Computers & Graphics 34 (2010) 441–448

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Technical Section

Simulating the appearance of sandy landscapes

Bradley W. Kimmel *, Gladimir V.G. Baranoski

Natural Phenomena Simulation Group, David R. Cheriton School of Computer Science, University of Waterloo, 200 University Ave. W., Waterloo, Ontario, Canada N2L 3G1

ARTICLE INFO

Article history:
Received 4 January 2010
Received in revised form
23 February 2010
Accepted 19 April 2010

Keywords:
Sand
Realistic rendering
Spectral reflectance
BRDF

ABSTRACT

Sand is one of the most complex materials found in nature. Undeniably the correct modelling of its appearance attributes (such as hue, lightness, and glossiness) is essential to the realistic image synthesis of a wide range of outdoor scenes. Despite this central role, to date, few simulation efforts have been specifically directed to this ubiquitous material. In this paper, we present a modular framework for simulating the appearance of sandy landscapes. It is based on the use of a comprehensive light transport model specifically designed for granular materials like sand, and robust numerical reconstruction methods. While the former provides the physical basis for the generation of predictive results, the latter add efficiency to entire simulation process by enabling the use of analytical formulae to represent the spectral and spatial (scattering related) appearance attributes of sand. The fidelity and usefulness of the proposed framework are demonstrated through several image sequences depicting sand appearance variations resulting from changes of mineralogical characteristics and environmental conditions.

1. Introduction

Sand is a natural particulate material found in a variety of environments, from beaches and deserts on Earth to dune fields on Mars. In order to render realistic images depicting such sandy landscapes, one has to carefully account for its optical properties. Although an extensive amount of research has been directed toward the simulation of material appearance [1], few efforts have been specifically aimed at sand. In this article, we present a compact framework for the efficient modelling of its appearance. Since our approach is geared up to the synthesis of predictive images of real landscapes, we employ physically meaningful parameters to characterize different sand samples and control the processes that determine their appearance attributes.

As any material, the appearance of sand is determined by the spectral and spatial distribution of light interacting with it [2]. The former affects colour-related characteristics, and it is usually measured in terms of reflectance. The latter affects characteristics such as reflection haze (glossiness) and retroreflection, and it is usually measured in terms of bidirectional reflectance distribution function (BRDF). The physically based modelling of these radiometric quantities is a challenging problem, however. The difficulties arise from the array of factors affecting the light interactions with sand, such as the shape, size and distribution of sand grains (particles), and the presence of different mineral contaminants and other inorganic substances, just to name a few.

In studies involving the geometric modelling of sand surfaces, such as the works by Valette et al. [3] or Onoue and Nishita [4,5], the rendering of sand is usually achieved through the use of texture maps or functions. Such an approach is bound by the availability of texture data with the resolution, illumination and appearance characteristics (e.g., hue and lightness) appropriate for the target scene. Few works have expanded upon this approach, however. For example, Kass and Miller [6] proposed the application of a wetness map to scale reflectance values in order to achieve the darkening effect caused by the presence of water in sand samples. Oren and Nayar [7] presented a generalization of the Lambertian model in order to simulate the BRDF of sand surfaces. They compared the output from their model to measurements performed on a sand sample. Their simulations, however, were restricted to the spatial domain, i.e., spectral reflectance data are required as input to their model. We remark that the measured spectral reflectance data sets currently available for sand are limited to a narrow range of illumination and environmental conditions. Subsequently, Jensen et al. [8] used an extension of the Henyey–Greenstein phase function to adjust the degree of forward scattering and achieve varying levels of wetness in sand samples. It has been demonstrated that this function bears no relation to the sand characterization parameters [9]. Hence, its use in this context precludes the generation of predictive images of sand samples. More recently, Soulié et al. [10] used 3D Voronoï diagrams to model the appearance of compact granular materials such as granite.

Recently, a spectral light transport model for sand (SPLITS) has been proposed [11]. This model uses Monte Carlo techniques to simulate light interactions with individual sand particles generated on the fly. It uses as input the physical parameters and

* Corresponding author.
E-mail addresses: bwkimmel@cs.uwaterloo.ca (B.W. Kimmel), gvgbaran@curumin.math.uwaterloo.ca (G.V.G. Baranoski).
URLS: http://www.npsg.uwaterloo.ca/people/brad (B.W. Kimmel), http://www.curumin.uwaterloo.ca/~gvgbaran (G.V.G. Baranoski).

0097-8493/$ - see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cag.2010.04.002

mineralogical characteristics describing a given sand sample, and outputs the radiometric quantities that provide its *measurement of appearance* [2], namely spectral reflectance and BRDF. Its predictability and accuracy have been illustrated by the close qualitative and quantitative agreement between modelled results and actual measured data. Similar to other stochastic models, SPLITS is computationally expensive since many trials are required to obtain asymptotically convergent results. Consequently, it is not suitable for online applications demanding high interactivity rates. It can, however, be used offline to compute accurate spectral and spatial radiometric data for sand under different illumination and environmental conditions. Such data can then be incorporated into a more efficient modelling framework.

This article presents such a framework, which builds upon reliable numerical reconstruction techniques to allow for the efficient prediction of the full measurement of appearance of sand samples. This framework takes as input the physical parameters describing the sample of interest and outputs its corresponding spectral BRDF, which, in turn, can be used on the realistic rendering of sandy landscapes. It is worth noting that the proposed framework can employ different models. The selection of SPLITS was motivated by the fact its target application is sandy landscapes, and this model is the most comprehensive model for light interaction with sand available in the literature. Although similar approaches have been used to model other key components of outdoor scenes, such as the formulation proposed by Preetham et al. [12] to fit data from simulations of atmospheric scattering responsible for the colour of the sky, to the best of our knowledge, no similar approach for modelling the appearance of sand has been proposed in the graphics literature to date.

2. Background

Sand is a particular type of soil. It is composed primarily of weathered rock, immersed in a medium composed of air and water called the *pore space* [14]. The mineral component of a soil is classified according to its particle size distribution [13]. A series of sieves are used to separate particles falling into size ranges as indicated in Table 1, with clay being the smallest, followed by silt and sand. The fraction of these soil separates by mass determines the soil *texture*, as indicated in Fig. 1. Particles larger than 2 mm in diameter are considered gravel and do not contribute to the determination of soil texture. Sand consists of at least 85% sand-sized particles by mass (Fig. 2).

The most common mineral constituent of sand is quartz [15]. Although quartz is colourless in pure form, its colour may be affected by the presence of trace amounts of contaminants [16]. As depicted in Fig. 3, these impurities are crucial in determining the appearance of sand. The constituent primarily responsible for determining the appearance of sand is iron oxide [18], which may be present in several forms. Hematite, or *red ochre*, imparts a red hue, and is often found in tropical climates. Goethite, one of the most common mineral colourants of soils, is responsible for the yellows and browns. Magnetite is black and is often present in beach and river

Table 1
Soil separates (particle size classes) defined by the United States Department of Agriculture [13].

Name	Range (mm)
Sand	0.05–2.0
Silt	0.002–0.05
Clay	< 0.002

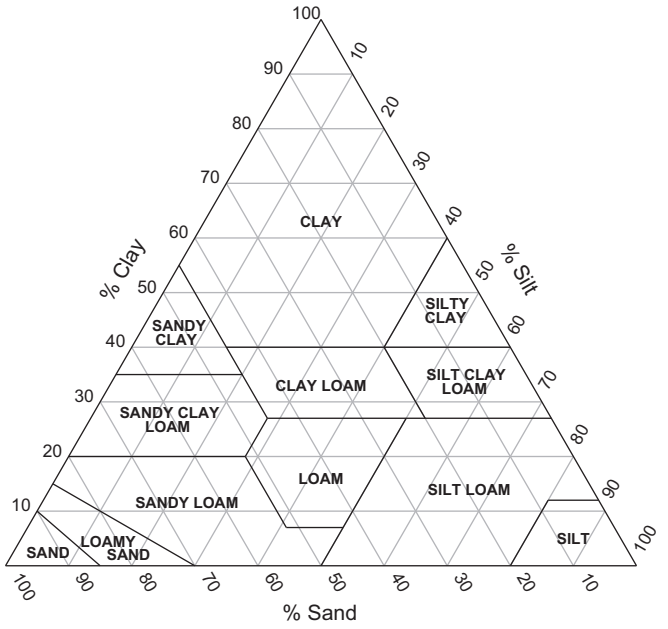

Fig. 1. The soil texture triangle depicting the percentages (by mass) of clay, silt and sand-sized particles comprising the various soil textural classes [13].

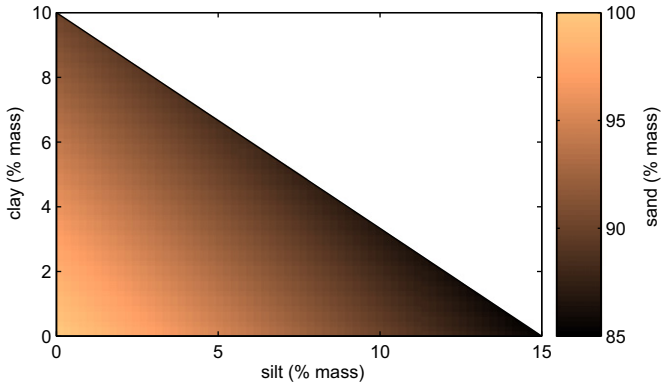

Fig. 2. The shaded region indicates the range for the relative proportions (by mass) of clay, silt, and sand-sized particles in sand [13].

sands [19]. Additionally, the presence of water darkens sand, primarily by reducing the contrast between the refractive index of the pore space and the refractive index of quartz, which is approximately 1.5 in the visible region of the electromagnetic spectrum. This reduction in contrast decreases the angle of refraction and reduces Fresnel reflection at quartz interfaces, thereby introducing a forward scattering bias. Other factors, such as particle size and shape, also influence the appearance of sand [20].

SPLITS [11] uses standard Monte Carlo techniques to simulate the appearance of a particulate medium, i.e., to predict its spectral BRDF. In its formulation, sand particles are modelled as randomly oriented spheroidal particles. Although the authors describe how arbitrary particle size and shape distributions may be used, it was found that approximation by prolate spheroids was sufficient to obtain good quantitative and qualitative reconstructions of the reflectance properties of sand [11]. These particles are randomly distributed throughout a medium of water and air (the *pore space*) contained within the half-space below a plane boundary. The particles themselves consist of quartz cores, possibly mixed with hematite or goethite, or covered in a thin coating of hematite or goethite in a kaolinite matrix. The model parameters include the concentrations of

various iron oxides: hematite, goethite, and magnetite, as well as the fraction of the pore space occupied by water, known as the *degree of saturation*. Additional parameters describe the geometrical arrangement of the mineral constituents [11].

Light interaction with the simulated sand medium proceeds using standard Monte Carlo techniques. Once a ray penetrates the extended boundary and enters the medium, light interaction is simulated with the particles contained therein. In a traditional ray tracing approach, these particles would be stored individually, potentially imposing a large memory footprint. Techniques such as spatial subdivision could be used to accelerate the ray–particle intersection process. Alternatively, the bulk scattering properties of the medium may be approximated using pre-computed phase functions [21]. Instead, the SPLITS model uses a hybrid approach, relying on stochastic methods to determine the location of the next ray-particle intersection, as illustrated in Fig. 4. The distance to and orientation of the nearest particle along the path is determined randomly. Light interaction with that particle is then simulated explicitly using standard techniques. The particle is subsequently discarded. This process is repeated until the ray is absorbed or is scattered back outside the boundary of the medium.

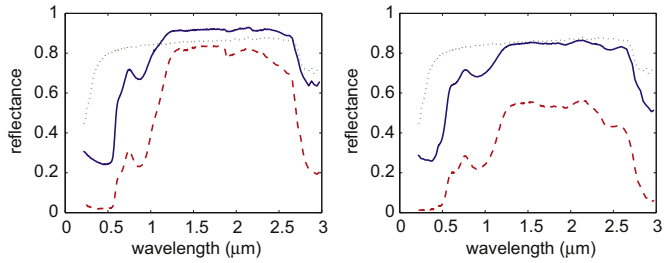

Fig. 3. A demonstration of the effect caused by iron oxide contamination in quartz. The dotted line in both plots represents the spectral reflectance of a pure quartz sample [17]. The dashed line represents the spectral reflectance of a pure iron oxide sample. The solid line represents the spectral reflectance of an artificially prepared mixture consisting of 98% (by mass) quartz and 2% iron oxide. *Left*: The iron oxide is hematite [17]. *Right*: The iron oxide is goethite [17]. The spectral features of hematite (respectively, goethite) are clearly visible in the mixed samples.

As with many techniques that use Monte Carlo methods, many repetitions are required for the result to converge within acceptable bounds. For SPLITS, or other Monte Carlo material models, on the order of 10^6–10^9 rays may be needed, depending on the requirements of the application. Although the precise timing varies according to the input parameters used and the number of trials needed, 10^8 rays typically requires 145 min using a single core of an Intel Xeon 2.8 GHz Quad-Core processor. To overcome this, the authors proposed an analytical formulation for predicting sand spectral signatures [22]. This formulation predicts the spectral directional-hemispherical reflectance [23,24],

$$\rho(\omega_i) = \int_{\Omega_o} f_r(\omega_i,\omega_o)\, d\omega_o^{\perp} \tag{1}$$

of a sand sample, given its physical and mineralogical characteristics, and is suitable for applications demanding high interactive rates. Whereas the Monte Carlo approach may require several hours to compute a spectral signature for a single sample, an analytical formulation is capable of computing several spectral signatures instantaneously in comparison. It does not, however, account for the spatial distribution of the reflected light (i.e., the BRDF), and is therefore not suitable for computer graphics applications. The framework proposed here builds upon the techniques described therein to allow for the prediction of the full measurement of appearance [2] of a given sand sample.

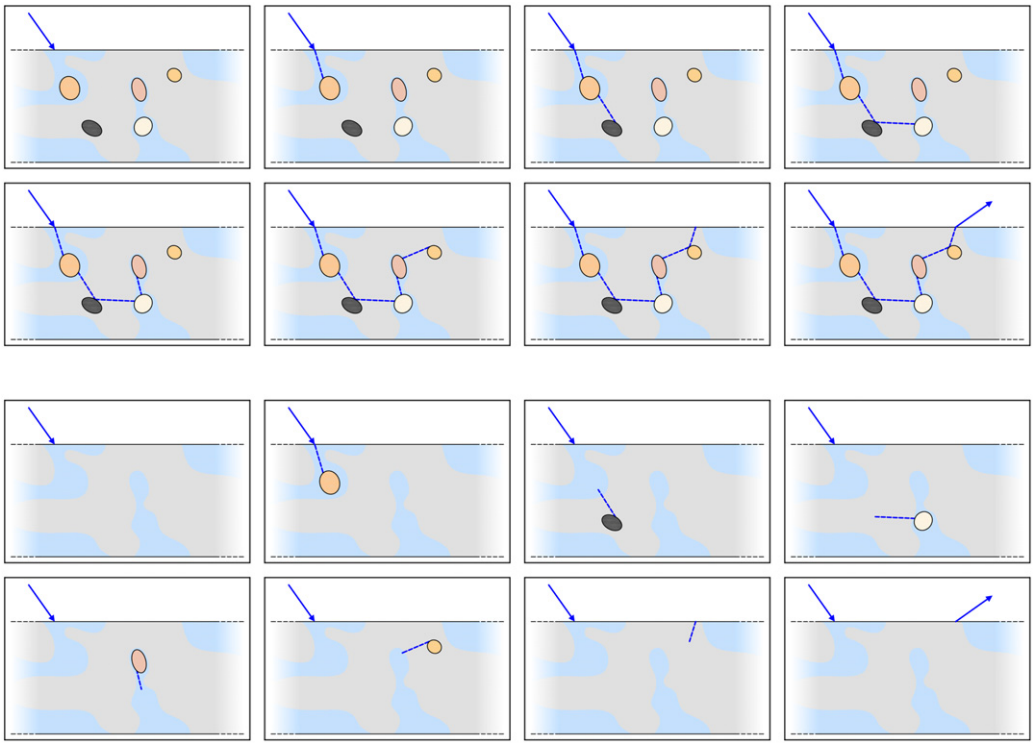

Fig. 4. A comparison between a traditional ray tracing approach and that employed by the SPLITS model [11] is depicted. *Top sequence*: For a traditional ray tracing approach, the particles are explicitly stored. Several ray-particle intersection tests may be required for each leg in the path. *Bottom sequence*: For the SPLITS model, the distance to the next particle is generated randomly according to a path length distribution. The orientation of the particle, as well as a point on the particle surface, are also generated randomly. Light interaction with the particle is simulated. The ray is either absorbed in the particle or scattered. In the latter case, the particle is discarded and a new one is generated. This process is repeated until the ray is absorbed or scattered outside of the medium.

3. Framework

The proposed analytical framework, illustrated in Fig. 5, is comprised of two main components. The spectral component predicts the directional-hemispherical reflectance based on the physical characteristics of the sand sample. This process is described in detail in [22]. For the sake of completeness, it is also outlined in this paper. The spatial component predicts the distribution of reflected light over the hemisphere as a function of incident direction, wavelength, and the physical characteristics of the sample. The resulting BRDF generated by the spatial component is scaled to integrate to 1 over the upper hemisphere for a zero degree incident angle. The final BRDF is then constructed by combining the BRDF generated by the spatial component with the directional-hemispherical reflectance given by the spectral component.

3.1. Spectral component

The purpose of the spectral component is to obtain a predictor for the spectral directional-hemispherical reflectance of sand samples, based on the SPLITS model, as a function of the physical and mineralogical properties of the sample. First, the characterization data for 3000 samples were chosen randomly from within the domain of the SPLITS model. The model was then applied to yield the corresponding directional-hemispherical reflectance curves, using an incident angle of zero degrees. Principal component analysis (PCA) and regression techniques were applied, as described below, to obtain the predictor. These techniques were selected due to their well known efficacy, which has been demonstrated in practical applications across many fields, including computer graphics (e.g., [25,26]). The variation in the directional-hemispherical reflectance with incident angle is accounted for by the spatial component of this framework to be described later.

3.1.1. Principal component analysis

Principal component analysis was applied to the directional-hemispherical reflectance curves generated by SPLITS. This process yields an orthogonal set of eigenvectors $\mathbf{u}_i$ indicating the directions of decreasing variance in the reflectance data. These eigenvectors span the same space as the original data and may be used as basis vectors to reconstruct the original reflectance curves. The analysis also yields the corresponding principal components that indicate the linear combination of the basis vectors required for this reconstruction.

Since the eigenvectors are ordered by decreasing variance in the original data, it is typically assumed that eigenvectors beyond a given threshold, k, represent noise in the data. These directions may therefore be discarded by projecting the original reflectance curves onto the space spanned by the first k eigenvectors $\mathbf{u}_i,\ldots,\mathbf{u}_k$. This process yields a predictor for the spectral direction-hemispherical reflectance, $\mathbf{r}$, of the sample as a function of corresponding principal components $\tilde{\mathbf{r}}$, having the form

$$\mathbf{r} \approx U_k \tilde{\mathbf{r}}, \tag{2}$$

where U_k is the $n \times k$ matrix having columns consisting of the first k eigenvectors, and n is the number of wavelengths sampled. In our experiments, we found 16 wavelengths, between 400 and 700 at 20 nm intervals, to be sufficient to accurately represent the reflectance curves produced by the SPLITS model. We also found that four eigenvectors were sufficient to reproduce over 99.9% of the variance in the training set.

3.1.2. Regression

Since the principal components $\tilde{\mathbf{r}}$ do not have meaningful physical or mineralogical interpretation themselves, it is necessary to relate these to the physical and mineralogical properties of the corresponding sand sample. To accomplish this, a non-linear regression analysis was performed. The original characterization data $\mathbf{x}$ is first mapped into a higher dimensional feature space via a non-linear function $\phi(\mathbf{x})$

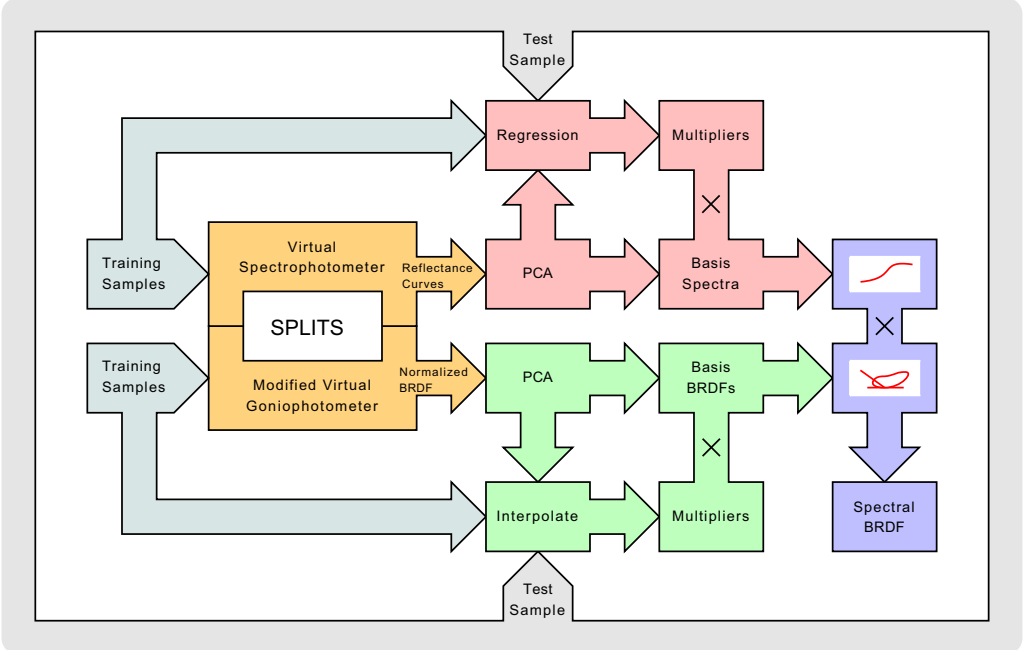

Fig. 5. Diagram illustrating the proposed framework for predicting the appearance of sandy surfaces. The top half represents the spectral component. The bottom half represents the spatial component. The two components are then combined to construct the spectral BRDF.

defined using a cubic monomial basis in the components of $\mathbf{x}$. Linear regression was then performed in this m-dimensional feature space. This yields an $m \times k$ matrix W which may be applied to obtain a predictor for the principal components $\tilde{\mathbf{r}}$ corresponding to a given set of characterization data. The predictor has the form

$$\tilde{\mathbf{r}} \approx W^t \phi(\mathbf{x}). \tag{3}$$

3.1.3. Reconstruction of the spectral reflectance curve

Combining Eqs. (2) and (3) yields

$$\mathbf{r} \approx U_k W^t \phi(\mathbf{x}). \tag{4}$$

This predicts the spectral directional-hemispherical reflectance curve $\mathbf{r}$, corresponding to a zero degree incident angle, for a sand sample having the specified set of characterization data $\mathbf{x}$.

3.2. Spatial component

In constructing a predictor for the spatial distribution of the light reflected from a sand sample having a given set of physical and mineralogical characteristics, we first restrict our attention to subset of the domain of the SPLITS model that includes the region of interest for a particular application. The characterization data for a number of samples are chosen, forming a mesh that covers this subset of the domain. A modification of virtual goniophotometry [27], described below, is applied to obtain the BRDF corresponding to each sample. Each BRDF is normalized to unit reflectance at normal incidence. Principal component analysis and interpolation, also described below, are applied to obtain a predictor for the normalized BRDF, given a new set of characterization data.

3.2.1. Virtual BRDF measurement

To compute the BRDF corresponding to a given a set of characterization data within the domain of the SPLITS model, a variation of virtual goniophotometry [27] is employed. For a virtual goniophotometer, the hemisphere above the point on the surface to be measured is divided into a number, r, of patches. A counter, n_i, is required for each patch for $i=1,\ldots,r$. A particular incident angle is selected and the model under consideration is applied N times. Typically, N is selected to be on the order of 10^8 rays, although this may vary depending on the nature of the BRDF and the desired accuracy. During each trial, the incident light may be scattered or absorbed. If it is scattered, the counter n_i corresponding to the patch containing the scattered direction is incremented. The BRDF for each outgoing patch i is estimated by

$$f_i = \frac{n_i}{N\omega_p},$$

where ω_p is the projected solid angle of the patch. This entire process is repeated for several incident angles to obtain a full BRDF.

In this framework, rather than fixing the incident direction, it is chosen at random, uniformly over the hemisphere, prior to each trial. The hemisphere of incoming directions is divided in an analogous manner to the hemisphere of outgoing directions. In our experiments, we divided each hemisphere into patches of equal solid angles. The hemisphere was divided into 31 stacks in the polar direction. The top stack (representing the normal direction) was not divided any further. The remaining stacks were each divided into 30 slices along the azimuthal direction, yielding a total of $r=901$ patches. Counters $n_{i,j}$ are required for each pair consisting of an incoming patch i and outgoing patch j, for $i=1,\ldots,r$, $j=1,\ldots,r$. If the incident ray is scattered, the corresponding counter $n_{i,j}$ is incremented. If the incident ray is absorbed, a separate counter $n_{i,0}$ is incremented. The BRDF corresponding to a pair (i,j) is estimated using

$$f_{i,j} = \frac{n_{i,j}}{\omega_p^{(j)} \sum_{k=0}^{r} n_{i,k}},$$

where $\omega_p^{(j)}$ is the projected solid angle of the outgoing patch.

Due to the nature of the SPLITS BRDF, the $O(r^2)$ data points generated by the above process may be greatly reduced. The isotropism of the SPLITS BRDF, as well as Helmholtz reciprocity, are exploited to reduce the error in the two dimensional array. This was accomplished by replacing each $f_{i,j}$ with an average of all the values $f_{i',j'}$ which are known to be identical to $f_{i,j}$. The resulting BRDF is then normalized to unit reflectance at normal incidence by computing

$$\hat{f}_{i,j} = \frac{f_{i,j}}{\sum_{k=1}^{r} \omega_p^{(k)} f_{i,k}}.$$

Finally, the two dimensional array $\hat{f}_{i,j}$ is reduced to a linear array $\mathbf{f}$ of the m unique values.

3.2.2. Principal component analysis

Principal component analysis was then applied to the linear arrays $\mathbf{f}_i$ representing the deflated BRDFs. This results in an orthogonal set of basis vectors $\mathbf{v}_i$ and the corresponding principal components $\tilde{\mathbf{f}}_i$ that indicate the linear combination of the basis vectors required to reconstruct the original BRDFs. As was done earlier, the eigenvectors beyond a given threshold ℓ are discarded. The first ℓ vectors $\mathbf{v}_1,\ldots,\mathbf{v}_\ell$ are then used in reconstruction. Given the principal components $\tilde{\mathbf{f}}$ corresponding to a new sample, the BRDF may be reconstructed by computing

$$\mathbf{f} \approx V_\ell \tilde{\mathbf{f}}, \tag{5}$$

and then expanding the resulting compressed BRDF to obtain the full BRDF, where V_ℓ is the $m \times \ell$ matrix $\mathbf{v}_1,\ldots,\mathbf{v}_\ell$ as its columns.

3.2.3. Interpolation of principal components

To determine the principal components to use for a new sample, recall that the selected domain points form a mesh. Thus, the characterization data $\mathbf{x}$ of the new sample may be expressed as an affine combination

$$\mathbf{x} = \sum_{i=1}^{n} \alpha_i \mathbf{x}_i, \quad \sum_{i=1}^{n} \alpha_i = 1 \tag{6}$$

of the selected domain points $\mathbf{x}_i$. We therefore combine the principal components $\tilde{\mathbf{f}}_i$ using the same interpolants. That is,

$$\tilde{\mathbf{f}} = \sum_{i=1}^{n} \alpha_i \tilde{\mathbf{f}}_i. \tag{7}$$

3.2.4. Reconstruction of the spatial component

Given the characterization $\mathbf{x}$ of a particular sand sample, we construct the spatial component of the BRDF first by determining the appropriate affine combination $\alpha_i^{(\mathbf{x})}$ of the $\mathbf{x}_i$ with which to represent $\mathbf{x}$ in the form of Eq. (6). We then combine the corresponding principal components using Eq. (7). Expanding the resulting linear array yields the estimated BRDF $f_{i,j}$ for each pair consisting of an incoming patch i and an outgoing patch j. To construct a continuous BRDF $f_{\mathbf{x}}(\mathbf{u},\mathbf{v})$ from these data points, the centroid $\mathbf{u}_i$ is computed for each patch i. We then set $f_{\mathbf{x}}(\mathbf{u}_i,\mathbf{u}_j)=f_{i,j}$. Values for other directions are computed by interpolation over the hemisphere.

3.3. Combining the spectral and spatial components

Using the spectral and spatial components of the proposed framework, we may now predict the spatial and the spectral

distributions of reflected light from a sand sample having defined physical and mineralogical characteristics. The spectral directional-hemispherical reflectance for zero degree incident light, described by Eq. (4), yields a spectral reflectance curve at 16 wavelengths between 400 and 700 nm, inclusively. The spatial distribution $f_{\mathbf{x}}$ for a sand sample with characterization data $\mathbf{x}$ is constructed as described above. It has the form of a BRDF that has been scaled so that

$$\int_{\Omega} f_{\mathbf{x}}(\mathbf{n},\mathbf{v})\cos\theta\, d\mathbf{v} = 1,$$

where $\mathbf{n}$ is the normal to the surface. We may now combine the spatial and the spectral components to yield a predictor,

$$U_k W^t \phi(\mathbf{x}) f_{\mathbf{x}}(\mathbf{u},\mathbf{v})$$

for the BRDF of a sand sample with physical characterization $\mathbf{x}$. This representation yields an explicit formula for the computation of the BRDF, enabling it to be incorporated into a wider variety of rendering algorithms than a pure Monte Carlo approach, ranging from basic path tracing [28] to more advanced techniques [29,21].

4. Results

Using the proposed framework, we can predict the appearance of a sand sample with given physical and mineralogical characteristics. To demonstrate the spectral aspect, we have used the proposed framework to colourize an image of a sandy surface, varying the characterization parameters within the domain of the SPLITS model. The proposed framework was also used to animate two scenes featuring sandy landscapes. The first portrays a desert scene as the Sun traverses the sky over the course of one day. The second depicts waterlines along a beach landscape which dry over several hours.

4.1. Spectral component

To demonstrate the spectral component of the proposed framework, a close up image of a sandy surface has been colourized. The physical and mineralogical properties have been varied within the domain of the SPLITS model.

4.1.1. Colourization

To give the image the appearance of a surface with the given characteristics, the spectral component of the framework was used for colourization. The corresponding reflectance curve, $r(\lambda)$, was computed according to the framework. Given the spectral power distribution $e(\lambda)$ for a particular illuminant, we may obtain the CIE XYZ tristimulus values,

$$X = \int_0^{\infty} e(\lambda) r(\lambda) x(\lambda) d\lambda,$$

$$Y = \int_0^{\infty} e(\lambda) r(\lambda) y(\lambda) d\lambda,$$

$$Z = \int_0^{\infty} e(\lambda) r(\lambda) z(\lambda) d\lambda,$$

where x, y, and z are the CIE colour matching functions [30]. For the images presented here, we used the CIE D50 illuminant. The CIE 1964 standard observer was used for the colour matching functions. The XYZ tristimulus values were then converted to standard RGB [30], first by performing the transformation

$$\begin{pmatrix} R' \\ G' \\ B' \end{pmatrix} = \begin{pmatrix} 3.2410 & -1.5374 & -0.4986 \\ -0.9692 & 1.8760 & 0.0416 \\ 0.0556 & -0.2040 & 1.0570 \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}.$$

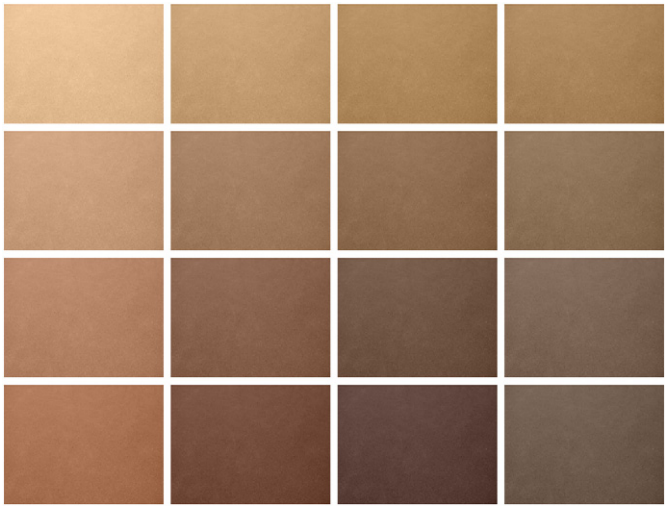

Fig. 6. Images depicting the variation of sand colour as the concentration of iron oxides and the hematite/goethite ratio are changed. The former parameter increases to the right, while the latter increases downward. The images were generated using the proposed framework.

Each of R', G', and B' was then clamped to the interval [0, 1] and gamma corrected using

$$R = 1.055R'^{\gamma} - 0.055,$$

(and similarly for B and G) where $\gamma = 1/2.4$.

The image was then colourized using the RGB triple. This was accomplished by computing the mean RGB values for the original image. This mean was then subtracted from each pixel and the RGB triple computed using the proposed framework was added.

4.1.2. Images

The techniques described above were used to colourize an image of a patch of a sandy surface. Figs. 6 and 7 depict the variations in colour corresponding to the spectral reflectance predicted using the proposed framework.

In Fig. 6, the concentration of iron oxide is varied from 0.001 to 0.05 in the horizontal direction. The hematite/goethite ratio,

$$\frac{\text{hematite}}{\text{hematite}+\text{goethite}},$$

varies from 0 to 1 in the vertical direction.

In Fig. 7, the hematite/goethite ratio is varied from 0 to 1 in the horizontal direction and the degree of saturation (water content) is varied from 0 to 1 in the vertical direction.

4.2. Animations

As a demonstration of potential applications for the proposed framework, we have created two animations: the first depicting a desert scene as the Sun traverses the sky over the course of a day, and the second depicting waterlines along a beach drying over time. Selected frames from these animations are presented in Figs. 8 and 9, respectively. The full animations are available with the online version of this manuscript or at the authors' website.[1]

4.2.1. Desert scene

An image of a desert landscape was used as the basis from which the animation was constructed. A rough approximation of the desert landscape (a flat surface) was used in a ray tracing simulation. The BRDF computed using this framework was

[1] http://www.npsg.uwaterloo.ca/misc/sand

sampled using a cosine distribution. For our purposes, we found this to be sufficient. However, at increasingly grazing angles, this approach may become inefficient. For such applications, one might opt to sample the hemisphere according to a closer approximation of the BRDF. This could be obtained using, for example, a combination of generalized cosine lobes [31] or a scheme based on elliptical contours [32]. The daylight model developed by Preetham et al. [12] was used to illuminate the scene, with the Sun passing over the sky over the course of a 12 hour period.

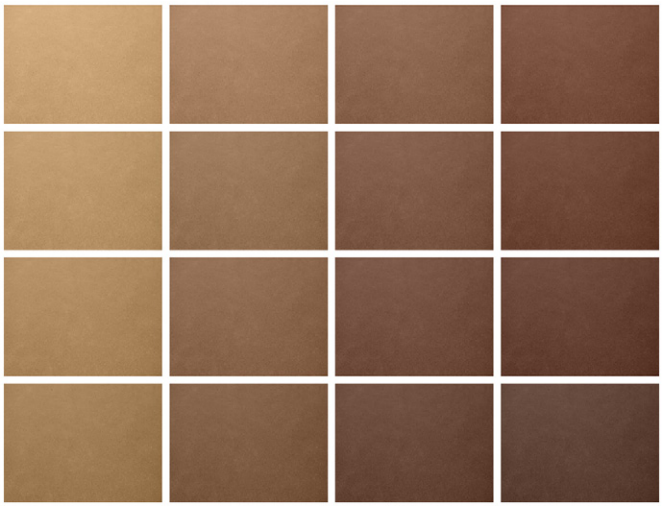

Fig. 7. Images depicting the variation of sand colour as the hematite/goethite ratio and the degree of saturation are changed. The former parameter increases to the right, while the latter increases downward. The images were generated using the proposed framework.

The resulting spectral image was converted to standard RGB using the same technique described above. The following process was then applied to incorporate the texture of the sand into the image. The ground in the original image was isolated and blurred heavily. The original image was then divided by the blurred image pixel by pixel, yielding a quotient image. Each frame in the animation was then multiplied by the resulting quotient image. More sophisticated techniques could be applied if required for a particular application.

4.2.2. Beach scene

To create the beach scene, two waterlines were created using Brownian motion curves and overlaid onto a flat surface representing the beach landscape. These were used to construct a texture map for controlling the degree of saturation (water content) of the sand. The degree of saturation is gradually decreased over the course of the day (as the Sun traverses the sky). It was assumed that the sand in this scene is otherwise homogeneous, and that the spatial distribution of the BRDF did not vary with wavelength. The adapted virtual goniophotometer was applied using 16 different values for the degree of saturation,

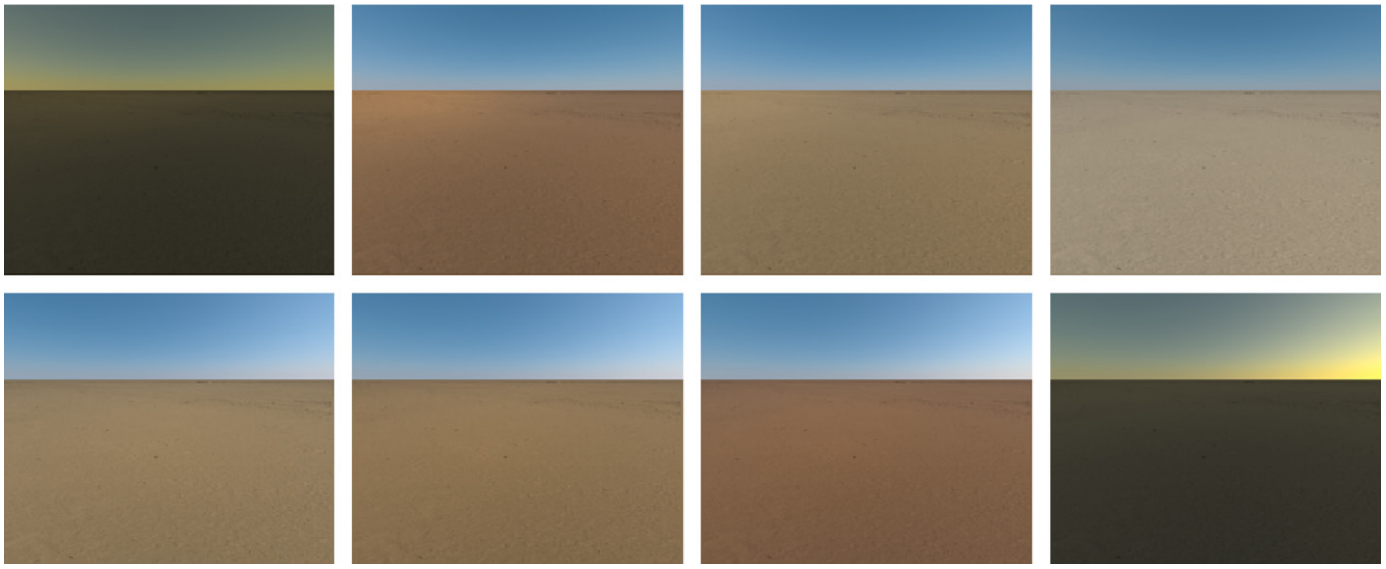

Fig. 8. Selected frames of an animation sequence depicting a desert scene as the Sun traverses the sky over the course of a day.

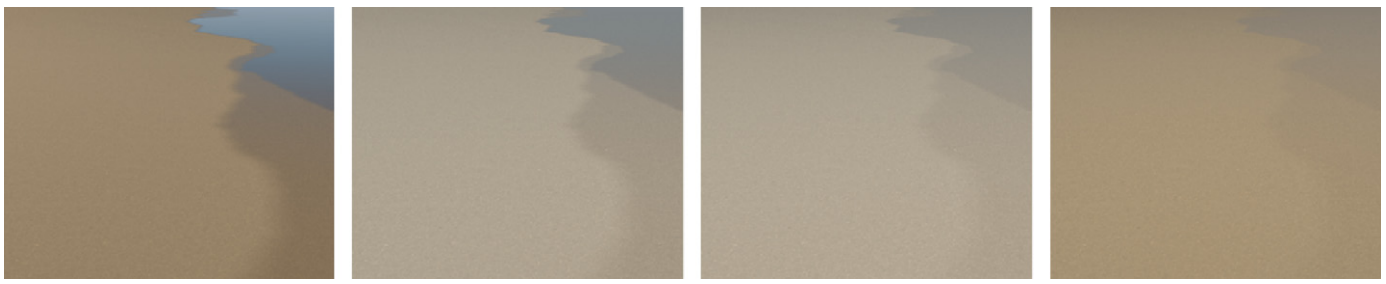

Fig. 9. Selected frames of an animation sequence depicting a beach scene as waterlines dry over several hours during the course of a day.

ranging from zero to one. The resulting BRDFs were reduced to three principal components using the techniques described above. A grainy texture was applied as a post-process to give the sand a more natural appearance.

5. Conclusion and future work

In this paper, we have presented an efficient framework for simulating the appearance of sandy landscapes. Its design is based on the use of high fidelity spectral and spatial data computed offline using a comprehensive Monte Carlo light transport model for sand [11]. This data are then reconstructed on demand using analytical formulae derived using robust numerical methods. This approach allows the simulation process to be controlled by physically meaningful parameters while enabling the predictable representation of sand appearance attributes at interactive rates. The effectiveness of the proposed framework was demonstrated through sets of rendered images depicting sand appearance changes triggered by different physical and environmental parameters.

The modular structure of the proposed framework allows a straightforward incorporation of application driven refinements. For example, to construct a predictor for the spatial distribution of the light reflected from a sand sample, we adapted virtual goniophotometric techniques and interpolated the resulting data directly. This approach was chosen for its relative simplicity and to achieve maximum flexibility to represent different forms of scattering. By exploiting the isotropism of the model in question, the memory footprint was reduced by a significant factor. Depending on the nature of the BRDF required for a particular application, however, one may opt for a different representation, such as spherical wavelets, or fitting to an analytical BRDF model.

As future work, we intend to further explore the modularity of the proposed framework in order to extend its scope of applications to the predictive rendering of other types of landscapes. For instance, more recent techniques in statistical analysis, such as independent component analysis [33], could be used in place of PCA. Additionally, despite the ubiquitousness of sand in the natural environment, there are many other classifications of soils that may be considered. In our research, we have accounted for the most common mineral constituents of sand. This framework, however, could be extended to include sands with other mineral compositions or to account for other contaminants, such as oil. We would also like to extend SPLITS itself to simulate additional granular materials. We remark that the main purpose of the proposed framework is generate results with a high accuracy/cost ratio. To achieve this goal, it uses as a basis of comprehensive light transport model (SPLITS) specifically designed and validated]for the material at hand (in this case sand), and employs analytical techniques to overcome the inherent computational costs of such a model. In the case of another natural material, SPLITS can be replaced by a model specifically designed and validated for this material. Although the reconstruction parameters (e.g., number of eigenvectors employed to represent the variance of the training set) may need to be adapted to the characteristics of a different material, the underlying framework structure remains the same. Ultimately, we expect to implement an online interactive system to generate predictable representations of appearance attributes for a number of natural materials under various environmental conditions.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.cag.2010.04.002.

References

[1] Dorsey J, Rushmeier H, Sillion F. Digital modeling of material appearance. Burlington, MA: Morgan Kaufmann/Elsevier; 2007.
[2] Hunter R, Harold R. The measurement of appearance, 2nd ed. New York, NY: John Wiley and Sons; 1987.
[3] Valette G, Prévost S, Lucas L, Léonard J. SoDA project: a simulation of soil surface degredation by rainfall. Computers and Graphics 2006;30(4):494–506.
[4] Onoue K, Nishita T. A method for modeling and rendering dunes with wind-ripples. In: Proceedings of the Pacific conference on computer graphics and applications, 2000. p. 427–8.
[5] Onoue K, Nishita T. Virtual sandbox. In: Proceedings of the Pacific conference on computer graphics and applications, 2003. p. 252–9.
[6] M. Kass, G. Miller, Rapid stable fluid dynamics for computer graphics. In: Computer graphics (SIGGRAPH proceedings), vol. 24, 1990. p. 49–57.
[7] Oren M, Nayar S. Generalization of Lambert's reflectance model. In: Computer graphics proceedings, Annual conference series, 1994. p. 239–46.
[8] Jensen H, Legakis, J, Dorsey J. Rendering of wet materials. In: Proceedings of the eurographics workshop on rendering, 1999. p. 273–82.
[9] Li Z, Fung A, Tjuatja S, Gibbs D, Betty C, Irons J. A modeling study of backscattering from soil surfaces. IEEE Transactions on Geoscience and Remote Sensing 1996;34(1):264–71.
[10] Soulié R, Mérillou S, Terraz O, Ghazanfarpour D. Modeling and rendering of heterogeneous granular materials: granite application. Computer Graphics Forum 2007;26(1):66–79.
[11] Kimmel B, Baranoski G. A novel approach for simulating light interaction with particulate materials: application to the modeling of sand spectral properties. Optics Express 2007;15(15):9755–77.
[12] Preetham A, Shirley P, Smits B. A practical analytic model for daylight. In: Computer graphics (SIGGRAPH proceedings), 1999. p. 91–100.
[13] Soil Science Division Staff, Soil Survey Manual, Soil Conservation Service, United States Department of Agriculture Handbook 18; 1993.
[14] Gerrard J. Fundamentals of soils. New York, NY: Routledge; 2000.
[15] Brady N. The nature and properties of soils, 8th ed. New York, NY: Macmillan Publishing Co., Inc.; 1974.
[16] Hunt G, Salisbury J. Visible and near-infrared spectra of minerals and rocks: I. silicate minerals. Modern Geology 1970;1(4):283–300.
[17] Clark R, Swayze G, Gallagher A, King T, Calvin W. The US Geological Survey digital spectral library: Version 1: 0.2 to 3.0 microns, United States Geological Survey Open File Report 93-592, 〈http://speclab.cr.usgs.gov〉; 1993.
[18] Farrant P. Color in nature: a visual and scientific exploration. London, England: Blandford Press; 1999.
[19] Hunt G, Salisbury J, Lenhoff C. Visible and near-infrared spectra of minerals and rocks: III. Oxides and hydroxides. Modern Geology 1971;2(3):195–205.
[20] Baumgardner M, Silva L, Biehl L, Stoner E. Reflectance properties of soils. Advances in Agronomy 1985;38:1–43.
[21] Morley R, Boulos S, Johnson J, Edwards D, Shirley P, Ashikhmin M, Premože S. Image synthesis using adjoint photons. In: Proceedings of graphics interface, 2006. p. 179–86.
[22] Kimmel B, Baranoski G. A compact framework to efficiently represent the reflectance of sand samples. IEEE Transactions on Geoscience and Remote Sensing 2009;47(11):3625–9.
[23] Nicodemus F, Richmond J, Hsia J, Ginsberg I, Limperis T. Geometrical considerations and nomenclature for reflectance. National Bureau of Standards, United States Department of Commerce; 1977.
[24] Glassner A. Principles of digital image synthesis, vol. 2. San Francisco, CA: Morgan Kaufmann Publishers, Inc.; 1995.
[25] Karni Z, Gotsman C. Compression of soft-body animation sequences. Computers & Graphics 2004;28(1):25–34.
[26] Gribble C, Brownlee C, Parker S. Practical global illumination for interactive particle visualization. Computers & Graphics 2008;32(1):14–24.
[27] Krishnaswamy A, Baranoski G, Rokne J. Improving the reliability/cost ratio of goniophotometric comparisons. Journal of Graphic Tools 2004;9(3):1–20.
[28] J. Kajiya, The rendering equation. In: Computer graphics (SIGGRAPH proceedings), vol. 20, 1986. p. 143–50.
[29] Veach E. Robust monte carlo methods for light transport simulation. PhD thesis, Stanford University; December 1997.
[30] Stone M. A field guide to digital color. Natick, MA: A K Peters; 2003.
[31] Lafortune E, Foo S, Torrance K, Greenberg D. Non-linear approximation of reflectance functions. In: Computer graphics proceedings, Annual conference series, 1997. p. 117–26.
[32] Donner C, Lawrence J, Ramamoorthi R, Hachisuka T, Jensen H, Nayar S. An empirical BSSRDF model. ACM Transactions on Graphics 2009;28(3):30. (1–10).
[33] Hyvärinen A, Karhunen J, Oja E. Independent component analysis. New York, NY: John Wiley & Sons; 2001.