


Published in IET Image Processing  
 Received on 4th December 2007  
 Revised on 28th April 2008  
 doi: 10.1049/iet-ipr:20070207

In Special Issue on Visual Information Engineering



# SimBIL: appearance-based simulation of burst-illumination laser sequences

A. Nayak<sup>1</sup> E. Trucco<sup>1,2</sup> A. Ahmad<sup>1</sup> A.M. Wallace<sup>1</sup>

<sup>1</sup>ERP Joint Research Institute on Signal and Image Processing, School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK

<sup>2</sup>School of Computing, University of Dundee, Dundee DD1 4HN, UK  
 E-mail: e.trucco@dundee.ac.uk


**Abstract:** A novel appearance-based simulator of burst illumination laser sequences, SimBIL, is presented and the sequences it generates are compared with those of a physical model-based simulator that the authors have developed concurrently. SimBIL uses a database of 3D, geometric object models as faceted meshes, and attaches example-based representations of material appearances to each model surface. The representation is based on examples of intensity–time profiles for a set of orientations and materials. The dimensionality of the large set of profile examples (called a profile eigenspace) is reduced by principal component analysis. Depth and orientation of the model facets are used to simulate time gating, deciding which object parts are imaged for every frame in the sequence. Model orientation and material type are used to index the profile eigenspaces and assign an intensity–time profile to frame pixels. To assess comparatively the practical merit of SimBIL sequences, the authors compare range images reconstructed by a reference algorithm using sequences from SimBIL, from the physics-based simulator, and real BIL sequences.

## 1 Introduction and related work

Burst illumination laser (BIL) imaging is a technique combining active laser illumination with time gating (or range gating). It can image objects up to several kilometres away in limited or even no-visibility conditions, for example, at night. For these reasons, BIL imaging has become increasingly important in defence and security applications [1–3]. We sketch in Section 2 the BIL imaging principle and describe a physics-based model in Section 4. The reader is referred to [2, 4–7] for detailed descriptions of theory and applications of BIL. In essence, a pulsed laser source illuminates the scene with a sequence of flashes. The pulses are synchronised with the shutter of a camera tuned to the laser wavelength. The laser returns from the scene are recorded with an adjustable delay setting, corresponding to the round-trip time of flight at increasing ranges (time gating). By stepping the delay so that a range of distances through the target workspace is covered, the system generates a sequence of frames sweeping the scene. This can be visualised as a visibility plane sweeping the workspace (Fig. 1a).

The public-domain literature on BIL image processing is small, partly because of strategic sensitivity, partly because sensors and data are not easily accessible. The literature available concentrates on target recognition from range data [1, 3, 8]. A key problem is validation, as, typically, only limited sets of BIL images are available to researchers: a simulator generating realistic BIL sequences would therefore be extremely desirable. Ideally, such a simulator should be inexpensive, easily available, and, most important for validation purposes, target algorithms should perform similarly on simulated and real sequences.

Our work addresses these needs and brings two contributions. First, we introduce a prototype appearance-based simulator, SimBIL, generating realistic BIL sequences given a database of 3D mesh models and a set of examples of material-specific intensity–time BIL profiles. To our best knowledge, SimBIL is the first ever simulator achieving realistic BIL sequences using examples instead of complex physical models. Second, we compare results of range image reconstruction from time-gated sequences using SimBIL, a physics-based simulator we have developed concurrently, and



**Figure 1** BIL imaging principle

*a* Illustration of visibility frame sweeping the workspace

*b* Illustration of the time-gated frames:  $x$  and  $y$  are pixel co-ordinates,  $f$  is the frame (depth) index

*c* Three markers for which the intensity–time profiles (intensity plotted against progressive frame number, i.e. time) are shown. Notice the shift in profiles (due to range gating) and variation in intensity levels (due to anisotropic illumination)

real BIL data, suggesting the suitability of SimBIL sequences for testing purposes. The reference problem chosen for validation is the reconstruction of a single range image [9–11]. The reason for this choice is the importance of 3D reconstruction, giving shape information useful for visualisation and target classification.

Appearance and example-based representations have become ubiquitous since early work addressing recognition and tracking [12–14]. Recent applications include graphics and animation [15], human body modelling [16], medical image understanding [17] and photometric stereo [18]. At the vision-graphics borderline, example-based systems have been reported for reflectance modelling [19] and new paradigms based on image priors have emerged for restoration and inpainting [20] as well as view synthesis [21, 22]. To our best knowledge, however, appearance representations have never been used before in the simulation of laser sequences.

The rest of this paper is organised as follows. Section 2 briefly describes the principles of BIL imaging and the concept of an intensity–time profile. The architecture of SimBIL is sketched in Section 3 and the underlying principles that govern the physical, model-based simulator are described in Section 4. Section 5 reports a few results of our experiments. Finally, we summarise our work in Section 6.

## 2 SimBIL imaging

This section offers a concise introductory account of BIL imaging. The level of detail is sufficient to support the presentation of SimBIL; a mathematical model of image formation is discussed in Section 4, in the context of physical simulation.

BIL imaging consists of a laser source, illuminating the scene with a single pulse, and a camera, that is synchronised with the start of the pulse. If the pulse starts at, say, time  $t_0$  (seconds) and the camera is switched on at time  $t_R$  (seconds), where  $t_R = t_0 + \Delta t$ , the depth  $R$  (metres) corresponding to the return registered by the image is given by  $R = c(\Delta t/2)$ , where  $c$  is the speed of light ( $3 \times 10^8$  m/s). This procedure

eliminates pulse returns from unwanted depths and atmospheric backscatter. Further, flashing the laser light in rapid succession, that is, at regularly increasing  $\Delta t$ , results in a sequence of intensity frames (Fig. 1a), capturing intensities reflected from increasing depths. In practice, each frame actually includes returns not from a single depth, but from a small range of depths. Therefore the scene space actually imaged in each frame is a narrow fronto-parallel volume (cuboid), not a plane. The thickness of this cuboid depends on the time for which the camera is on (gate width).

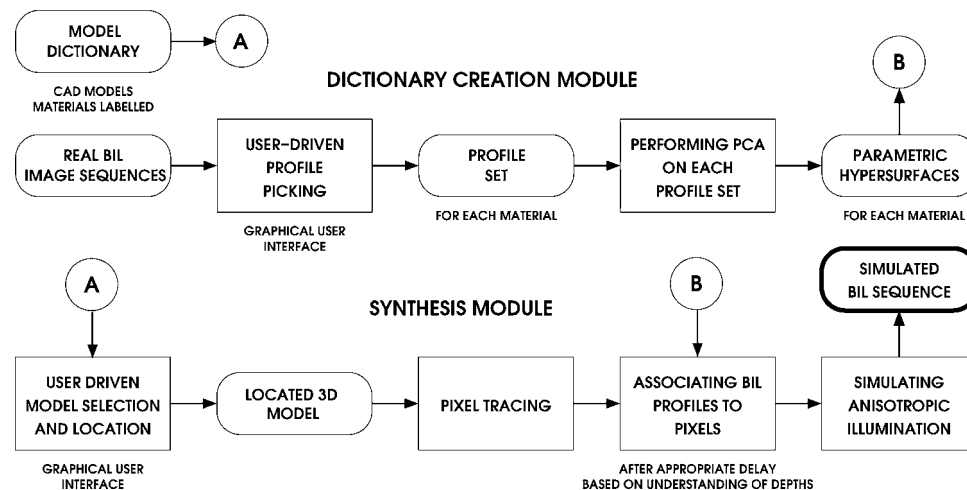
Fig. 1b illustrates the concept of a BIL sequence as a series of time-gated images, each one obtained for a given time-gating setting. The intensity–time profile (hereafter referred simply as profile) associated with a generic pixel, describing the evolution of the intensities at that pixel through the sequence, is a function of time gating. To clarify, Fig. 1c shows the profiles corresponding to three pixel locations (marked as 1, 2 and 3) in a frame from a real BIL sequence (suggested in the figure as a series of images). The plot shows the intensity values  $\hat{I}(x, y, f)$  for the  $f$ th frame in the sequence; the hat (^) indicates measurement, that is, the observed intensity at a given location. The profiles rise when the swept plane crosses a surface point reflecting the laser beam into that pixel. For notational simplicity, we shall drop  $f$  wherever immaterial to the discussion, generally denoting a single, complete profile as  $\hat{I}_s$ , where  $s = (x, y)$ . We notice that there is a depth-dependent delay in the rise of the profile, as the time-gating window sweeps the scene front to back, and a variation in intensity levels in the three profiles, which is due to anisotropic illumination (Section 3.2.4).

## 3 SimBIL architecture

The SimBIL architecture (Fig. 2) comprises two modules, described below: dictionary creation and synthesis.

### 3.1 Dictionary creation module

The dictionary creation module (Fig. 2, top) provides 2-fold knowledge to the system:



**Figure 2** SimBIL architecture

Rectangular boxes indicate processes, rounded boxes data or results

- a database of 3D CAD models, each of which is a triangular mesh with facets labelled by material type (Fig. 3a);
- parametric hypersurfaces of profile examples for each material computed from the dictionary of examples.


The former is a collection of wavefront OBJ models with modified material attributes for each facet. The latter is a contribution of this paper and is described in detail.

Examples of intensity–time profiles are selected manually from the pixels of real sequences using a GUI, and organised into a dictionary of examples by material and orientation (facet normal). Example profiles are collected for any given pair material orientation; for example, the current SimBIL prototype uses about 1155 examples for the material metal. This collection is the systems model of the radiometric response of the BIL sensor given a specific material, and this constitutes the major departure from conventional, physics-based modelling adopted in laser-imaging simulation. We discuss merits and limits of this approach in Section 6.

Extending the notation introduced in Section 2, we denote each time–profile sample as  $\hat{I}_{s_n, \gamma}^m$ , where  $n$  identifies the sample number. Recall that we collect many examples of intensities for each pair orientation material  $(\gamma, m)$ , that is, a small patch of a given material observed in a given orientation with respect to the sensor. The complete set of profiles,  $\mathbf{p}^m$ , obtained by imaging a patch of a given material in a number of orientations, is called the profile set. To reduce dimensionality, we transform this set into a profile eigenspace by principal component analysis (PCA). We denote  $\mathbf{p}^m$  as:

$$\mathbf{p}^m = [\hat{I}_{s_1, 1}^m, \hat{I}_{s_2, 1}^m, \dots, \hat{I}_{s_N, 1}^m, \dots, \hat{I}_{s_1, \Gamma}^m, \hat{I}_{s_2, \Gamma}^m, \dots, \hat{I}_{s_N, \Gamma}^m] \quad (1)$$

where  $\Gamma$  is the number of orientations and  $N$  the number of samples for each orientation, for a given material  $m$ . To build the profile eigenspace, the average of all profiles in



**Figure 3** Three-dimensional CAD models

- a Triangulated wavefront OBJ model of a Land Rover  
 b Examples of parametric hypersurfaces computed for metal door side  
 c Examples of parametric hypersurfaces computed for glass windscreen  
 For visualisation, only the three most important dimensions are shown

the profile set  $\mathbf{p}^m$ ,

$$\omega^m = \frac{1}{N\Gamma} \sum_{n=1}^N \sum_{\gamma=1}^{\Gamma} \hat{\mathbf{I}}_{s_n, \gamma}^m \quad (2)$$

is first subtracted from each profile to ensure that the largest eigenvector represents the most significant variation mode [14]. A new profile set,  $\mathbf{P}^m$ , is obtained after this process. We apply PCA to  $\mathbf{P}^m$ , obtaining a set of eigenvalues  $\{\lambda_i^m | i = 1, 2, \dots, \ell\}$  where  $\{\lambda_1^m \geq \lambda_2^m \geq \dots \geq \lambda_\ell^m\}$ , and a corresponding set of eigenvectors  $\{\mathbf{e}_i^m | i = 1, 2, \dots, \ell\}$ . We have found empirically that  $4 \leq \ell \leq 7$  eigenvectors describe a profile with sufficient detail, in two practical senses: first, the synthetic sequences generated are realistic; second, the sequences generated allowed successful completion of the target task (reconstruction).

The  $\ell$  eigenvectors define the profile eigenspace, that is, the example-based representation of BIL intensities when imaging known materials at known orientations. The eigenspace parameters are patch orientations and sample index. Each profile sample is projected into the profile eigenspace by subtracting the average profile,  $\omega^m$ , and then finding the dot product of the result with each of the eigenvectors

$$\mathbf{H}_{s_N, \Gamma}^m = [\mathbf{e}_1^m, \mathbf{e}_2^m, \dots, \mathbf{e}_\ell^m]^T [\hat{\mathbf{I}}_{s_n, \gamma}^m - \omega^m] \quad (3)$$

Finally, a parametric hypersurface,  $\mathbf{H}^m(\alpha^m, \beta^m)$ , is obtained by fitting cubic splines to the discrete points,  $\mathbf{H}_{s_N, \Gamma}^m$ , obtained by projecting all profile samples into the profile eigenspace. Here,  $\alpha^m$  and  $\beta^m$  are continuously varying parameters for sample number and orientation, respectively, used to obtain the profile set. All other symbols are as described.

To illustrate, two parametric surfaces, for the metal door and glass window of a Land Rover, are shown in Fig. 3. For display, only the first three most significant dimensions of each profile eigenspace are shown. The dictionary-creation module simply

stores the parametric hypersurface ( $\mathbf{H}^m(\alpha^m, \beta^m)$ ), the eigenvectors ( $\{\mathbf{e}_i^m | i = 1, 2, \dots, \ell\}$ ) and the average ( $\omega^m$ ) for each material  $m$ .

## 3.2 Synthesis module


The synthesis module generates the complete BIL sequence of a given model in a target orientation (Fig. 2, bottom). The synthesis module has four main blocks, presented below.

### 3.2.1 User-driven model selection and placement:

First, the modified CAD model corresponding to the target vehicle is selected from a model dictionary. We have developed an interactive model selection tool; a GUI allows the user to rotate, translate and scale the selected CAD model to bring it into the 3D pose necessary to achieve the desired viewpoint. The posed 3D object is then used by a pixel tracing process, sketched next, to identify the image pixels laying on thin 3D slices (the thin viewing volumes corresponding to successive time-gated frames) at constant, fixed depth.

### 3.2.2 Pixel tracing:

Pixel tracing aims to determine the model point, including its associated material and orientation, with a pixel in each sequence frame. The user selects a resolution for the frames in the simulated sequence, that is, the number of pixels along the two image dimensions. The selected model pose provides the position and orientation of the model relative to the camera, that is, the synthetic viewpoint. In pixel tracing, a ray is traced from an image pixel to intersect the surface of the 3D model. Pixel tracing works basically as an inverse orthographic projection from the image plane to the CAD model. This is adequate given the very large stand-off distance relative to the target size; absolute size can be ignored in simulations. The inverse projection associates each image pixel with a CAD model point (and facet), and depth, facet orientation and material label are recorded at each pixel (Fig. 4a). An intermediate representation, essentially a depth map from the desired viewpoint



**Figure 4** Pixel-tracing a Land Rover model

For visualisation, the image plane resolution is selected as  $16 \times 9$  pixels  
Pixel intensities encode depth (the brighter the closer)

augmented with material information, is now built. As an example, the depth data obtained after pixel tracing the Land Rover model of Fig. 3a, for a given viewpoint are shown in Fig. 4b. Data are shown as an intensity image (brighter pixels are closer to the camera).

### 3.2.3 Associating BIL profiles to pixels and depth:

After associating image pixels with points on the 3D model, the next step is to associate a time profile to each pixel. This involves selecting a suitable profile (or combination thereof) from the profile eigenspaces, and shifting it in time to illuminate the corresponding scene point,  $p$ , at the correct time, as the time-gated imaging plane sweeps the scene. In other words, the profile must be shifted so that its peak (highest return) occurs when the sweeping plane intersects the object point imaged by  $p$ .

To understand how this is done, recall that pixel tracing determines the depth of the scene point imaged by each frame pixel, so that a complete depth map can be built for the object in the user-defined pose. This depth map is now used to define the time shift of the BIL intensity profile associated with each pixel (Fig. 1c). As described in Section 3.1, a precomputed parametric hypersurface,  $H^m(\alpha^m, \beta^m)$ , is available for each material  $m$ , modelling the BIL image intensities generated by patches of that material in every possible orientation,  $\beta^m$ .

Once material and orientation are fixed at each pixel from the augmented depth map, a 1D hypercurve is defined on the hypersurface. This hypercurve spans all the samples collected for the pair  $m, o$  (and interpolated points on the hypersurface).

To generate the intensity of the actual pixel in all the sequence frames, we select a sample profile at random along the hypercurve, so that the sample number,  $\alpha^m$ , satisfies  $1 \leq \alpha^m \leq N$ . As described earlier,  $N$  is the number of profile samples for each orientation for a material  $m$ .

### 3.2.4 Simulating anisotropic illumination:

For short-range images, such as the ones in our real data sets, the laser beam may not diverge enough to illuminate the whole scene uniformly. This results in a typical intensity pattern fading away from the centre of the laser beam. A very simple approximation of this effect can be achieved by superimposing to the image an illumination cone with circular normal cross-section centred in the laser source. We therefore modulate intensities in each frame of the sequence with a 2D Gaussian to simulate this illumination cone. With longer-range sequences, the intensity pattern would be less noticeable, as the standard deviation of the Gaussian approximating the profile becomes larger. More sophisticated approaches can, of course, be adopted, inspired by the solutions adopted to compensate for non-uniform illumination patterns in various application domains, for example, underwater [23] or retinal image processing [24].

## 4 Physical simulation of BIL imagery

To simulate the pulsed coherent BIL imaging system, illustrated schematically in Fig. 5, we use models of laser propagation similar to those described in [5, 7, 25]. In general, there are two principal degrading effects that we can include in physical simulation of BIL imagery: speckle induced by the interaction of the rough surface with the laser source, and the effect of turbulence introduced by the atmospheric propagation path, usually modelled by a refractive index structure constant, conventionally indicated as  $C_n^2$ . Turbulence-induced image degradation can be characterised by the length of exposure; at short exposures small atmospheric eddies cause blurring, and at longer exposures larger eddies cause ‘dancing’ in the image plane. The main cause is the dynamic variation of refractive index caused by the temperature variation in the atmosphere. Physical modelling of turbulence is beyond the scope of this paper and is indeed an open problem in general because of the considerable complexity of modelling the eddy structure at different latitudes, at different times of day, in different weather conditions and so on. The appearance method described in Section 3 provides an antidote to this modelling, in that turbulence effects can be captured in the training imagery. In our physical simulation, we consider only the generation of speckle, neglecting turbulence, which is still applicable when the propagation path is short, for example, a few hundred metres, or the refractive index structure constant is low, for example, on a clear, early morning at a northern latitude [6]. Thus, we can use our physical simulation as a reference for the appearance-based method. We give a verbal account of the simulation here, and refer the reader interested in the mathematical treatment to [25].

The laser output pulse has a Gaussian spatial profile in  $(x, y)$  at the exit pupil. The beam diverges spatially as it propagates through the atmosphere. At the target, the laser radiation is diffracted by a surface that is modelled by first- and second-order statistics. The reflected radiation is then propagated from the target to the receiver lens/aperture and image plane. Since we are observing the target far from the aperture, we





Figure 5 Simplified system used for the physics model-based approach

can model the received diffraction pattern using the Fraunhofer model [25]. This gives an expression for the field distribution immediately before the aperture, which approximates the diffraction pattern in the observation plane  $(x, y)$  by the Fourier transform of the field at spatial frequency  $x/\lambda z$ , where  $\lambda$  is the laser wavelength and  $z$  the stand-off distance from sensor to target. After transmission through a circular aperture, and considering the phase changes introduced by the imaging lens, the image of the target in the image plane is described as the superposition of the entire wave field produced by all the target points. In effect, the convex lens Fourier transforms the incident field to give an expression for the image of the target that has the form of a Fraunhofer diffraction pattern in the image plane. If the object is regarded as a collection of points that scatter the waves in all directions, and if we know the response generated by a unit-amplitude point source applied at the object coordinates in the image plane, then the field produced by the target is the convolution of the point-spread function of the unit-amplitude point with the geometric approximation of the object field in the image plane. The square of the resulting field describes the intensity distribution in the image plane.

On the same physical basis, the coherent imaging system is defined as the convolution of the normalised coherent transfer function, with the geometric approximated incoherent object field in the image plane. We also include the spatial frequency filtering of the CCD array, described by a sinc function of the ratio  $f_s/f_{s0}$ , where  $f_s$  is the spatial frequency and  $f_{s0}$  is defined by the pixel spacing. Once again we refer the reader to [25] for the mathematical details.

As the laser is pulsed, the source signal has a finite temporal width, as shown in the profiles illustrated in Fig. 1c. Typically, this temporal width is of the order of 10–33 ns, corresponding to a depth of field of 3–10 m. We consider typical laser sources with a short coherence length of the order of 1 mm; beyond this distance, which is much shorter than the depth of field of the imaging system, the observed pattern de-correlates rapidly. Therefore the simulated image is acquired over a defined frame integration time, in which the relative intensities follow the temporal profiles of Fig. 1c. The latter can be defined parametrically by a Gaussian function of time, fitted to these profiles, or by a look-up table formed from a weighted average of a subset of these profiles. In practice, the received image is created by multi-look averaging; different speckled images are generated and combined using the propagation model described above and target surfaces with defined, statistical roughness parameters.

We can now write the procedure for physical simulation of a range-gated BIL image sequence. An example of synthetic frames is shown in Fig. 6. We use two image files as source data: an intensity image of the target scene and the range information at each pixel. The simulation then requires the physical model parameters described in the preceding paragraphs.



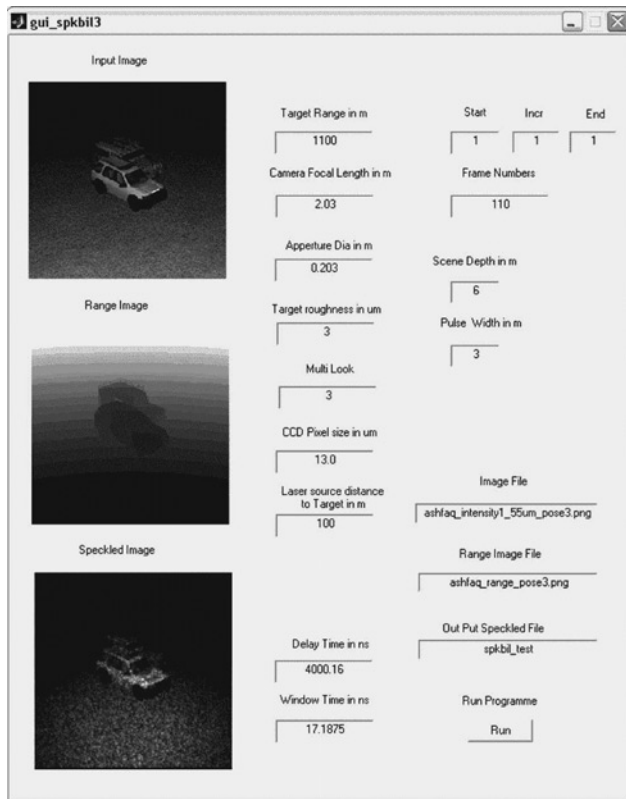
**Figure 6** Three frames from a sequence generated by the physics-based simulator: saloon car in front of Land Rover (see text for parameter values)

1. Model coherent radiation source, for example, as a Gaussian beam field at the target after propagation through a known distance from the source, as defined by stand-off distance and range information.
2. Define target reflectance map using intensity image file.
3. Generate the Gaussian-correlated rough surface, using statistical parameters to define roughness and an exact surface roughness profile if available.
4. Introduce the phase shifts in the reflected wave field from the rough surface.
5. Using the point spread function of the imaging system, evaluate the field distribution. The square of this field will provide the intensity in the image plane.
6. Filter this intensity field by the MTF to obtain the digital image.
7. To simulate a multi-look intensity image, integrate over several samples each of which is defined by the temporal profile, using independent random rough surfaces.

Fig. 7 shows the GUI developed for the physics-based simulator. Fig. 6 shows three frames from an example of the physical simulation process, in this case of a saloon car in front of a Land Rover. The range to the target is 1100 m, the approximate scene depth is 16 m, and the focal length and aperture of the imaging system are 2.03 and 0.203 m, respectively. The image plane pixel separation is 26  $\mu\text{m}$  and the target roughness parameter is 3  $\mu\text{m}$ . The image is integrated over seven looks, each of which has a pulse width of 3 m.

## 5 Experimental results

This section presents the results achieved with SimBIL and the physical BIL simulator described above. The overall purpose of the experiments was to assess the quality of SimBIL sequences in two senses. First, the sequences must look realistic, that is, convincingly similar to those generated by a real sensor. Second, the results of algorithms run on SimBIL sequences must be similar to those achieved with sequences generated by a physics-based simulator or real BIL sensors.



**Figure 7** Screenshot of the GUI developed for the physics-based simulator

Real BIL imagery is difficult to obtain and, to our best knowledge, no public data sets exist. The sequences we used were provided by BAE Systems Avionics (now SELEX S&AS). The targets were vehicles and flat boards placed about 300 m from the imaging and laser source platform. It was possible to image targets in a small number of orientations only, as the sensor could not be moved and the vehicles had obviously to rest on the road. We collected about 2630 profiles representing seven different materials from three vehicles, a relatively small number of examples but the largest feasible data set given the constraints (availability of sensor from our industrial partner, absence of public-domain BIL sequences), and sufficient for proof-of-concept testing. However, the much smaller number of samples compared with the number of pixels to be generated contributes to the low-speckle appearance of current SimBIL sequences. Profiles for each material on average were available in 15 orientations. The same material appears in multiple orientations for a single vehicle pose, as different parts of the frame are oriented differently.

We show three representative frames from each sequence of a Land Rover in Fig. 8. The top row shows frames from a SimBIL sequence, the middle row a sequence from the physics-based simulator (with an added camouflage pattern) and the bottom row a real BIL sensor sequence. The Land Rover model used for SimBIL was only a low-resolution facet model of the type of Land Rover in the real-sensor




**Figure 8** Representative frames from BIL sequences of a Land Rover

Top row: SimBIL sequence; middle row: sequence from a physical model-based simulator; bottom row: real BIL sequence

Notice the Land Rover model used for generating simulating sequences differed from the real Land Rover used during the real BIL data acquisition

sequence. Saturation was avoided in the simulated BIL sequences by excluding pixels with saturated profiles at the dictionary creation stage. With a real BIL sensor, saturation can be avoided by adjusting the gain.

Fig. 9 shows two frames from a real sequences of painted flat boards approximately front-parallel to the sensor. The saturation present in the real sequences was deliberately




**Figure 9** Representative frames from BIL sequences of fronto-parallel flat boards

Top row: real sequence; Bottom row: SimBIL sequence

avoided in the simulated ones, but would be easy to simulate, if needed, by a simple alteration of the dynamic range. Notice the simulated anisotropic illumination (Section 3.2.4).

Qualitatively (by visual inspection), the SimBIL sequences resemble the real ones, but the speckle is less pronounced than in real data. In the real sequence, the turbulence effects were not very significant as the range was short, but there was some likely beam wander in addition to the uneven illumination. For the physics-based simulator, we were able to obtain first- and second-order image statistics very similar to the real data by selecting optical system parameters appropriate to the real sensor and using plausible roughness data for the vehicle surface, as the real roughness could not be measured.

To further validate comparatively SimBIL and the physics-based simulator, we processed the data with a depth reconstruction algorithm we have used previously for real BIL data [11]. This was based on reversible jump Markov chain Monte Carlo (RJMCMC) techniques. We were able to explore a parameter space of variable dimension, representing the number of returns that may be observed in the field of view of a single pixel in the image. Multiple returns are possible because this may span more than one surface at longer range, or the laser may transmit and reflect at a semi-transparent surface, for example, the window of the Land Rover. For a given dimension, MCMC is a powerful simulation algorithm that allows us to find a stationary distribution which is the posterior distribution of the parameters given by the data (target distribution). These parameters are the amplitude and range (time of flight) to the given surface or surfaces. In general, the RJMCMC technique has been shown to be much more sensitive and accurate to weak returns from surfaces that have poor reflectance, but in this paper, we use the method in a comparatively constrained situation, that is, we display only the strongest return at each pixel location. The resulting range images obtained from three SimBIL sequences are shown in Fig. 10; at this stage, we can only present this as a visual comparison. The surfaces of the Land Rover are all recovered and distinguished easily, and we have already used simulated depth and infrared imagery of this nature to investigate the performance of our depth reconstruction and vehicle identification algorithms [11].



**Figure 10** Reconstructed range images

- a Range images obtained from SimBIL
- b Range images obtained from physics model-based simulator
- c Range images obtained from real BIL sequences

## 6 Discussion and future work

We have presented a novel technique to simulate BIL sequences, based on appearance examples, and a prototype implementation, SimBIL. SimBIL departs from the conventional approach for simulating sequences in the non-visible spectrum, typically using complex physical models, in that it is based on examples of time-intensity profiles. Simplified models, capturing a wide range of shapes, have proven sufficient to generate sufficiently realistic sequences, although refinements may be necessary in specific cases (e.g. classification based on details). Regions of different materials are identified offline in a database of representative models.

In addition, we have reported a physics-based simulator of BIL sequences that is easy to use, requires only a depth image and an intensity image that can be generated by a standard graphics package, or a more complex atmospheric modelling package such as CameoSim and gives results that are very comparable to real BIL data at ranges of the order of 100–500 m, depending on the atmospheric conditions. We have used such data in our evaluation of depth reconstruction algorithms, where we can thus provide ground truth for the surface distribution.

To further summarise the results of our experiments, we can say that SimBIL sequences compare well with real ones, as well as with sequences generated from our physics-based BIL simulator, for a range of vehicles. Results from range reconstruction suggest that SimBIL sequences can support the validation of computer vision algorithms, but more development and analysis are required to determine suitability for specific tasks. Our choice of range reconstruction as a target algorithm was dictated by the applicative importance of reconstructing the range from BIL sequences; in addition, successful reconstruction of range data makes it possible to deploy a large corpus of well-established algorithms for range image analysis.

The advantages of using appearance examples as model of radiometry of surfaces are mainly two-fold: avoiding the sometimes exceeding intricacies of physical modelling of some phenomena (e.g. turbulence), and guaranteeing realistic results as the generated sequences are based on real video material. These are the typical promises of image-based rendering, the inspiring paradigm behind SimBIL [26]. The limit of this approach is typical of learning approaches: a sufficient volume of training data must be available. Also, some characterisation of the spatial dependencies between pixel profiles, in the spirit of recent work on image synthesis [22], could benefit the visual quality of the synthesised textures.

Further work can further improve the realism of SimBIL sequences, making them more similar to real-sensor sequences. However the present prototype, in our opinion, demonstrates successfully that realistic sequences can be achieved via an example-based system, based on real-sensor data and therefore holding the promise of high realism.



Future work should also address the selection of profiles from the profile eigenspaces (currently a simple random selection), and accessing larger volumes of examples to improve generality.

## 7 Acknowledgments

We thank Sergio Hernández-Marín for reconstructing the range data from our simulated sequences. A. Nayak was sponsored by BAE Systems and SELEX S&AS (formerly BAE Systems Avionics) grant within the BAE Strategic Alliance programme.

## 8 References

- [1] HUTCHINSON B.A., GALBRAITH R.L., STANN B.L., DER S.Z.: 'Simulation-based analysis of range and cross-range resolution requirements for identification of vehicles in ladar imagery', *Opt. Eng.*, 2003, **42**, (9), pp. 2734–2745
- [2] LETALICK D., AHLBERG J., ANDERSSON P., CHEVALIER T., GRONWALL C., LARSSON H., ET AL.: '3-D imaging by laser radar and applications in preventing and combating crime and terrorism'. Proc. RTO SCI Symp. Systems, Concepts and Integration (SCI) Methods and Technologies for Defence Against Terrorism, London, October 2004
- [3] ZHENG Q., DER S.Z., MAHMOUD H.I.: 'Model-based target recognition (ATR) in pulsed ladar imagery', *IEEE Trans. Image Process.*, 2001, **10**, (4), pp. 565–572
- [4] CHUROUX P., BESSON C.: 'Burst illumination laser imaging system for long range observation', *Proc. SPIE*, 1997, **3707**, pp. 582–589
- [5] CHUROUX P., BESSON C., BOUZINAC J.P.: 'Model of a burst imaging lidar through the atmosphere', *Proc. SPIE*, 2000, **4035**, pp. 324–331
- [6] DUNCAN S., COPLEY J., HARVEY G., HUMPHREYS D., GONGLEWSKI J., BAKER I.: 'Advances in laser gated imaging in an airborne environment', *Proc. SPIE*, 2006, **6206**
- [7] EDOUARD D., CHUROUX P., FLAMANT P.H.: 'Burst Illumination Imaging Lidar: intensity correlation function in the image plane', *Proc. SPIE*, 2002, **4723**, pp. 189–197
- [8] VERLY J.G., DELANOY R.L.: 'Model-based automatic target recognition (ATR) system for forward-looking ground-based and airborne imaging laser radars (LADAR)', *Proc. IEEE*, 1996, **84**, (2), pp. 126–162
- [9] ANDERSSON P.: 'Long-range three-dimensional imaging using range-gated laser radar images', *Opt. Eng.*, 2006, **45**, (3), pp. 034301-1–034301-10
- [10] BUSCK J.: 'Underwater 3D optical imaging with a gated viewing laser radar', *Opt. Eng.*, 2005, **44**, (11), pp. 116001-1–116001-7
- [11] HERNANDEZ-MARIN S., WALLACE A., GIBSON G.: 'Bayesian analysis of lidar signals with multiple returns', *IEEE Trans. Pattern Anal. Mach. Intell.*, 2007, **29**, (12), pp. 2170–2180
- [12] BLACK M.J., JEPSON A.D.: 'EigenTracking: robust matching and tracking of articulated objects using a view-based representation'. Proc. European Conf. Computer Vision, 1996
- [13] COOTES T.F., COOPER D., TAYLOR C.J., GRAHAM J.: 'Active shape models: their training and application', *Comput. Vis. Image Underst.*, 1995, **61**, (1), pp. 38–59
- [14] MURASE H., NAYAR S.K.: 'Visual learning and recognition of 3D objects from appearance', *Int. J. Comput. Vis.*, 1995, **14**, pp. 5–24
- [15] BLANZ V., POGGIO T., VETTER T.: 'Reanimating faces in images and videos', *Eurographics*, 2003, **22**, (3), pp. 641–650
- [16] ALLEN B., CURLESS B., POPOVIC Z., HERTZMANN A.: 'Learning a correlated model of identity and pose-dependent body shape variation for real-time synthesis'. Proc. ACM SIGGRAPH/Eurographics Symp. on Computer Animation, 2006, pp. 147–156
- [17] ROBERTS M.G., COOTES T.F., ADAMS J.E.: 'Vertebral morphometry: semi-automatic determination of detailed vertebral shape from DXA images using active appearance models', *Investigative Radiol.*, 2007, **41**, (12), pp. 849–859
- [18] HERTZMANN A., SEITZ S.M.: 'Shape and materials by example: a photometric stereo approach'. Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition, June 2003, vol. 1, pp. 533–540
- [19] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: 'A data-driven reflectance model'. Proc. SIGGRAPH, 2003, pp. 759–769
- [20] ROTH S., BLACK M.: 'Fields of experts: a framework for learning image priors'. Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition, 2005, vol. 2, pp. 860–867
- [21] NAVARATNAM R., FITZGIBBON A.W., CIPOLLA R.: 'Semi-supervised learning of joint density models for human pose estimation'. Proc. XVII British Machine Vision Conf., Edinburgh, 2006
- [22] WOODFORD O.J., REID I., TORR P., FITZGIBBON A.W.: 'Fields of experts for image-based rendering'. Proc. XVII British Machine Vision Conf., Edinburgh, 2006
- [23] KHAMENE A., NEGAHDARIPOUR S.: 'Motion and structure from multiple cues: image motion, shading flow, and

stereo disparity', *Comput. Vis. Image Underst.*, 2003, **90**, (1), pp. 99–127

[24] GRISAN E., GIANI A., CESARACCIU E., RUGGERI A.: 'Model-based illumination correction in retinal images'. Proc. IEEE Int. Conf. Engineering in Medicine and Biology (EMBC06), New York, August–September 2006

[25] GOODMAN J.: 'Speckle phenomena in optics: theory and applications' (Roberts & Co, 2006)

[26] SZELISKI R.: 'Stereo algorithms and representations for image-based rendering'. Proc. British Machine Vision Conf., September 2005