

STOCKPILE MONITORING USING LINEAR SHAPE-FROM-SHADING ON PLANETSCOPE IMAGERY

Marie d'Autume¹*, Alexandra Perry², Jean-Michel Morel¹, Enric Meinhardt-Llopis¹, Gabriele Facciolo¹

¹ Université Paris-Saclay, CNRS, ENS Paris-Saclay, Centre Borelli, 94235, Cachan, France

² Kayrros SAS

KEY WORDS: Shape-from-Shading, time series, satellite imagery, stockpiles monitoring

ABSTRACT:

The storage and management of stockpiles of materials is a fundamental process in large scale activities such as mining, civil engineering, and in the management of waste landfill sites. Following the evolution of stocks has always been important, and advancements in remote sensing technology are not only facilitating this, but also making it possible in near real-time. Nowadays, this monitoring appears to be performed almost exclusively using UAV based techniques. This paper proposes to apply a simple Shape-from-Shading method on low resolution satellite images provided by PlanetScope to monitor the evolution of the volume of stockpiles. The proposed Shape-from-Shading formulation makes it possible to handle occluding objects in the scene. The loss in accuracy due to the low resolution of PlanetScope images is well compensated by the daily revisit frequency and by the fact that spaceborne acquisitions require no human supervision. With satellites it is also easy to follow simultaneously several stocking sites all over the world. We test our method on two coal storage sites and demonstrate that the stockpiles are well detected.

1. INTRODUCTION

The storage and management of stockpiles of materials is a fundamental process in large scale activities such as civil engineering, mining and extraction, and in the management of waste landfill sites. The ability to follow in real time the evolution of stocks is more and more important today. For instance, following the evolution of coal stocks over two years in European ports allowed to monitor the switch from coal to gas as countries seek to reduce their carbon footprint to combat climate change (Hodges, 2019). On a more mundane level, knowing the evolution of coal stocks on a day-to-day basis makes it possible to anticipate the evolution of prices.

Continuous monitoring of stockpile volumes of goods is typically conducted manually using topographic surveying techniques (Fawzy, 2015). However, such methods have limitations regarding safety and access to the sites. Current trends to estimate these volumes use remote sensing tools such as airborne lidar scanners (Zhao, 2016), stereo-photogrammetry from planes and UAVs (He et al., 2019, Fawzy, 2015, Uysal et al., 2015, Raeva et al., 2016, Tucci et al., 2019, Salim et al., 2017), and even photogrammetry from satellite imagery (Schmidt et al., 2015, Demirel et al., 2011, Ge et al., 2007, Salim et al., 2017).

All these techniques are more precise, faster, safer and cheaper than the manual surveying measurements undertaken until a few years ago. The sheer number of commercial actors proposing UAV-based solutions proves the popularity of these methods. However, they fall short when the goal is to survey a large number of stocking areas around the world with a high frequency.

Satellite images have the advantage over UAV-based imagery that they require no supervision and can have a high revisit time. They also allow surveillance in areas of difficult access. As with

* Corresponding author:

marie.de-masson-d-autume@ens-paris-saclay.fr

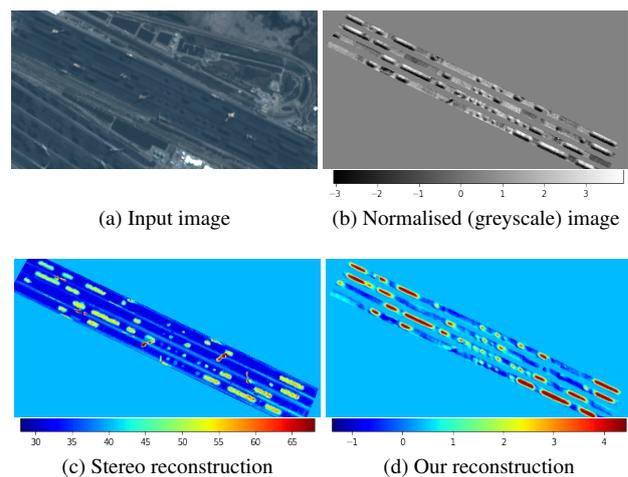


Figure 1. Stockpile monitoring with Shape-from-Shading. (a) shows the input PlanetScope image (3m pixel resolution), (b) the ROI after normalisation, (c) stereo reconstruction from SkySat imagery (0.8m resolution) and (d) output of the proposed method using Shape-from-Shading.

UAV imagery it is possible to acquire stereo pairs/triplets from high resolution satellites to recover 3D surface models (Facciolo et al., 2017, de Franchis et al., 2014). These techniques could be used to monitor the evolution of stockpiles with satisfying results. High resolution satellite stereo imagery is however expensive, which limits the applicability of this approach for continuous monitoring.

New micro-satellite constellations enable unprecedented systematic monitoring applications thanks to their wide coverage and short revisit capabilities. The PlanetScope constellation (Planet Team, 2017), for instance, provides daily global coverage at a resolution of 3.5m. These types of products are not adapted to do stereo reconstruction, nevertheless the stockpiles are clearly visible in the images.

As the goal of this work is to study the evolution of stocks, micro-satellites provide a good trade-off between accuracy and revisit time. Thus, in this paper we propose to use Shape-from-Shading techniques (Horn, Brooks, 1986) on these lower resolution PlanetScope images. This provides a much cheaper monitoring solution. Although the recovered shapes are less accurate than the ones acquired by stereo, the information they provide is enough to study the overall dynamics of the stocking sites (as illustrated in Fig.1). The use of mixed boundary conditions, as opposed to pure Dirichlet conditions, allows us to handle parts of the image that do not comply with the Lambertian model. We demonstrate on two sites that this technique allows us to automatically monitor stockpiles. We also include a validation of our volume computation using stereo-based methods on some of the dates.

The rest of the paper is organised as follows. In Section 2 we review the problem and the literature on shape-from-shading. In Section 3 we present the proposed method. Experiments are shown in Section 4 and we conclude in Section 5

2. SHAPE FROM SHADING

The Shape-from-Shading (SFS) problem consists in recovering the three-dimensional shape of a surface from one image of that surface. One of the main interests of this problem is that the input data are minimal as a single image is used. This is an important difference with other three-dimensional reconstruction methods such as stereo.

This problem was first formulated by Horn (Horn, 1970) as finding the solution of a nonlinear first-order Partial Differential Equation (PDE) called the brightness equation

$$I(x, y) = R(\mathbf{n}(x, y)). \quad (1)$$

Under the assumption of orthographic projection, the visible part of a scene is, up to a scale factor, a graph $z = h(x, y)$ where (x, y) is an image point. The unit normal \mathbf{n} can be easily expressed as

$$\mathbf{n}(\mathbf{x}) = \frac{1}{\sqrt{1 + |\nabla h|^2}} \left(-\frac{\partial h}{\partial x}, -\frac{\partial h}{\partial y}, 1 \right). \quad (2)$$

The light source is assumed to be at infinity with a unit direction vector $\mathbf{d} = (\alpha, \beta, \gamma)$. For a Lambertian surface of uniform albedo A ,

$$R(\mathbf{n}(x, y)) = A \mathbf{d} \cdot \mathbf{n}(x, y). \quad (3)$$

Assuming an albedo equal to 1, (1) can be written

$$I(x, y) \sqrt{1 + |\nabla h|^2} - \gamma = -\alpha \frac{\partial h}{\partial x} - \beta \frac{\partial h}{\partial y}. \quad (4)$$

This equation can be generalised to consider perspective projections (Lee, Kuo, 1994), a light source no longer at infinity but at the centre of projection (Prados, Faugeras, 2005) or non Lambertian scene (Bakshi, Yang, 1994). It has been proved in (Belhumeur et al., 1999) that when the lighting direction and the albedo of the surface are unknown, the same image can be obtained by a continuous family of surfaces.

A large number of numerical methods has been proposed for solving the Shape-from-Shading problem since its introduction in the 70s. Two review papers (Zhang et al., 1999, Durou et al., 2008) have extensively described these methods. They have

classified them into three main classes: the methods of resolution of PDEs, the optimisation methods and the methods approximating the brightness equation.

The first class contains characteristic strips expansion and approximation of viscosity solutions. The use of photogrammetric cues for 3D reconstruction was first proposed in (Van Diggelen, 1951) and was later developed in (Horn, 1975). This method has been essentially used for the theoretical study of the number of solutions of class C^2 of the eikonal equation that corresponds to a frontal light source at infinity. The notion of "viscosity solutions" is a more recent approach introduced in (Rouy, Tourin, 1992). Several algorithms have been proposed to compute these solutions, among which finite difference numerical methods (Rouy, Tourin, 1992, Lions et al., 1993) and Markov Chain approximation (Oliensis, Dupuis, 1993).

The optimisation methods are based on the variational formulations. They differ in the choice of the unknown, the regularisation term of the functional and the method of minimisation. A major reference for this class of methods is (Horn, Brooks, 1986). The method we use in our experiments belongs to this class.

Finally the third class covers the local methods and the linear methods. Local methods compute the normal at each point of the image independently of the other points. They usually assume the surface to be locally spherical (Hayakawa et al., 1994) or cylindrical (Wildey, 1986). Linear methods make a global (Kozera, Klette, 2000) or a local (Tsai, Shah, 1994) linear approximation of the reflectance function (3). A more detailed description of these different classes of methods can be found in the surveys previously mentioned.

3. OUR METHOD

In order to apply our method we first remove the useless images containing clouds or that are too blurry (Anger et al., 2019). Detecting clouds in images from pushbroom satellites can be achieved by exploiting the parallax effect on the different channels (Dagobert et al., 2019a). However, this method is not adapted to our input images as PlanetScope satellites do not use pushbroom sensors. But, since we are in possession of a time series of the same scene it is possible to use the algorithm proposed in (Dagobert et al., 2019b). Its approach exploits the redundancy of information acquired during revisits. On the other hand, the input images are not perfectly aligned. We therefore align them with the phase correlation method described in (Leprieux et al., 2007) and studied in detail in (Rais, 2016). Once the images are aligned, we delineate the coal storage areas. This has to be done manually but only once for each site.

3.1 Continuous Shape-from-Shading Formulation

Like most shape-from-shading methods we assume a Lambertian scene, an orthographic view and a far light source. The metadata of the satellite image provide the azimuth ϕ and elevation θ of the sun from which the light direction $\mathbf{d} = (\alpha, \beta, \gamma)$ can be computed:

$$\begin{cases} \alpha = \sin \phi \cos \theta \\ \beta = \cos \phi \cos \theta \\ \gamma = \sin \theta. \end{cases} \quad (5)$$

The goal is to find a digital elevation model $h(x, y)$ compatible with the observed data. As data, we have the orthoimage $I(x, y)$, a domain of interest Ω and a boundary condition for h given on the boundary $\partial\Omega$ of the domain. For the sake of simplicity—and because it suits our concrete application—we will always assume homogeneous Dirichlet boundary conditions ($h = 0$ on $\partial\Omega$); however, the techniques described below can be easily adapted to a boundary condition with arbitrary values. Furthermore, we also allow a Neumann-type boundary condition on a small part $\Gamma \subsetneq \Omega$ of the boundary, that is $\nabla h \cdot \mathbf{n} = 0$, where \mathbf{n} is the normal vector to $\partial\Omega$.

The shape-from-shading equation for a Lambertian model with albedo A and constant ambient light B has the following form

$$\begin{cases} I = A \frac{-\alpha h_x - \beta h_y + \gamma}{\sqrt{1 + h_x^2 + h_y^2}} + B & \text{on } \Omega \\ h = 0 & \text{on } \partial\Omega \setminus \Gamma \\ \frac{\partial h}{\partial \mathbf{n}} = 0 & \text{on } \Gamma. \end{cases} \quad (6)$$

This is a nonlinear first-order equation that can be studied by the method of characteristics. In the general case it has no unique solution, or it may even fail to have a solution at all, and further conditions are required. Here we propose a particular simplification of this equation that is suitable for our application.

3.2 Linearised Shape-from-Shading

The proposed method relies on the following two assumptions:

1. The sun is low in the horizon (strictly lower than 45°);
2. The entire terrain is lit by the sun (no shadows).

Notice that these assumptions taken together imply that the maximal slope of the terrain is 45° . These assumptions are of course not valid in general, but they are quite realistic for the application to heliosynchronous images of loose material stockpiles (whose angle of repose is always smaller than 45° , attained for wet sand). These assumptions allow for a great simplification of our equation, leading to the linearised solver. The first consequence of the assumptions is that the typical “inversion ambiguity” of shape-from-shading models is not possible: you cannot invert the whole terrain profile of a solution and obtain another solution like when sun is very high. The second consequence, regarding small slopes, is that we can assume that $\|\nabla h\| \ll 1$ so that $\sqrt{1 + h_x^2 + h_y^2} \approx 1$, which turns the nonlinear equation into a, mostly trivial, linear equation. The proposed technique, described below for the discrete case, consists in solving this linear equation and then iteratively refine the solution by solving a sequence of analogous linear problems that approach the non-linear one.

Setting $\sqrt{1 + h_x^2 + h_y^2} = 1$, the shape from shading equation (6) becomes linear. After arranging the terms we get the following problem:

$$\begin{cases} \vec{\alpha} \cdot \nabla h = f & \text{on } \Omega \\ h = 0 & \text{on } \Omega \setminus \Gamma \\ \mathbf{n} \cdot \nabla h = 0 & \text{on } \Gamma, \end{cases} \quad (7)$$

where $\vec{\alpha} = (\alpha, \beta)$ is the projected direction of the sun and $f = (I - B)/A - \gamma$ is the new data term.

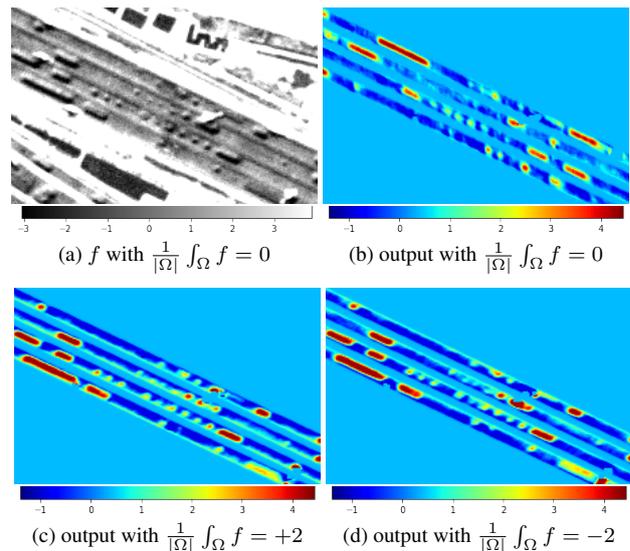


Figure 2. The output of the shape-from-shading is robust to changes to the mean of f . The mean normalised image is shown in (a) and the resulting surface in (b). In (c) and (d) we modified this mean to observe that it results in a slight slope on the reconstructed surface. Although some differences can be observed along the boundaries of the bands the stockpiles are barely impacted.

The shape-from-shading equation is based on the assumption that the intensity depends only on the surface normal, illumination and albedo, with the albedo often put to one when unknown. For our input images we can see that the intensity is also strongly dependent on the atmospheric conditions as seen in the top left columns of Figures 4 and 5.

From (1) and (3) we have

$$I = \frac{1}{\sqrt{1 + |\nabla h|^2}} (-\alpha \frac{\partial h}{\partial x} - \beta \frac{\partial h}{\partial y} + \gamma). \quad (8)$$

Introducing the variable $f(I, h, \gamma) = I\sqrt{1 + |\nabla h|^2} - \gamma$ we rewrite it as

$$f(I, h, \gamma) = -\alpha \frac{\partial h}{\partial x} - \beta \frac{\partial h}{\partial y}. \quad (9)$$

Note that for any point on an horizontal surface we have $|\nabla h| = 0$, thus $f(I, h, \gamma) = 0$. By assuming the ground to be horizontal and the stockpiles to be more or less regular, enforcing a zero mean for f ensures the integrability condition. Figure 2 shows that the reconstruction is robust to changes of the mean. If the values corresponding to the ground in f are zero the ground appears flat in the output image. If this doesn't hold true the SFS will reconstruct the ground as a slope. This is visible only when comparing the values along the sides of the lanes of stored coal.

While we do not give any proof of the convergence or soundness of this simplified procedure for larger slopes (when the approximation $\|\nabla h\| \ll 1$ does not hold), the experiments show that it consistently produces surfaces of a reasonable shape, and quantitatively comparable to those obtained by stereo methods.

3.3 Study of the one-dimensional case

The linearised equation (7) is a first-order linear equation on a compact domain. In general, these equations never have a solu-

tion unless the data satisfies a strict integrability condition. As we explain below, the integrability condition can be enforced by applying an affine contrast change to the input image. Any first-order PDE can be rewritten as a family of independent ODE along curves that cover the domain, called *characteristic curves of the equation*. In our case, the characteristic curves are straight line segments in the direction of the sun. Along each of these curves, the ODE to solve has the following form (we consider first the homogeneous boundary condition):

$$\begin{cases} u'(x) = f(x) & x \in [a, b] \\ u(x) = 0 & x \in \{a, b\}, \end{cases}$$

where $f(x)$ is the intensity at the position x and $u(x)$ is the height profile to be found. In general, this equation has no solution that satisfies both boundary conditions, as can be seen by integrating on the whole domain $[a, b]$, which gives $\int_a^b f = \int_a^b u' = u(b) - u(a) = 0$, so the integral of f on $[a, b]$ must vanish for this equation to have a solution. The condition $\int_a^b f = 0$ is the integrability condition required on the input data. The solution is then found trivially since the equation is directly integrable: $u(x) = \int_a^x f$.

The fact that the image intensities are known up to an arbitrary contrast change comes in handy: we can change f so that it satisfies the integrability condition.

The general case (in 1D) is thus given by an equation of the following form

$$\begin{cases} u'(x) = Rf(x) + S & x \in [a, b] \\ u(a) = p \\ u(b) = q, \end{cases}$$

where R and S are constants that we can choose so that the integrability condition holds. The equation is directly integrable and the solution is

$$u(x) = R \int_a^x f + S(x - a) + p.$$

Notice that $u(a) = p$ is automatically satisfied, and the condition $u(b) = q$ is written as $q = R \int_a^b f + S(b - a)$, which determines the unique possible value of S :

$$S = \frac{q - p}{b - a} - R \frac{1}{b - a} \int_a^b f.$$

The value of $R \neq 0$ is thus arbitrary, as it amounts to fixing the (unknown) albedo of the surface. We can set it to 1 for simplicity.

To solve the linearized problem we use a variational approach with the regularisation term proposed in (Brooks, Horn, 1985). This leads to minimising an energy of the following form

$$E(h) = \int_{\Omega} \|\vec{\alpha} \cdot \nabla h - f\|^2 + \lambda \int_{\Omega} \|\nabla h\|^2. \quad (10)$$

The Euler-Lagrange equation of this quadratic functional is

$$D^2(h)(\vec{\alpha}, \vec{\alpha}) + \lambda \Delta h = \vec{\alpha} \cdot \nabla f \quad (11)$$

The term $D^2(h)(\vec{\alpha}, \vec{\alpha})$ is the second-derivative of h in the direction $\vec{\alpha}$, which does not “see” variations of h perpendicular to $\vec{\alpha}$.

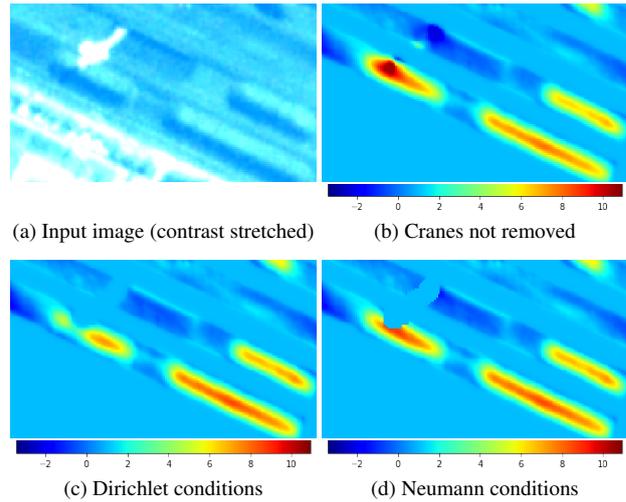


Figure 3. The necessity of mixed boundary conditions. The cranes seen in (a) violate the constant albedo hypothesis which results in an erroneous reconstruction (b). The result in (c) is obtained by imposing a Dirichlet condition everywhere and in particular along the occluding objects boundaries, which affects the estimation of the occluded stockpile. The result in (d) uses Neumann conditions along the occluding objects boundaries resulting in better reconstruction.

These variations are smoothed-out by the laplacian term, which is an isotropic regularizer. This Euler-Lagrange PDE is a linear elliptic equation that has a unique solution provided there is at least one Dirichlet datum on each connected component of $\partial\Omega$.

3.4 The effect of occluding objects and the necessity of mixed boundary conditions

An important hypothesis of our SFS method is that the albedo is constant and that there are no shadows. This is not true in the presence of occluding objects. In our experiments these occluding objects are cranes, that tend to be much brighter than the coal stockpiles we observe. A simple thresholding of the intensity images detects them well in most cases. Once detected, their boundaries are treated differently from the manually drawn boundaries of the lanes. Assuming that the ground on which the coal is deposited is flat it is coherent to enforce Dirichlet boundary conditions. On the other hand when a crane is above a stockpile the reconstructed height should not be imposed to be zero on that part of the boundary (Dirichlet boundary conditions), but instead to have zero normal derivative (Neumann boundary conditions).

The importance of using these mixed boundary conditions is noticeable in Figure 3, particularly on the middle lane. The top row corresponds to the normalised image and the reconstruction without removing the cranes. This leads to a depression on the ground of the second lane and a peak on the bottom lane. Imposing boundary Dirichlet conditions everywhere (bottom left of the figure), the former aberrations disappear but the left stockpile appears carved in the middle. This comes from enforcing a zero height near the middle of the pile. Adding Neumann conditions on the occluded part (bottom right) leave the reconstruction more freedom along the crane boundaries and the results are much more realistic (bottom right).

3.5 Discrete formulation

To ease the discretisation of equation (7), and especially the Neumann boundary condition, we use the algebraic graph formalism for finite difference schemes (Lézoray, Grady, 2012). Briefly, the whole image domain is modelled as a grid graph of size $W \times H$. This graph has $n = W \times H$ vertices and $m = 2WH - W - H$ edges. Notice that $m \approx 2n$. The region of interest Ω is thus a subset of the vertex set, identified by a binary vector $\mathbf{1}_\Omega \in \mathbf{R}^n$, and the boundary condition is specified by the knowledge of $h \in \mathbf{R}^n$ outside of Ω . Scalar fields correspond to real-valued functions defined over the vertices (i.e. vectors of \mathbf{R}^n), and vector fields to functions defined on the edges (i.e. vectors of \mathbf{R}^m). We define the following three matrices: The identity matrix $I \in \mathcal{M}_{n,n}(\mathbf{R})$, the mask operator of the region of interest $M = \text{diag}(\mathbf{1}_\Omega) \in \mathcal{M}_{n,n}(\mathbf{R})$, and the signed incidence matrix of the graph, $\nabla = \mathcal{M}_{m,n}(\mathbf{R})$. The matrix ∇ is the operator for computing first-order derivatives by forward differences.

Notice that if $u \in \mathbf{R}^n$ represents a scalar field, then its discrete gradient is $\nabla u \in \mathbf{R}^m$, which is a vector field according to the above convention. From these operators we define the divergence operator $-\nabla^\top \in \mathcal{M}_{n,m}(\mathbf{R})$ and the Laplacian $\Delta = -\nabla^\top \nabla \in \mathcal{M}_{n,n}(\mathbf{R})$. We also define a *centering operator* $C = \frac{1}{2}|\nabla| \in \mathcal{M}_{n,m}$. It can be checked readily that the pointwise scalar product of two discrete vector fields $F, G \in \mathbf{R}^m$ corresponds naturally to $F \cdot G = C^\top(F \odot G)$, where \odot is the component-wise product (Hadamard product) on \mathbf{R}^m . This allows to compute the directional derivatives that appear in the continuous equation. Notice that the Hadamard product can be written as a matrix-vector product $F \odot G = \text{diag}(F)G$. To enforce Neumann boundary conditions, we remove the lines of the matrix ∇ that corresponds to the edges that we want to remove.

This formalism allows a direct translation of the linear equations and variational problems defined above in a continuous setting. For example, the linear shape-from-shading equation $f = \vec{\alpha} \cdot \nabla u$ is discretised as $f = C^\top \text{diag}(S_\alpha) \nabla h$, where $S_\alpha \in \mathbf{R}^m$ is a (constant) vector field in the direction $\vec{\alpha}$. This is a square linear system with matrix $N_\alpha := C^\top \text{diag}(S_\alpha) \nabla \in \mathcal{M}_{n,n}(\mathbf{R})$. The matrix N_α computes the discrete directional derivative in the direction of $\vec{\alpha}$. The Dirichlet boundary conditions g are linear constraints that can be easily added to the problem:

$$\begin{cases} MN_\alpha h = Mf \\ (I - M)h = (I - M)g \end{cases} \quad (12)$$

or $(MA + I - M)h = Mf + (I - M)g$. This is a linear problem on the unknown h . As happens in the continuous setting, this problem has no solution (its matrix is singular) unless the data term happens to satisfy a further integrability condition. Worse than that, even when there is a solution, the computations along each characteristic line are performed independently, resulting in a structured noise on the solution. To solve this problem we prefer to use the variational formulation, that corresponds to the associated normal equations of the linear problem, and where it is easier to add a regularisation term.

The energy functional $E(h)$ is discretised using the same formalism, to give the following discrete energy:

$$E(h) = \|N_\alpha h - f\|^2 + \lambda \|\nabla h\|^2$$

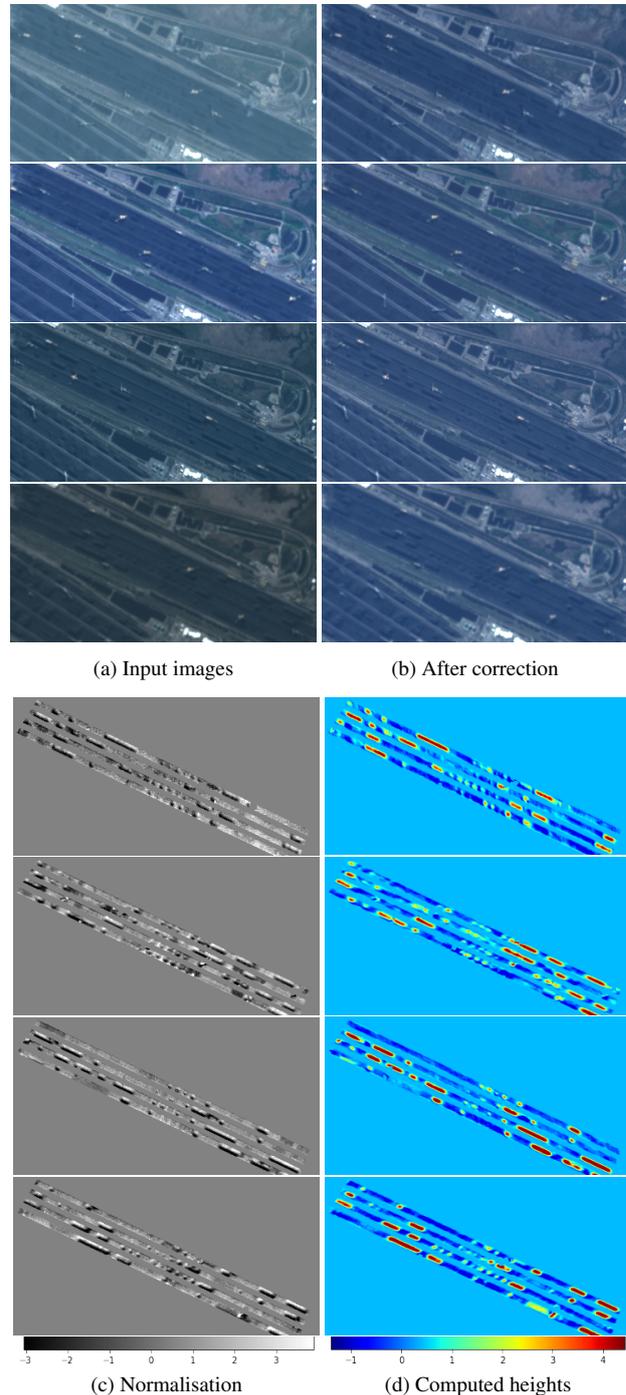


Figure 4. Input images before and after relative radiometric correction (Hessel et al., 2020), normalised images and predicted heights from our shape-from-shading method for the first site (2.58×1.58 km).

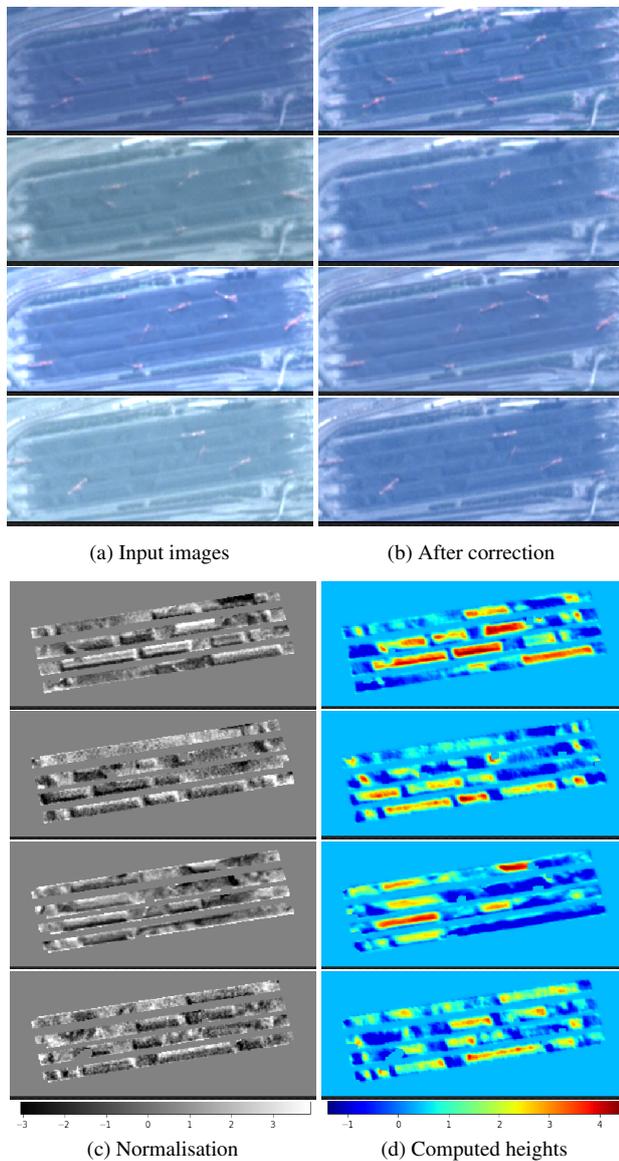


Figure 5. Input images before and after relative radiometric correction (Hessel et al., 2020), normalised images and predicted heights from our shape-from-shading method for the second site, more challenging because smaller (1.17×0.46 km).

this is a positive-definite quadratic form on the variable h . The unique minimum is attained at the point h solution of $(N_{\alpha}^T N_{\alpha} - \lambda \Delta)h = N_{\alpha} f$. Notice that, apart from the regularisation term depending on λ , this is just the normal equation of the first-order discrete problem. The proposed method consists in the solution of the regularised problem by this equation. The whole method can be implemented in about 10 lines of Matlab/Octave code.

So far, we have explained how to solve the linear shape-from-shading equation, which is just an approximation. To deal with the *nonlinear* problem we use the following strategy: first we solve the linear problem to obtain an initial approximation h^0 . Then we re-weight the data term f on the linear problem by the function $\sqrt{1 + (h_x^0)^2 + (h_y^0)^2}$, which results in a new linear problem of the same form that has a unique solution h^1 . Thus we compute a sequence h^n of approximate solutions that, if it converges, the limit must be the solution of the nonlinear problem. In practice, it seems that after a couple of iterations the solution does not change much.

4. EXPERIMENTS

For our experiments we use PlanetScope (Planet Team, 2017) imagery. The PlanetScope constellation is made of approximately 130 small satellites (form factor of $10 \times 10 \times 30$ cm) imaging the entire Earth's landmass every day. The satellites carry 3-band or 4-band frame cameras and fly at 475 km on sun-synchronous orbits whose constant local solar time is between 9:30 and 11:30 am. The provided images are orthorectified with a GSD resolution of 3m.

Our dataset consists of a pair of temporal series of two sites over seven months, which amounts to about 100 images per site after removing cloudy days. We do not have access to the ground truth for the sites we studied. Still we were provided DEM reconstruction of the first site using stereo photogrammetry (de Franchis et al., 2014) obtained from SkySat imagery along with volumes (in cubic meters) computed from these DEMs. These reconstructions will serve as references for our validation.

4.1 Qualitative evaluation

We applied our method to the time series on the two sites and we observed that the stockpiles were correctly estimated in most cases. The images for the first site (Figure 4) cover an area of 2.58×1.58 km while the images for the second site (Figure 5) only cover an area of 1.17×0.46 km. Thus the second site is much more challenging and we don't expect the same quality of results for both sites. Indeed our method struggles in reconstructing the shapes of the second site that presents much smaller stockpiles. By looking at the enhanced images we can confirm that it is actually quite hard to see the stockpiles in the images.

4.2 Analysis of time series

The ground is lower than the rails along which the cranes move. During the reconstruction we assumed the height of these rails to be zero. The first step to compute volumes is to add a global additive constant to restore the ground level to zero. The vertical scale of our SFS reconstruction is arbitrary but coherent in the whole time-serie. We add a multiplicative constant to this scale such that the maximum altitude of our stockpiles matches the one from the reference stereo reconstructions. The volumes

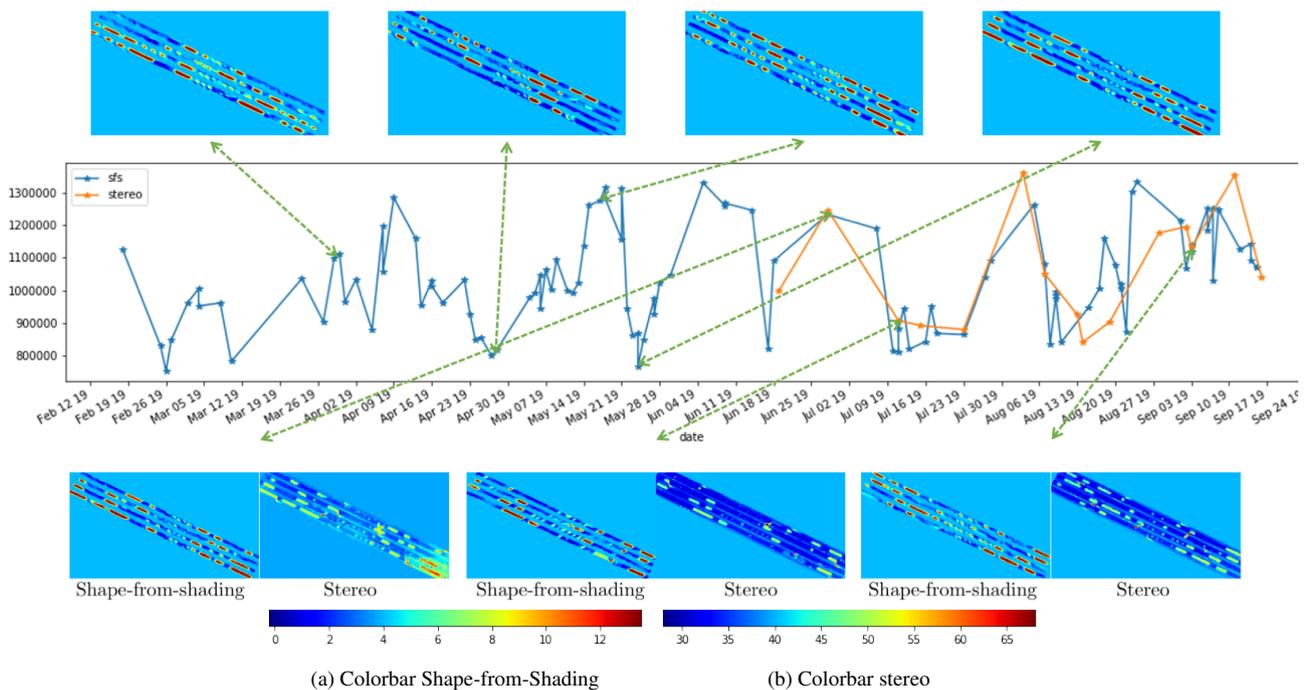


Figure 6. Volume tracking in a time series. The plot represents the estimated volumes using the algorithm described in Section 3 on a time series of PlanetScope images spanning 7 months. The Shape-from-Shading heights are scaled by a constant so that the maximum altitude of our stockpiles matches the one from a reference stereo reconstruction (about 15m). The results for some highlighted dates are shown above the plot. The results of our method are consistent with the volumes obtained using stereo reconstruction and are obtained with a much higher frequency. The images shown below the plot allow the comparison on 3 dates of our result with 3D reconstructions. Note that the altitudes in the 3D reconstructions are expressed relative to the ellipsoid.

are then computed from the reconstructed shapes by integration on the ROI after clipping the first and last percentile to get rid of aberrant values. Then we multiply this result by 9, which is the size of a pixel for a GSD resolution of 3m.

In Figure 6 we represent the volumes estimated using the algorithm described in Section 3 on a time series of PlanetScope images spanning 7 months. The results for some highlighted dates are shown above the plot. We can see that the estimated volume is lower when less stockpiles are detected. For the stereo reconstructions we only have access to three months of data with about one acquisition per week. The images shown below the plot allow the comparison on 3 dates of the result of the 3D reconstruction with our method. Note that the PlanetScope and the SkySat images are often acquired at different hours or days. So the volumes obtained by stereo reconstruction are only an approximation of the ground truth. Taking this into account we see from Figure 6 that our estimated volumes are quite realistic.

5. CONCLUSION

Nowadays stockpile monitoring appears to be performed almost exclusively using UAV based techniques. In this paper we have shown that applying a simple Shape-from-Shading method on low resolution PlanetScope images can be sufficient to monitor the evolution of the volume of these stockpiles. The unquestionable loss in accuracy is well compensated in our opinion by the much higher revisit frequency and less human intervention. It is also easy to follow simultaneously several stocking sites all over the world. Thus this method appears very well suited to the problem of large scale stocks monitoring. In future work,

more sophisticated Shape-from-Shading algorithms should be considered and the detection of occluding objects could be improved. Our method needs to be tested on more stocking sites. Access to more stereo reconstructions would also allow a quantitative evaluation of our results in addition to our qualitative evaluation.

ACKNOWLEDGEMENTS

Work partly financed by the Office of Naval Research grant N00014-17-1-2552, CNES MISS project, DGA Astrid project «filmer la Terre» n° ANR-17-ASTR-0013-01.

REFERENCES

- Anger, J., de Franchis, C., Facciolo, G., 2019. Assessing the sharpness of satellite images: Study of the planetscope constellation. *IGARSS*, IEEE.
- Bakshi, S., Yang, Y.-H., 1994. Shape from shading for non-lambertian surfaces. *ICIP*, IEEE.
- Belhumeur, P. N., Kriegman, D. J., Yuille, A. L., 1999. The bas-relief ambiguity. *IJCV*, 35(1), 33–44.
- Brooks, M. J., Horn, B. K. P., 1985. Shape and source from shading.
- Dagobert, T., Grompone von Gioi, R., de Franchis, C., Morel, J.-M., 2019a. Cloud Detection by Disparity Phase Analysis for Pushbroom Satellite Imagers. *IPOL preprint*.

- Dagobert, T., Morel, J.-M., de Franchis, C., Grompone von Gioi, R., 2019b. Visibility detection in time series of planet-scope images. *IGARSS*, IEEE.
- de Franchis, C., Meinhardt-Llopis, E., Michel, J., Morel, J.-M., Facciolo, G., 2014. An automatic and modular stereo pipeline for pushbroom images. 2(3), 49–56.
- Demirel, N., Emil, M. K., Duzgun, H. S., 2011. Surface coal mine area monitoring using multi-temporal high-resolution satellite imagery. *IJCG*, 86(1), 3–11.
- Durou, J.-D., Falcone, M., Sagona, M., 2008. Numerical methods for shape-from-shading: A new survey with benchmarks. *CVIU*, 109(1), 22–43.
- Facciolo, G., De Franchis, C., Meinhardt-Llopis, E., 2017. Automatic 3d reconstruction from multi-date satellite images. *CVPR Workshops*.
- Fawzy, H., 2015. The Accuracy of Determining the Volumes Using Close Range Photogrammetry. *IOSR-JMCE*, 12, 10–15.
- Ge, L., Chang, H.-C., Rizos, C., 2007. Mine subsidence monitoring using multi-source satellite SAR images. *PE & RS*, 73(3), 259–266.
- Hayakawa, H., Nishida, S., Wada, Y., Kawato, M., 1994. A computational model for shape estimation by integration of shading and edge information. *Neural Networks*, 7(8), 1193–1209.
- He, H., Chen, T., Zeng, H., Huang, S., 2019. Ground Control Point-Free Unmanned Aerial Vehicle-Based Photogrammetry for Volume Estimation of Stockpiles Carried on Barges. *Sensors*, 19(16), 3534.
- Hessel, C., Grompone R., M. J.-M., G., F., P., A., , de Franchis C., 2020. Relative Radiometric Normalization Using Several Automatically Chosen Reference Images for Multi-Sensor, Multi-Temporal Series. *ISPRS Annals*.
- Hodges, J., 2019. Coal is piling up in europe as utilities prefer natural gas. <https://www.bloomberg.com/news/articles/2019-09-27/coal-is-piling-up-in-europe-as-utilities-prefer-natural-gas>. Accessed: January 2020.
- Horn, B. K. P., 1970. Shape from shading: A method for obtaining the shape of a smooth opaque object from one view.
- Horn, B. K. P., 1975. Obtaining shape from shading information.
- Horn, B. K. P., Brooks, M. J., 1986. The variational approach to shape from shading. *CVGIP*, 33(2), 174–208.
- Kozera, R., Klette, R., 2000. Evaluation of numerical solution schemes for differential equations. *Performance Characterization in Computer Vision*, Springer.
- Lee, K. M., Kuo, C. C. J., 1994. Shape from shading with perspective projection. *CVGIP*, 59(2), 202–212.
- Leprince, S., Barbot, S., Ayoub, F., Avouac, J.-P., 2007. Automatic and precise orthorectification, coregistration, and sub-pixel correlation of satellite images, application to ground deformation measurements. *TGRS*, 45(6), 1529–1558.
- Lézoray, O., Grady, L., 2012. *Image processing and analysis with graphs: theory and practice*. CRC Press.
- Lions, P.-L., Rouy, E., Tourin, A., 1993. Shape-from-shading, viscosity solutions and edges. *Numerische Mathematik*, 64(1), 323–353.
- Oliensis, J., Dupuis, P., 1993. A global algorithm for shape from shading. *ICCV*, IEEE.
- Planet Team, 2017. Planet Application Program Interface: In Space for Life on Earth. San Francisco, CA. <https://api.planet.com>.
- Prados, E., Faugeras, O., 2005. Shape from shading: a well-posed problem? *CVPR*, IEEE.
- Raeva, P. L., Filipova, S. L., Filipov, D. G., 2016. Volume Computation Of A Stockpile-a Study Case Comparing Gps And Uav Measurements In An Open Pit Quarry. *ISPRS Archives*, 41, 999–1004.
- Rais, M., 2016. Fast and accurate image registration. Applications to on-board satellite imaging. PhD thesis.
- Rouy, E., Tourin, A., 1992. A viscosity solutions approach to shape-from-shading. *Journal on Numerical Analysis*, 29(3), 867–884.
- Salim, P. M., Jais, M. A., Sahriman, N., Samad, A. M., Abbas, M. A., Maarof, I., Tarmizi, N., 2017. Monitoring quarry areas using remote sensing techniques. *CSPA*, IEEE.
- Schmidt, B., Malgesini, M., Turner, J., Reinson, J., 2015. Satellite monitoring of a large tailings storage facility. *Tailings and Mine Waste Management for the 21st Century*.
- Tsai, P.-S., Shah, M., 1994. Shape from shading using linear approximation. *Image and Vision computing*, 12(8), 487–498.
- Tucci, G., Gebbia, A., Conti, A., Fiorini, L., Lubello, C., 2019. Monitoring and Computation of the Volumes of Stockpiles of Bulk Material by Means of UAV Photogrammetric Surveying. *Remote Sensing*, 11(12), 1471.
- Uysal, M., Toprak, A. S., Polat, N., 2015. DEM generation with UAV Photogrammetry and accuracy analysis in Sahitler hill. *Measurement*, 73, 539–543.
- Van Diggelen, J., 1951. A photometric investigation of the slopes and the heights of the ranges of hills in the Maria of the moon. *Bulletin of the Astronomical Institutes of the Netherlands*, 11, 283.
- Willey, R. L., 1986. Radarclinometry for the Venus radar mapper. *PE&RS*, 52, 41–50.
- Zhang, R., Tsai, P.-S., Cryer, J. E., Shah, M., 1999. Shape-from-shading: a survey. *TPAMI*, 21(8), 690–706.
- Zhao, S., 2016. 3D real-time stockpile mapping and modelling with accurate quality calculation using voxels. PhD thesis.