensional bucket grid structure.*

1. Introduction

1.1. Variable-lighting photometric stereo

In *variable-lighting photometric stereo* (VLPS), the goal is to determine the 3D geometry of a scene from a list of $m \geq 3$ monochromatic digital photos $S_1, .. S_m$, all taken with different lighting conditions but with the same pose and viewpoint. As shown by R. J. Woodham in 1980 [8], by analyzing the m pixel intensities $S_i[p]$ at any image point p, one can recover the unit vector $\vec{s}[p]$ that is perpendicular to surface element which is visible at p. The third dimension (depth) can then be recovered by integration of this data. This problem has attracted a lot of attention in recent years [3, 5, 7, 11, 2, 10].

To perform the above analysis, one must have enough information about the *bidirectional radiance distribution function* (BRDF) of the surface, and about the light field Φ_i in each image S_i. The BRDF of the scene's surface at point p of the image is a function $\sigma[p](\vec{n}, \vec{u}, \vec{v})$ that gives the apparent brightness of the surface when oriented with normal $\vec{n}$, viewed from the direction $\vec{v}$, and illuminated with unidirectional light of unit intensity flowing in the direction $\vec{u}$. (Note that we include the geometric light spread factor $\max\{0, -\vec{u} \cdot \vec{n}\}$ in the BRDF itself.)

1.2. Gauge-based VLPS

In the *gauge-based* variant of VLPS, the BRDF information is indirectly given by m images $G_1, .. G_m$ of a *light gauge*, a sample object of known shape and color; where each gauge image G_i is taken under the same lighting conditions as the corresponding scene image S_i. (Depending on the application, it may be convenient to include the gauge object as part of the scene itself. In that case, each G_i will be just a sub-image of S_i.)

In this paper, we assume that all images $S_1, .. S_m$ have been geometrically corrected, trimmed, and aligned, so that each point p on their common domain $\mathcal{S}$ corresponds to the same point on the scene's visible surface. The same condition is assumed for the gauge images $G_1, .. G_m$, whose common domain will be denoted by $\mathcal{G}$. We also assume a linear pixel value scale (with samples directly proportional to physical light intensity). Finally, we assume that the surface normal vector $\vec{g}[q]$ is known for every point q on the gauge's images.

In its basic form, gauge-based VLPS is viable only if all visible scene and gauge surfaces have the same finish everywhere, except for variations in intrinsic color. That is, the BRDF $\sigma[p]$ of the scene at each image point p, and the BRDF $\gamma[q]$ of the gauge at any point q, must be multiples

of some fixed BRDF $\bar{\beta}$:

$$\begin{aligned} \sigma[p](\vec{n},\vec{u},\vec{v}) &= \overset{*}{s}[p]\,\bar{\beta}(\vec{n},\vec{u},\vec{v}) \\ \gamma[p](\vec{n},\vec{u},\vec{v}) &= \overset{*}{g}[p]\,\bar{\beta}(\vec{n},\vec{u},\vec{v}) \end{aligned} \tag{1}$$

The constant factors $\overset{*}{s}[p]$ and $\overset{*}{g}[q]$ in these formulas are the *intrinsic lightness* or *albedo* of the scene and gauge surface, respectively, at those points. Observe that the gauge's albedo $\overset{*}{g}[q]$ and normal direction $\vec{g}[q]$ must be known for all $q \in \mathcal{G}$; typically, one uses a spherical gauge with uniform albedo, preferably white ($\overset{*}{g}[q] = 1$ everywhere).

Another necessary condition for gauge-based VLPS is that the BRDF $\bar{\beta}$ must be dominated by wide-angle scattering, with no mirror-like reflection or sharp glossy scattering. The standard example is the Lambertian BRDF,

$$\bar{\beta}(\vec{n},\vec{u},\vec{v}) = \max\{0, -\vec{u}\cdot\vec{n}\} \tag{2}$$

However, almost any BRDF $\bar{\beta}$ will do, as long as it doesn't have impulse-like components (sharp peaks or ridges). For simplicity, we will also assume that the the images are taken under nearly parallel projection and illuminated by distant light sources; so that the viewing direction $\vec{v}$ and the lighting conditions are the same at every point of $\mathcal{S}$ or $\mathcal{G}$.

Gauge-based VLPS can be extended to multichannel (e.g. RGB trichromatic) images, yielding a single normal map $\vec{s}[p]$ but a different albedo map $\overset{*}{s}_\lambda[p]$ for each spectral band λ. The latter provide the illumination-independent *intrinsic color* of the scene at each pixel.

1.3. Fundamental equations

The key idea of gauge-based VLPS is that the intensity of each point on a scene photo S_i or a gauge photo G_i can be analyzed into the product of two factors: the *intrinsic albedo*, $\overset{*}{s}$ or $\overset{*}{g}$ (that depends only on the surface's material and finish) and the *lighting factor*, (that depends only on the index i and on the surface's slope, $\vec{s}$ or $\vec{g}$). Specifically,

$$\begin{aligned} S_i[p] &= \overset{*}{s}[p]\,L_i(\vec{s}[p]) \\ G_i[q] &= \overset{*}{g}[q]\,L_i(\vec{g}[q]) \end{aligned} \tag{3}$$

Here, each L_i is the *shading function* implied by the lighting field Φ_i and the BRDF $\bar{\beta}$. It maps each unit vector $\vec{n}$ to the apparent lightness of a white surface perpendicular to $\vec{n}$, and is given by the formula

$$L_i(\vec{n}) = \int_{\mathbb{S}^2} \Phi_i(\vec{u})\bar{\beta}(\vec{n},\vec{u},\vec{v})\,\mathrm{d}\vec{u} \tag{4}$$

The factor $\Phi_i(\vec{u})$ is the intensity of the light flow that is incident on the surface from direction $\vec{u}$. The uniform lighting hypothesis allows us to assume that $\Phi_i(\vec{u})$ does not depend explicitly on the position on the surface, but only on the light source's direction $\vec{u}$.

Note that the value of $\Phi_i(\vec{u})$ is irrelevant for directions $\vec{u}$ that point outwards from the local surface, or for points that are not on the surface; so it is indeed plausible to have a single function Φ_i for the whole scene, independently of the local surface orientation. In particular, this lighting model allows attached shadows, and is adequate for scenes consisting of a single mostly convex object. On the other hand, this model cannot account for projected shadows, radiosity effects, or sources with uneven light distribution.

Note that, in this model, the intrinsic albedos $\overset{*}{g}$ and $\overset{*}{s}$ and the normals $\vec{s}$ and $\vec{g}$ are distinct from each other and vary from point to point, but are same for all i; whereas the shading functions L_i are the same for S_i and G_i and are constant over each image, but are different for each i.

If formulas (3) hold, then we can determine the normal $\vec{s}[p]$ at a point p of the scene images by finding a point in the gauge images that reacts in the same way as p to changes in lighting directions, except for the albedos $\overset{*}{s}[p]$ and $\overset{*}{g}[q]$. More precisely, we must find $q \in \mathcal{G}$ such that the m-vectors

$$\begin{aligned} \boldsymbol{S}[p] &= (S_1[p], S_2[p], \ldots, S_m[p]) \\ \boldsymbol{G}[q] &= (G_1[p], G_2[p], \ldots, G_m[p]) \end{aligned} \tag{5}$$

are multiples of each other. The vectors $\boldsymbol{S}[p]$ and $\boldsymbol{G}[q]$ are called the *observation vectors* (OVs) of the points p and q. Having located the matching gauge point q, we can recover the normal vector $\vec{s}[p]$ and albedo $\overset{*}{s}[p]$ of the scene at p by the formulas

$$\vec{s}[p] = \vec{g}[q] \qquad \overset{*}{s}[p] = \frac{|\boldsymbol{S}(p)|}{|\boldsymbol{G}(q)|}\overset{*}{g}(q) \tag{6}$$

This method will fail if there are two points q', q'' on the gauge images which have different normals (that is, $\vec{g}[q'] \neq \vec{g}[q']$) but collinear OVs (that is, $\boldsymbol{G}[q'] = \alpha\boldsymbol{G}[q'']$ for some scalar α). To avoid this problem, the number of images m must be at least 3, and the light fields $\Phi_1, .. \Phi_m$ must be sufficiently varied to break any such ambiguities. We will assume that this condition is satisfied in what follows.

1.4. The table-lookup step

The most time-consuming part of gauge-based VLPS, as outlined above, is locating the point $q \in \mathcal{G}$ such that $\boldsymbol{S}[p]$ is a multiple of $\boldsymbol{G}[q]$. If neither vector is zero, this is equivalent to matching the *observation signatures* $\boldsymbol{s}[p]$ and $\boldsymbol{g}[q]$, defined by

$$\boldsymbol{s}[p] = \frac{\boldsymbol{S}[p]}{|\boldsymbol{S}[p]|} \qquad \boldsymbol{g}[q] = \frac{\boldsymbol{G}[q]}{|\boldsymbol{G}[q]|} \tag{7}$$

Here $|\cdot|$ is any norm of $\mathbb{R}^m$, e.g. the Euclidean norm

$$|\boldsymbol{X}| = \sqrt{\textstyle\sum_{i=1}^{m} X_i^2} \tag{8}$$

Note that the position q is not meaningful by itself; it is only used to associate the observation signature $\boldsymbol{g}[q]$ to the normal $\vec{g}[q]$ and to the OV modulus $\mathring{g}[q] = |\boldsymbol{G}[q]|$. Therefore, we can replace the gauge images by a *signature table*, an unordered set of triplets

$$\mathfrak{T} = \{ (\boldsymbol{g}[q], \vec{g}[q], \mathring{g}[q]) : q \in \mathcal{G} \} \tag{9}$$

The computation of $\vec{s}[p]$ then becomes a *closest-match table look-up problem*, where we look for the element $(\boldsymbol{g}, \vec{g}, \mathring{g})$ of the table $\mathfrak{T}$ that minimizes the distance $\mathrm{dist}(\boldsymbol{g}, \boldsymbol{s}[p])$.

The brute-force solution to this problem would be to scan the table $\mathfrak{T}$, computing $\mathrm{dist}(\boldsymbol{g}, \boldsymbol{s}[p])$ for each signature $\boldsymbol{g}$ in it, while keeping track of the closest-matching entry. However, in order to provide a good coverage of all possible normal directions, the table $\mathfrak{T}$ must have tens of thousands of entries. Since the lookup must be repeated for each pixel of the scene domain $\mathcal{S}$, it may take tens of minutes to process a single set of scene images with this method.

1.5. Previous work

Several techniques have been proposed in the literature to speed up the table search step. Woodham himself [9, 10] used a regular m-dimensional grid spanning the hypercube $[0, 1]^m$, with 2^b cells along each axis, for some bit count b. In the preprocessing phase, each gauge observation vector $\boldsymbol{G}[q]$ was quantized with b bits per coordinate, yielding the m-tuple of indices of some grid cell where the associated normal vector $\vec{g}[q]$ was stored. (Woodham assumed uniform albedos $\mathring{s} = \mathring{g}$, so there was no reason to normalize the observation vectors.) In the lookup phase, each scene OV $\boldsymbol{S}[p]$ was mapped to a table cell in the same way, and the desired normal $\vec{s}[p]$ was recovered from the grid. One obvious disadvantage of this method is the size of the grid: 2^{mb} entries, which is about 250,000 for $m = 3$ and $b = 6$.

Later works have proposed other general m-dimensional nearest-point algorithms for this task. For instance, Hertzmann and Seitz [4] use *approximate nearest neighbour*(ANN) of Arya et al. [1]; while Zhong and Little [11] use the *locally sensitive hashing* of Indyk and Motwani [6]. However, all these methods have a common shortcoming: they consider the set of all gauge signatures $\mathbb{G} = \{ \boldsymbol{g}[q] : q \in \mathcal{G} \}$ to be a generic cloud of points scattered in m-dimensional space, and therefore use general m-dimensional nearest-neighbor search algorithms — which are inherently expensive in space and/or time [6].

2. Fast table searching with 2D bucketing

We now describe an algorithm to locate the best-matching entry in the signature table, that exploits the peculiar shape of the set $\mathbb{G}$ to achieve very fast look-up at a modest space cost.

2.1. Shape of the signature table

The key observation for our improved method is that *the set $\mathbb{G}$ of all gauge signatures is essentially a two-dimensional subset of $\mathbb{R}^m$*. Therefore, we can reduce the problem to a two-dimensional nearest-point search, which can be solved very efficiently by a two-dimensional bucket grid scheme.

To understand the key observation above, note that, because of formulas (3) and (7), the observation signatures $\boldsymbol{s}[p]$ and $\boldsymbol{g}[q]$ can be expressed as $\boldsymbol{l}(\vec{s}[p])$ and $\boldsymbol{l}(\vec{g}[q])$, respectively; where $\boldsymbol{l}$ is the *lighting signature function*

$$\boldsymbol{l}(\vec{n}) = \frac{\boldsymbol{L}(\vec{n})}{|\boldsymbol{L}(\vec{n})|} \tag{10}$$

and $\boldsymbol{L}(\vec{n}) = (L_1(\vec{n}), .. L_m(\vec{n}))$. Note that the function $\boldsymbol{l}$, that maps surface normals to lighting signatures, is defined only on the hemisphere H of $\mathbb{S}^2$ consisting of the normal directions that deviate less than 90 degrees from the viewing direction $\vec{v}$. On the other hand, a good gauge object must provide a fairly dense and uniform sampling of H (which is why spheres are normally used for that purpose). It follows that the set of gauge signatures must be a fairly dense and uniform cover of $K = \boldsymbol{l}(H)$, the range of the function $\boldsymbol{l}$.

Now, given our assumption that the gauge's BRDF $\bar{\beta}$ lacks the sharp spikes of mirror-like reflection, the shading factors $L_i(\vec{n})$ given by formula (4) are continuous functions of the surface normal $\vec{n}$. In fact, L_i is typically fairly smooth, with just a few broad and hardly-distinguishable maxima. Furthermore, the gauge-based VLPS problem is solvable if and only if the function $\boldsymbol{l}(\vec{n})$ is invertible, i.e. for every point $\boldsymbol{v}$ of $\mathbb{S}^{m-1}$ there is at most one direction $\vec{n}$ such that $\boldsymbol{l}(\vec{n}) = \boldsymbol{v}$. If this condition holds, the range K of $\boldsymbol{l}$ is an embedding of the hemisphere H into $\mathbb{S}^{m-1}$. Finally, since the observation signatures are contained in the positive orthant of $\mathbb{R}^m$, the width of K, as seen from the origin of $\mathbb{R}^m$, is at most 90 degrees.

From these considerations, intuition suggests (and experience confirms) that the range K of $\boldsymbol{l}$ is a relatively flat patch of a 2-dimensional manifold (surface) immersed in $\mathbb{S}^{m-1}$; and that the gauge's observation signatures $\mathbb{G}$ must be distributed over K with fairly uniform density.

2.2. The 2D bucketing scheme

In our method, the signature table $\mathfrak{T}$ is preprocessed as follows. We first compute the centroid $\boldsymbol{b}$ of $\mathbb{G}$ (seen as a set of points of $\mathbb{R}^m$), and two orthogonal unit vectors $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^m$ that define its directions of maximum extent. These vectors are found by computing the $m \times m$ coordinate moment matrix M of the displacements $\boldsymbol{g} - \boldsymbol{b}$, for all $\boldsymbol{g} \in \mathbb{G}$, and taking the eigenvectors associated to its two largest eigenvalues. The point $\boldsymbol{b}$ and the vectors $\boldsymbol{u}, \boldsymbol{v}$ define

a two-dimensional affine subspace P of $\mathbb{R}^m$, the *signature projection plane*, which is roughly coplanar with the set $\mathbb{G}$. The orthogonal projection onto P of a given observation signature $\boldsymbol{g}$ will be denoted by $\downarrow \boldsymbol{g}$.

Next, we choose a regular grid of $N \times N$ square cells on the projection plane P. This grid is centered on the point $\boldsymbol{b}$, has its sides parallel to the vectors $\boldsymbol{u}$ and $\boldsymbol{v}$, and is barely large enough to contain the projection $\downarrow \boldsymbol{g}$ of any observation signature $\boldsymbol{g}$ in $\mathfrak{T}$. More precisely, the grid side is $2R$, where

$$R = \varepsilon + \max \{ |(\boldsymbol{g} - \boldsymbol{b}) \cdot \boldsymbol{u}|, |(\boldsymbol{g} - \boldsymbol{b}) \cdot \boldsymbol{v}| : \boldsymbol{g} \in \mathbb{G} \} \quad (11)$$

for some small safety margin ε. Having chosen the grid, we build, for each cell $C[i,j]$, a linked *bucket list* $B[i,j]$ of all table entries $(\boldsymbol{g}, \vec{g}, \overset{\bullet}{g})$ whose signatures $\boldsymbol{g}$ project onto that cell. We also compute the corresponding *bucket mean* $\mu[i,j]$, defined as the barycenter of all signatures $\boldsymbol{g}$ in the list $B[i,j]$; and the *bucket radius* $\rho[i,j]$, defined as the the maximum Euclidean distance from $\mu[i,j]$ to any signature $\boldsymbol{g}$ in that list. See figure 1.

The two-dimensional shape of $\mathbb{G}$ means that the entries in $B[i,j]$ are fairly close to each other, even if their mean distance from the plane P is large compared to the cell size. This property remains true even when m is greater than 3.

Once the bucket grid has been constructed, the scene signatures $\boldsymbol{s}[p]$ are looked up with procedure 1 below. Its steps are explained in sections 2.3 through 2.5.

Procedure 1 (Table lookup) *Given an observation signature $\boldsymbol{s}$, finds the entry tmin $\in \mathfrak{T}$ whose signature field $\boldsymbol{g}$ is most similar $\boldsymbol{s}$.*

1. $i \leftarrow \lfloor N((\boldsymbol{g} - \boldsymbol{b}) \cdot \boldsymbol{u} + R)/(2R) \rfloor$;
2. $j \leftarrow \lfloor N((\boldsymbol{g} - \boldsymbol{b}) \cdot \boldsymbol{v} + R)/(2R) \rfloor$;
3. *dmin* $\leftarrow +\infty$;
4. For each pair (r,s) in Δ, in order, do
   5. If *dmin* $\leq \delta \| (r,s) \|$, return *tmin*.
   6. $(i',j') \leftarrow (i,j) + (r,s)$;
   7. If $0 \leq i' < N$ and $0 \leq j' < N$, then
      8. If *dmin* $> \operatorname{dist}(\boldsymbol{s}, \mu[i',j']) - \rho[i',j']$, then
         9. For each $t = (\boldsymbol{g}, \vec{g}, \overset{\bullet}{g})$ in $B[i',j']$, do
            10. Set $d \leftarrow \operatorname{dist}(\boldsymbol{s}, \boldsymbol{g})$;
            11. If $d <$ *dmin*, set *dmin* $\leftarrow d$ and *tmin* $\leftarrow t$.
12. Return *tmin*.

2.3. Bucket grid searching

In order to locate the entry closest to a given observation signature $\boldsymbol{s}$, we compute the indices $(i,j) = h(\boldsymbol{s})$ of the cell that contains its projection $\downarrow \boldsymbol{s}$ (steps 1–2). We then search for the entry *tmin* of $\mathfrak{T}$ whose signature $\boldsymbol{g}$ is closest to $\boldsymbol{s}$ in the list $B[i,j]$, and then, if necessary, in nearby buckets $B[i',j']$, in some appropriate order (steps 3–4). Note that some buckets may be empty, and that the best match to the query $\boldsymbol{s}$ may not be in bucket $B[i,j]$ — even if that bucket is non-empty.

The bucket parameters $\mu[i,j]$ and $\rho[i,j]$ allow us to quickly skip over buckets that cannot possibly contain a better match to the query signature $\boldsymbol{s}$. More precisely, we should examine a bucket $B[i',j']$ only if the query signature $\boldsymbol{s}$ is closer to the bucket's bounding ball than to the best match $\boldsymbol{x}$ found so far (step 8); namely, only if

$$\operatorname{dist}(\boldsymbol{s}, \boldsymbol{x}) > \operatorname{dist}(\boldsymbol{s}, \mu[i',j']) - \rho[i',j'] \quad (12)$$

We will call condition (12) the *bucket scan criterion*.

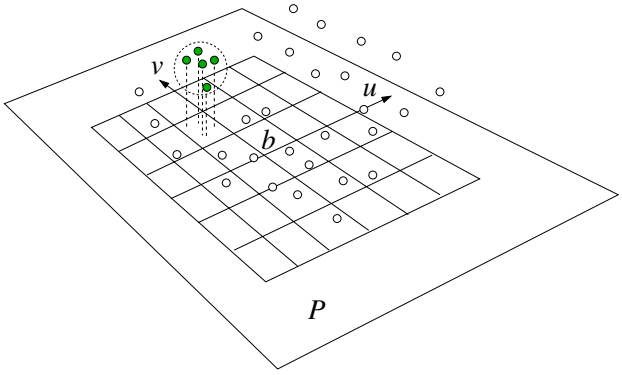

Figure 1. The two-dimensional bucketing algorithm, for $m = 3$, showing some observation signatures in $\mathbb{G}$ (small circles), a bucket list $B[i,j]$ (small gray circles), and the enclosing sphere (dotted circle) defined by the bucket's centroid $\mu[i,j]$ and radius $\rho[i,j]$.

2.4. Bucket search order and early return

The bucket scan criterion (12) will often allow us to skip a bucket $B[i',j']$ without examining its entries. However, if we were to apply this criterion for all buckets, individually, the running time would still be proportional to the number N^2 of buckets in the grid — which, as in any hashing scheme, is expected to be proportional to the size of the table. In that case, the bucket grid search would improve on the brute-force solution only by a constant factor, at best.

To avoid scanning the whole grid, we enumerate the buckets $B[i',j']$ in a specific order, starting with the bucket $B[i,j]$ that contains the query and then moving gradually away from it (step 4). A second criterion (step 5) then allows us to abandon the search as soon as we detect that none

of the buckets still to be scanned can possibly contain a better match than the one found so far. Typically, this happens after scanning only a small fraction of the bucket array.

More precisely, consider two signatures $\boldsymbol{s}'$ and $\boldsymbol{s}''$ that project orthogonally to P into cells $C[i', j']$ and $C[i'', j'']$, respectively. It is easy to see that

$$\operatorname{dist}(\boldsymbol{s}', \boldsymbol{s}'') \geq \operatorname{dist}(C[i', j'], C[i'', j'']) \tag{13}$$

In this formula, $\operatorname{dist}(C[i', j'], C[i'', j''])$ is the minimum distance between the two cells, seen as subsets of P. This distance is

$$\operatorname{dist}(C[i', j'], C[i'', j'']) = \delta \|(i' - i'', j' - j'')\| \tag{14}$$

where $\delta = 2R/N$ is the grid mesh size, and

$$\|(r, s)\| = \sqrt{(\max\{0, |r| - 1\})^2 + (\max\{0, |s| - 1\})^2} \tag{15}$$

Note that $\|(r, s)\|$ is a bit smaller than the Euclidean norm $|(r, s)| = \sqrt{r^2 + s^2}$.

We conclude that a bucket $[i', j']$ can be ignored if the cell distance bound (13) excludes the possibility that a better match can be found within it; that is, if

$$\operatorname{dist}(\boldsymbol{s}, \boldsymbol{x}) \leq \delta \|(i' - i, j' - j)\| \tag{16}$$

Note that condition (16) is weaker than condition (12). However, condition (16) depends only on the current match $\boldsymbol{x}$ and the difference $(i' - i, j' - j)$ between the cell indices. Therefore, if we scan the buckets (i', j') in such an order that $\|(i' - i, j' - j)\|$ is increasing, we can stop the search as soon as that condition is satisfied (step 5).

For that purpose, as part of the table preprocessing we precompute a list Δ of all vectors (r, s) in $\{-N + 1.. N - 1\} \times \{-N + 1.. N - 1\}$, sorted by increasing value of $\|(r, s)\|$ (and breaking ties by $|(r, s)|$). For each query signature $\boldsymbol{s}$, we enumerate the buckets $B[i', j']$ by taking each displacements (r, s) from the ordered list Δ and computing $(i', j') \leftarrow (i, j) + (r, s)$ (step 6), provided that i' and j' lie in $\{0, .. N - 1\}$ (step 7). See figure 2.

2.5. Analysis

The average computation cost of algorithm 1 is roughly $Bb + Dd + O(1)$, where b is the average number of buckets examined in each call (step 5), d is the average number of table entries tested (step 10), and B, D are the costs associated to those two operations.

In the extreme case when $N = 1$, we will have $b = 1$ and $d = |\mathfrak{T}|$, which is equivalent to a brute-force search of $\mathfrak{T}$. As N increases, d will usually decrease towards 1, because the test of step 5 will get satisfied before the procedure finds the second non-empty bucket. At the same time, b will increase immediately to about 10, because $\|(r, s)\|$ is zero for the first nine pairs (r, s) in the list Δ. Thereafter, b will grow slowly in proportion to N^2, because the procedure will have to skip Increasingly more empty buckets before finding the first non-empty one.

This analysis indicates that there will be an optimal value of N which minimizes the running time. The optimum depends on the cost ratio B/D. In our tests, we found that the total time was minimized when N was about $2\sqrt{|\mathfrak{T}|}$ (an average of 0.25 entries per bucket).

8	5	4	4	4	5	8
5	4	1	1	1	4	5
4	1	0	0	0	1	4
4	1	0	0	0	1	4
4	1	0	0	0	1	4
5	4	1	1	1	4	5
8	5	4	4	4	5	8

(a)

45	41	33	27	34	42	46
37	21	17	11	18	22	38
29	13	05	03	06	14	30
25	09	01	00	02	10	26
31	15	07	04	08	16	32
39	23	19	12	20	24	40
47	43	35	28	36	44	48

(b)

Figure 2. (a) The squared cell distance function $\|(r, s)\|^2$, and (b) the bucket scan order implied by the list Δ, for the 7×7 cells nearest to the starting cell (at center).

3. Experiments

To measure the actual performance of our bucketing scheme, we used synthetic images produced by ray-tracing. (While synthetic images are not acceptable for validating complete VLPS algorithms, they are adequate to test the speed of the signature look-up step, since they provide a set $\mathbb{G}$ of realistic size and shape.)

In these tests, the scene consisted of a hemispherical smiley-like mask with convex eyes and concave mouth (both in low relief in order to avoid projected shadows), with various shades of matte gray finish. See figure 3(a). The gauge object was a sphere with white Lambertian finish; see figure 3(b).

The gauge images were subsampled to provide a signature table $\mathfrak{T}$ with 10219 entries. In all tests, the lighting setup was a single point source located very far from the scene, and the camera field-of-view was narrowed to provide near-parallel image projection. We used 6 different input image sets, varying the camera-to-light angle θ (either 10 or 45 degrees) and the number of input images m (either 3, 5, or 30). For each data set, we used bucket grids with two different sizes $N \times N$, either 202×202 or 143×143 (corresponding to average entry-to-bucket ratios $\kappa = |\mathfrak{T}|/N^2$ of 25%, 50%). We also processed each image set with $N = 1$, which is essentially equivalent to the brute-force nearest-match algorithm. While the larger values of N greatly reduced the

(a) (b) (c)

Figure 3. Scene (a) and gauge (b) used in the tests. Figure (c) is a 3D view of the height map obtained by integrating the scene slopes computed by the method described in this paper.

running time, the observation signatures returned by the table look-up procedure were always identical to those of the brute-force version.

Table 1 shows various average cost metrics for each table look-up operation: the number b of buckets $B[i', j']$ that were examined, the number d of table entries that were actually tested (i.e., the number of evaluations of $\mathrm{dist}(\boldsymbol{s}, \boldsymbol{g})$), and the look-up time t in microseconds. The tests were run on a standard PC with a 3GHz clock. The absolute time t obviously depends on the implementation.

Table 1. Average costs and operation counts of the table look-up procedure for various values of θ, m, and N. The entries with $N = 1$ represent sequential table search (without any bucket-grid speed-up).

θ	m	N	κ	t	d	b
10°	3	202	0.25	20.5	6.8	12.4
10°	3	143	0.50	22.9	11.8	11.2
10°	3	1	—	3987.5	10219.0	1.0
45°	3	202	0.25	18.0	3.5	10.0
45°	3	143	0.50	18.3	6.4	10.0
45°	3	1	—	3985.3	10219.0	1.0
10°	5	202	0.25	22.3	6.4	11.7
10°	5	143	0.50	25.1	11.1	10.9
10°	5	1	—	5620.1	10219.0	1.0
45°	5	202	0.25	29.0	10.5	45.2
45°	5	143	0.50	28.7	12.2	28.9
45°	5	1	—	5606.3	10219.0	1.0
10°	30	202	0.25	58.4	9.7	11.4
10°	30	143	0.50	76.4	16.7	10.8
10°	30	1	—	26637.9	10219.0	1.0
45°	30	202	0.25	74.9	12.5	51.2
45°	30	143	0.50	78.9	14.1	32.3
45°	30	1	—	26605.5	10219.0	1.0

Figures 4 and 5 show the sizes of the bucket lists $B[i, j]$ in two of those tests (with $m = 5$, $\theta = 10°$ and $m = 30$, $\theta = 45°$, respectively). In both cases we had $N = 202$, and therefore $\kappa = 25\%$. Note that, in most cases, the observation signatures are distributed fairly evenly over a substantial fraction of the grid.

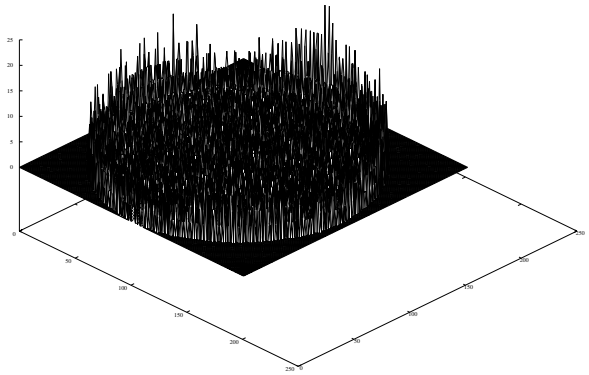

Figure 4. Bucket list lengths for $m = 5$ and $\theta = 45°$. The longest bucket has 8 entries.

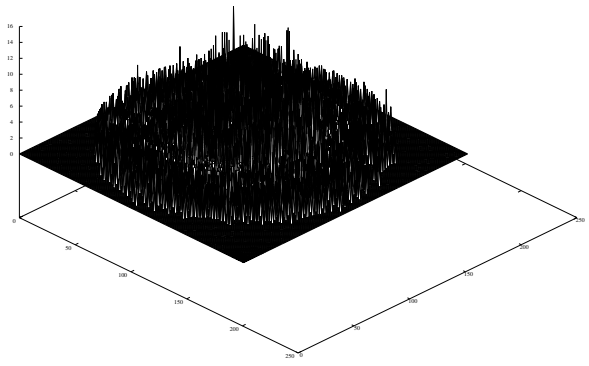

Figure 5. Bucket list lengths for $m = 30$ and $\theta = 10°$. The longest bucket has 5 entries.

4. Conclusions and future work

Our bucket-grid-based algorithm always yields the best-matching entry in the signature table. With the proper choice of the grid size N, the algorithm is considerably

faster than brute-force search (by orders of magniture) even for large values of m.

The two-dimensional grid that we use is consideraly more efficient, in both time and space, than the general m-dimensional nearest-neighbor data structures previously considered for this problem. Thanks to the optimal alignment of the grid, we obtain compact spherical enclosures for each bucket, which allow us to eliminate an entire bucket with a single distance comparison. Moreover, the 2D structure means that we need to scan only a few buckets (about 10) surrounding the hashed cell. Moreover, unlike previous grid schemes, our method always yields the best matching entry in the table (and not just a close approximation thereof).

We have restricted the input to monochromatic images only to simplify the exposition; but our 2D bucket-grid method works equally well for color images. If each image has c spectral bands (color channels), the *color observation vectors* $\boldsymbol{S}[p]$ and $\boldsymbol{G}[q]$ are the concatenation of c monochromatic OVs with m components each. As before, in order to recover the scene normal $\bar{s}[p]$ at a point p, we look for for a gauge point q such that the color signatures $\boldsymbol{s}[p]$ and $\boldsymbol{g}[q]$ match; except that the color signatures are obtained from the color OVs by normalizing the OV in each band separately. The color signatures then become points of the space $(\mathbb{S}^m)^c \subseteq \mathbb{R}^{mc}$; but they are still a 2-dimensional manifold in that space, and therefore can be organized by a single 2-D bucket grid.

Acknowledgements

This project was partly supported by research grants from CNPq (304581/2004-6 and 301016/92-5), CAPES, FAPESP and FAPERJ.

References

[1] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Y. Wu. An optimal algorithm for approximate neareast neighbor searching in fixed dimensions. *Journal of the Association of Computing Machinery*, 45(6):891–923, 1998.

[2] S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in presence of highlights and shadows. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25(10):1239–1252, Oct. 2003.

[3] R. Basri, D. Jacobs, and I. Kemelmacher. Photometric stereo with general unknown lighting. *International Journal of Computer Vision*, 72(3):239–257, 2007.

[4] A. Hertzmann and S. M. Seitz. Shape and materials by example: A photometric stereo approach. In *Proceedings IEEE CVPR 2003*, volume 1, pages 533–540, June 2003.

[5] A. Hertzmann and S. M. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying BRDFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1254–1264, Aug. 2005.

[6] P. Indyk and R. Motwani. Approximate nearest neighbour: Towards removing the curse of dimensinality. In *Proceedings of the 13th Annual ACM Symposium on Theory of Computing (STOC'98)*, pages 604–613, 1998.

[7] L. Shen, T. Machida, and H. Takemura. Efficient photometric stereo for three-dimensional surfaces with unknown BRDF. In *Proceedings of the 5th International Conference on 3-D Digital Imaginga nd Modeling (3DIM'05)*, pages 326–333, 2005.

[8] R. J. Woodham. Photometric method for determining suface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.

[9] R. J. Woodham. Determining surface curvature with photometric stereo. In *Proceedings of the 1989 IEEE International Conference on Robotics and Automation*, volume 1, pages 36–42, May 1989.

[10] R. J. Woodham. Gradient and curvature from the photometric stereo method, including local confidence estimation. *Journal of the Optical Society of America, Series A*, 11(11):3050–3068, 1994.

[11] L. Zhong and J. J. Little. Photometric stereo via locality sensitive high-dimension hashing. In *Proceedings of the Second Canadian Conference on Computer and Robot Vision (CRV'05)*, pages 104–111, 2005.