

UNCLASSIFIED



Australian Government
Department of Defence
Defence Science and
Technology Organisation

Further Studies into Synthetic Image Generation using CameoSim

A. Brady and S. Kharabash

Intelligence, Surveillance and Reconnaissance Division
Defence Science and Technology Organisation

DSTO-TR-2589

ABSTRACT

Previous work by DSTO described the synthetic image generation process using CameoSim and RadThermIR; this report describes the progress made on the modelling of atmospheres, terrains and sensors. A major source of error in the synthetic image generation process previously identified has been resolved. The error concerns the use of interpolation in CameoSim when it uses MODTRAN to incorporate atmospheric effects in synthetic images. A software tool has been created which iteratively reduces the interpolation errors until they are insignificant. Validation studies are currently being planned in the visible to shortwave infrared. In preparation of the validation effort a study of BRDF models has been completed, which includes the physical plausibility of models, how measured data is fitted to models and how well CameoSim samples each model.

RELEASE LIMITATION

Approved for public release

UNCLASSIFIED

UNCLASSIFIED

Published by

*Intelligence, Surveillance and Reconnaissance Division
DSTO Defence Science and Technology Organisation
PO Box 1500
Edinburgh South Australia 5111 Australia*

*Telephone: (08) 7389 5555
Fax: (08) 7389 6567*

*© Commonwealth of Australia 2011
AR-015-073
August 2011*

APPROVED FOR PUBLIC RELEASE

UNCLASSIFIED

UNCLASSIFIED

Further Studies into Synthetic Image Generation using CameoSim

Executive Summary

DSTO has acquired CameoSim for the purpose of generating realistic synthetic images over visible and infrared wavebands; it is a complicated program with large number of settings, so much so that the effort in understanding how the models work is just as significant as the building of scenes. As we strive for the best results that CameoSim is capable of producing, we have improved upon our processes and developed new tools to facilitate better use of CameoSim.

High fidelity results depend upon two things, having the correct data and using it well. Whilst the importance of the former is no surprise, we have found great improvements can still be made in the latter. There are many inputs to CameoSim that rely upon the user's judgement to determine that the quality is 'high enough'. Whilst the judgment of an expert is irreplaceable, this can be assisted by having software that produces objective measures of error and methods of optimising the settings to achieve a specified level of quality.

Validation studies are currently being planned to test the performance of CameoSim in the visible to shortwave infrared. In preparation of this, a study of the model inputs has been completed, with the aim of providing direction for collecting field data at the appropriate precision and incorporating it effectively into the model.

The ability to generate synthetic imagery is complementary to the collection of imagery with real sensors. The advantages of synthetic imagery are that the elements within a scene can be controlled and set to be whatever is desired. Whilst there are limitless scenes that can be generated, the most frequently beneficial will be scenes designed to answer 'what if' questions, by constructing a scene and then changing a single feature, e.g. the range of the target, the resolution of the sensor, etc. The modelling capability outlined in this report enables DSTO to provide advice on the performance of a range of Electro-Optic (EO) sensors from visible through long wave IR cameras, including hyper spectral systems. This has the potential to impact on the support DSTO provides across all the services on the performance of such sensors, in maritime, land, airborne and even spaced-based environments.

UNCLASSIFIED

UNCLASSIFIED

UNCLASSIFIED

UNCLASSIFIED

Authors

Andrew Brady

Intelligence, Surveillance and Reconnaissance
Division

Andrew Brady received a B.C.S.T. (Hons.) degree in 2002 and a B.E. (Hons.) in 2006 both from the University of Sydney. Andrew joined DSTO in 2006 and has been involved in work in synthetic image generation, thermal modelling and sensor modelling for infrared sensor studies.

Sergey Kharabash

Intelligence, Surveillance and Reconnaissance
Division

Sergey Kharabash received a B.S. (Hons) degree in 2005 from Flinders University. Sergey joined DSTO in 2007 and has been involved in work in synthetic image generation; BRDF measurement and fitting and scene modelling.

UNCLASSIFIED

UNCLASSIFIED

UNCLASSIFIED

Contents

1. INTRODUCTION.....	1
2. MODELLING THE TERRAIN	2
2.1 Terrain Geometry.....	2
2.2 Vertex Normal Interpolation	3
2.3 Flora Geometry	4
2.4 Planting.....	7
2.5 Materials	7
2.6 Procedural Textures.....	9
2.7 Material Classification.....	9
2.8 Layering Materials.....	10
2.9 Thermal Sensitivity	10
3. MODELLING SPECTRAL ATMOSPHERES	12
3.1 Processing Time Issues	13
3.2 Range Normalisation	14
3.3 Interpolation	15
3.4 Geometry Optimisation.....	15
3.5 Example.....	17
4. MODELLING BRDFS	19
4.1 Theory	19
4.2 CameoSim BRDFs.....	20
4.2.1 Ashikhmin Shirley (isotropic).....	20
4.2.2 Cook Torrance (uncoated/coated).....	21
4.2.3 Lorentzian Peaks Model	21
4.2.4 Sandford Robertson	21
4.3 Process of Verifying CameoSim Importance Sampling.....	22
4.4 Verification Results.....	23
5. RENDERING SETTINGS	25
5.1 The Settings to be Optimised	25
6. POST-PROCESSING IMAGERY.....	27
6.1 Accounting For Atmospheric Scattering and Turbulence.....	27
6.2 Sensor Effects.....	28
6.3 Image Processing Software Choices.....	28
6.4 Future Work	29
7. CONCLUSIONS.....	30
7.1 Ongoing Issues	30

8. ACKNOWLEDGMENTS.....	32
9. REFERENCES	33
APPENDIX A: SAMPLE IMAGERY	35
APPENDIX B: ATMOSPHERE GENERATION SPECIFICS	38

1. Introduction

CameoSim is widely used at DSTO. Weapon Systems Division (WSD) originally acquired the model in 2004 and has used it extensively in applications to assess EO seeker performance. Maritime Platforms Division (MPD) and Electronic Warfare and Radar Division (EWRD) have also used CameoSim in various studies into ship signatures and the performance of electro-optic (EO) threat warning type systems. Intelligence Surveillance and Reconnaissance Division (ISRD) acquired CameoSim in late 2004 and has used it extensively to generate synthetic imagery as would be observed through a variety of EO and infrared (IR) sensors. This work has included the generation of synthetic imagery in support of studies into the Thermal Imager on the Australian Army Armed Reconnaissance Helicopter (ARH) [Carr and Brady, 2007]. More recently ISRD has used CameoSim to generate hyperspectral imagery in visible (VIS) through short wave infrared (SWIR) bands to model the Woodside Army Barracks.

The report referred to above consists of a description of all the software tools acquired, a step-by-step account of the construction of the models, and identification of areas where model improvements were necessary. Those model improvements can roughly be divided into two cases: not having sufficient data to model a scene element correctly, and not being able to obtain high enough simulation fidelity with the software available.

This report addresses some of the latter issues regarding simulation fidelity and records the changes that have been made to the modelling process. The aims are to use CameoSim to obtain a high level of simulation fidelity given the data available and to identify the practical or fundamental limitations in the process.

As a case study, the task of modelling a varied 8×8 km region of rural Northern Territory (NT) was undertaken. The use of CameoSim (versions 5.8 to 5.12) to generate high fidelity VIS, medium wave infrared (MWIR) and long wave infrared (LWIR) synthetic imagery with attempts made to generate imagery for both the wet and the dry seasons. The scenes include active and inactive BRDM2s, 2S6s (Tunguska) and Australian Army Land Rovers, which were constructed for a previous report [Carr and Brady, 2007]. The atmosphere was modelled to allow the observer to be positioned anywhere between 0.02-25 km from the targets and 0.001-8km above the ground.

2. Modelling the Terrain

The main requirement for selecting a suitable terrain region was to provide a modelling challenge in terms of geological, topographical and floral contrasts. The selected region is from the Northern Territory and features a 450 m escarpment, a valley, varying numbers and sizes of trees and a dry river bed. The region is centred at S12°54'15" E132°55'45" and is 4'30"x4'30" or 8.333x8.123km North-to-South and East-to-West, respectively. The size of the region allows high polygonal detail for close up views without requiring excessive number of polygons for the overall region. Earth curvature was ignored, since for an 8km range the drop due to earth curvature is only 5m. However any imagery of the horizon is misleading because the terrain terminates where it should continue on.

2.1 Terrain Geometry

Digital terrain elevation data of level 2 type (DTED2) was imported and meshed in MultiGen Creator v3.1. The CAT meshing algorithm was used to convert 72,900 height posts to a mesh of 65,000 triangles. The geometry was saved using Creator's native open flight, *.FLT, file format. Creator was used because it offers a convenient interactive environment for constructing and manipulating 3D geometry and the CAT algorithm because it uses more polygons in angulated regions and less on the flats. Creator's CAT algorithm creates several levels of detail for real time rendering, these were removed together with unnecessary hierarchy. Some smoothing was applied along escarpment edges where a lack of polygons resulted in a jagged appearance.

An overhead RGB 2721x2652 pixel image of the same size as the terrain was captured from Google Earth. The image was mapped to the terrain in Creator using its corner points. Image mapping conveniently imports to a CameoSim CGF file through the **cs-flt2cgf** converter. This image was later used to classify the terrain with materials based on colour.

We observed a difficulty in creating a high fidelity central region in terms of higher polygon detail and higher image resolution, due to the manual effort required to stitch the two regions together and due to a notable discontinuity in materials after material classification. Thus this effort was abandoned.



Figure 1: Screen capture from Creator of the terrain geometry and the Google Earth image applied

2.2 Vertex Normal Interpolation

Creator's FLT file format and CameoSim's CGF format both allow each vertex to have its own vertex normal where each may or may not point in the same direction as the vertices for adjacent polygons. These are used to render smooth surfaces where the effective surface normal at any point on a polygon is an interpolation between its vertex normals. Creator was used to co-align all vertex normals for the terrain, with *Triangle Area Weighting* feature enabled, to give the terrain a smooth appearance.




Figure 2: Example of a sphere made of 72 faces, rendered with no vertex alignment, with vertex alignment and a rendering of a true curved surface. The schematic shows 2 adjacent polygons with vertex normals (gray arrows) and the effective surface normals (black arrows).

Conversely vertex normal interpolation in general is a way of fooling the ray tracer into thinking the surface faces a direction that it does not. This can create some physical problems. One problem is where the reflected ray is fired internally into the surface because the surface normal dictates that such a ray should be possible and a hypothetical problem follows of what one should do with such a ray. Arguably however the benefits of using vertex interpolation outweigh the costs.




Figure 3: A curved surface (left) in cross section approximated by flat polygons (right) with co-aligned vertices and the problems of internally fired reflected rays that may result.

CameoSim's **flt2cgf** converter is programmed to limit the import of vertex normals that deviate by more than 45° from the surface normal. If a normal deviates by more than this then all the normals for that polygon are replaced with the surface normal, which reveals the polygon edges. Some manual labour was required to ensure that none of the normals in the FLT model deviated by more than 45° . It would have been preferred if CameoSim substituted such normals with the nearest acceptable 45° normal to partially hide the edges of these polygons.

2.3 Flora Geometry

Only a few eucalyptus tree models were created using OnyxTreeBroadleaf 6.0 to populate the entire region. OnyxTree allows good control over the structure of the trees and the number of polygons used to describe the structure.

Table 1: Basic description of the three tree models used to populate the region.


Tree	Width (m)	Height (m)	Leaves	Leaf resolution (polygons)	Trunk resolution (polygons)	Polygons	File size (Mb)
1	2	2	506	1	1	778	0.1
2	8	9	18728	2	8	61700	6.5
3	9	15	29894	2	8	71026	7.5

To create a more diverse range of tree size and shape the models were randomly scaled in height and width. However closer inspection showed excessive scaling produces many unnatural looking trees when scaled too far from their original dimensions. More trees and shrubs of different sizes could have been generated with relatively little effort.

Grass for our terrain was modelled as a procedural texture, rather than as geometric blades of grass. The texture looks plausible when viewed from above, but has several limitations when viewed from glancing angles. That is, it does not obscure lower parts of vehicles as it would in real life, especially in NT where in the wet season grasses can grow to more than 2 m tall. Figure 5 shows other features of geometric grass, when viewed from high observer elevation the gaps in the grass reveal the soil beneath, but when viewed at low observer elevations the gaps are obscured, which is not true of a texture. Also the scene is different when the sun is 'in front' or 'behind' the viewer, which is not true of a texture.

Planting 4 m² sparse overlapping clusters of grass with 1-3 polygons per blade was trialled on a small region in a separate project. Planting the entire 64 km² region was found to require too much memory due to the large number of planting instances required while planting only around vehicles would produce a transition that is all too obvious. The best way to handle grass in the future may be to model say 100 m² clusters. Only experimentation can say if there exists a cluster small enough to accurately hug angulated terrain and large enough to need only a feasible number of planting instances. Another important limitation of geometric grass is that unlike textured grass we currently have no easy means of planting it with spatial variations observed in overhead imagery.

Sun Front



Sun Behind




Figure 4: The scene changes as the viewer's elevation changes and the azimuth between the sun and the viewer changes. For the 'Sun Front' case the relative azimuth is 180° and in the 'Sun Behind' case it is 0°




Figure 5: Viewing geometric and textured grass from different observer elevation angles and relative azimuth angles between the observer and the sun

2.4 Planting

Trees were planted onto the terrain using CameoSim's tree planting scheme where 31 unique planting regions were manually defined and assigned the parameters: tree models, planting density, relative weightings, height and width scaling and variability among others. This can be a slow and clunky process if gradual variations in vegetation are sought. The planting process could be greatly improved with an improved GUI to make the process more interactive, or with an image processing tool that would identify tree locations and tree sizes with some constraints and weightings set by the user.




Figure 6: Screen shot of the terrain and the 31 planting regions

The terrain CGF file unnecessarily saves a copy of each tree for each region creating many copies of the same tree. Because of this, terrain CGF files used a large amount of disk space, around 450Mb. Ideally one copy of the tree should be saved and be referenced by multiple planting regions. If more tree models were created the terrain CGF file may become larger than the available system RAM, which may hamper or prohibit rendering.

2.5 Materials

The materials whose reflectance curves are shown in Figure 7 were used to classify the terrain. Soil, Grass, Scrub, Bark and Leaf were taken from a previous DSTO report [Boyd, 1995]. Sandstone, Sandstone Varnishes and Lake Bed are standard CameoSim Materials.

The thermal properties of materials that were changed were the transpiration factor and the convection characteristic length. Sap flow for Eucalyptus species in various regions of the NT have been reported by [Eamus et al, 1999]. To set the transpiration factor:

1. The sap flow is equated to the evaporated water flow
2. The power needed to vaporise that amount of water is determined
3. The transpiration factor that results in that power is chosen

We used the above process to work out transpiration factors in the dry and wet seasons. Whilst the sap flows in these seasons were not greatly different, the transpiration factors were very different due mainly to the humidity difference.

Table 2: Thermal properties of terrain materials

Material	Transpiration Factor		Convection Coefficient (m)	Density (kg/L)	Heat Capacity (kJ/kg K)	Thickness (m)	Thermal Conductivity (W/m K)
	Dry Season	Wet Season					
Soil	0	0.9	1	1.75	1.36	∞	0.9
Bark	0	0.9	1	0.80	2.41	0.2	0.21
Grass	0	0.9	0.004	0.545	2.64	0.04	0.14
Scrub	0	0.9	0.01	0.545	2.64	0.005	0.14
Leaf	0.02	0.9	0.005	0.545	2.38	0.002	0.17
Lake Bed	0.5	0.5	10	0.545	2.64	∞	0.14
Sandstone	0	0	10	2.7	0.92	∞	2.01
Sandstone Varnished	0	0	10	1.7	0.92	∞	0.4




Figure 7: Reflectance curves of dry season materials used to classify the terrain

2.6 Procedural Textures

Procedural textures were created to provide high frequency spatial variations in materials and 'bump' across the terrain for close up imagery. CameoSim provides a number of mathematical expressions for how materials can be mixed and how to generate the bump. The Octaves give the texture features at different scales; the Phase algorithm breaks up regular patterns; the Filter adds complexity to the texture; Bump changes the surface normals which creates the illusion of curvature.

Table 3: Some of the parameters used to create procedural textures

Texture	Base	Octaves	Phase	Octaves	Filter	Bump
Grass + Soil	Quick 2D	10	-	-	Linear	0.03
Scrub + Grass	Quick 2D	10	-	-	Linear	0.1
Scrub + Sandstone	Quick 2D	10	-	-	Linear	0.1
Sandstone	Quick 3D	1	Quick3D	3	-	1
Varnished Sandstone	Quick 3D	1	Quick3D	3	-	1
Lakebed	Quick 2D	3	-	-	-	-

The material mixes and the bump chosen for each texture were based on visual aesthetics only and required a significant amount of trial and error. No justification of the chosen parameters was attempted.

2.7 Material Classification

The CameoSim image based texture editor, cs-ctxedit, was used to transform the Google RGB image into a material map for the terrain, which in CameoSim terminology is referred to as a classification. The image was broken into three regions, namely escarpment, river and flat lands. This was done by creating two black-and-white selection images in MS Paint. cs-ctxedit uses the two images to assign each pixel to one of three layers where each layer is classified independently from one another. Once classified a pixel is assigned a mixture of bland materials and/or procedural textures. cs-ctxedit determines the mixture based on the colour of the pixel and a colour assigned by the user to each material.

Some limitations of the process are:

- Classification is limited to three band images
- Layers cannot transition gradually one to another because greyscale layer selection images are not allowed and thus sharp 1 pixel transition regions are observed
- Some issues were observed with using compressed RGB textures for classification where otherwise non obvious compression artefacts become obvious after classification
- A common misclassification problem is that grass is mapped to trees in the image, so it is preferable that somehow the trees be removed from the image before classification, however this proved too difficult for this task
- Another common misclassification was for shadows in the image to be incorrectly classified with the least reflective materials

For future projects we intend to use ENVI for material classification as it has a very large number of classification algorithms and will allow us to use the same classifications in CameoSim and DIRSIG.

2.8 Layering Materials

After specifying material layer thicknesses, CameoSim internally subdivides the layers during its temperature calculations. We have found that the degree of automatic subdivision is sometimes insufficient, in that replacing a layer with many smaller sub-layers can alter the temperature predictions. Since temperature calculations are relatively quick, there is no reason not to do this, especially for the topmost or thick layers where top to bottom temperature difference can be significant.

2.9 Thermal Sensitivity

We performed a sensitivity analysis to determine the necessary measurement fidelity of each thermal variable in the model. Temperatures were calculated for every material in the project, at every altitude, slope and azimuth in the thermal atmosphere. In turn each property of each material was both increased and decreased by 10% and new temperatures were calculated. The maximum temperature difference was calculated for each variable for each material. For all the materials in this database, solar absorptivity and thermal emissivity were the most sensitive. Note that 'Density' and 'Specific heat' are identical in all cases, as whilst CameoSim requires these variables to be entered separately, the conduction formula depends upon volumetric heat capacitance, which is the product of the two. For many materials measurement of the combined property volumetric heat capacitance will be far easier to obtain than the two separate properties, in which case this measurement can be entered into either density or specific heat with the other set to 1.

Table 4: Maximum temperature difference (°C) when a thermal property is changed by 10%

Material	Absorptivity	Emissivity	Char Len	Density	Conductivity	Specific heat
Grass	0.8	0.1	0.1	0.0	0.0	0.0
Lakebed	1.1	0.1	0.1	0.0	0.1	0.0
Sandstone	1.7	0.6	0.2	0.1	0.2	0.1
Scrub	0.8	0.2	0.1	0.0	0.1	0.0
Soil	1.6	0.4	0.2	0.5	0.1	0.5
Thin bark	2.1	0.6	0.3	0.3	0.0	0.3
Thick bark	2.0	0.5	0.2	0.2	0.1	0.2
Leaf	0.5	0.1	0.1	0.0	0.1	0.0

Table 5: Average temperature difference (°C) when a thermal property is changed by 10%

Material	Absorptivity	Emissivity	Char Len	Density	Conductivity	Specific heat
Grass	0.3	0.1	0.0	0.0	0.0	0.0
Lakebed	0.4	0.1	0.0	0.0	0.1	0.0
Sandstone	0.8	0.4	0.1	0.0	0.1	0.0
Scrub	0.4	0.1	0.1	0.0	0.0	0.0
Soil	0.8	0.3	0.1	0.4	0.1	0.4
Thin bark	0.7	0.2	0.1	0.1	0.0	0.1
Thick bark	0.7	0.1	0.1	0.1	0.0	0.1
Leaf	0.2	0.0	0.0	0.0	0.0	0.0

3. Modelling Spectral Atmospheres

CameoSim uses MODTRAN to predict propagation of electro-optical radiation over given lines of sight (LOS) for given atmosphere conditions. MODTRAN calculates the quantities:

- Lightshine – direct irradiance from the sun incident on a sun facing surface
- Sky Radiance – emitted and scattered radiance from the sky
- Path Radiance – emitted and scattered radiance over a finite length path
- Transmission – the fraction not attenuated nor scattered out of the path

CameoSim does not call MODTRAN for every ray traced when rendering an image, as this is too computationally expensive. Instead CameoSim calls MODTRAN for a set of rays, and uses interpolation to determine the quantities of every other ray during rendering. This interpolation can be a significant source of error. The rays over which MODTRAN is called are specified by the user with vectors for the following parameters:

- Solar Elevation – angle between the sun and the horizon
- Observer Altitude – initial altitude of path (in reverse-ray-tracing context)
- LOS Elevation – angle between path and the horizon
- LOS Range – path length
- Solar Observer Angle – angle between path and a path to the sun

All combinations of these parameters define either 2-, 4-, 5- or 4-dimensional grids for the quantities Lightshine, Sky Radiance, Path Radiance and Transmission respectively (refer to Table 6). The set of values for each parameter which defines the grid is also referred to as the spectral atmosphere geometry. The grid is not perfectly regular due to various boundary conditions and a mapping of LOS Range to Normalised LOS Range (refer to Section 3.2).

Table 6: Which quantities (top) depend on which parameters (left) are shown. All quantities are defined spectrally at MODTRAN's native cm^{-1} resolution, using $\pm \frac{1}{2} cm^{-1}$ bin averages.

	Lightshine (W/m ²)	Sky Radiance (W/m ² .sr)	Path Radiance (W/m ² .sr)	Transmission
Solar Elevation (deg)	X	X	X	X
Observer Altitude (km)	X	X	X	X
LOS Elevation (deg)		X	X	X
LOS Range (km)			X	X
Solar Observer Angle (deg)		X	X	

CameoSim also pre-calculates Sky Shine and Ground Shine, which are irradiances from an unobscured sky and unobscured ground which is modelled as a single material flat surface. These quantities are used only to render surfaces marked for 'quick' radiosity, which effectively ignore any surrounding objects. An additional parameter of Target Altitudes defines where these additional quantities are calculated. This report does not go into the detail of this as we avoid using quick radiosity.

3.1 Processing Time Issues

The current approach of CameoSim to pre-compute atmospheric quantities over a regular grid is simple in principle but can result in a large number of MODTRAN calls that in some cases cannot be processed in a reasonable time. This is partly due to two inefficiencies imposed by this method. Firstly only a fraction of the generated atmospheric data is likely to be used for rendering. Secondly placing a high density of points where atmospheric quantities are likely to vary or at a region of interest cannot be achieved without also adding points at other regions.

As an example the MWIR atmosphere in this report, which was generated by automatically adding points where interpolation errors were more than 10%, took several days to generate on a PC with two 64-bit AMD 246 Opteron processors. This computation time could be many times longer if rendering in the visible domain where multiple Solar Elevations are needed or a lower tolerance of interpolation errors is required. Processing time can then be further elevated by several orders of magnitude if MODTRAN higher fidelity scattering models with azimuth dependence are invoked.

An example of a spectral atmosphere geometry where many quantities of Path Radiance and Transmission are generated but not used is an airborne observer looking down at the ground. Here we need a sufficient number of paths that go from the observer to the ground however in doing so most paths are not used (see Figure 8). The problem is further compounded if the observer is required to span a range of altitudes as this requires additional LOS ranges and LOS elevations to ensure a significant number of paths at any altitude.




Figure 8: A Cartesian plot (left) and a polar plot (right) showing a grid of LOS Elevations and LOS Ranges for the quantities of Path Radiance and Transmission, for an observer at 2 km above ground. Highlighted in green are the only quantities used for a rendering of the ground. Here points going below ground have been truncated to ground altitude, however we are not sure how exactly this boundary condition is handled in CameoSim.

Another reason why some points are redundant is that atmospheric properties change quickly at some combinations of parameters, but not at others, however all combinations are generated. For example to accurately capture the Sun's halo it is useful to define several small Solar Observer Angles and LOS Elevations about the Solar Elevation (see Figure 9). Here small Solar Observer Angles are useful when paired with LOS Elevations about the Sun Elevation however other combinations of these parameters are less useful.




Figure 9: Plot showing a grid of LOS Elevations and Solar Observer Angles for a Sun Elevation of 30 deg. Highlighted in green is a region of expected high variability associated with the Sun's halo and highlighted in red is a region with unnecessarily high point density. The boundary highlighted in yellow shows the limit of meaningful parameter combinations. Points outside the boundary had their Solar Observer Angles increased or decreased to put them inside the boundary.

Ideally CameoSim should call MODTRAN to calculate quantities at points where they are used and more so in regions of significant change. However difficult, such an "on demand" adaptive irregular grid may provide a significant speed improvement to CameoSim's atmospheric model and allow processing of more accurate atmospheres that could use higher fidelity MODTRAN settings.

3.2 Range Normalisation

Atmospheric density decreases exponentially with altitude and the concentration of absorbents also decreases with altitude, thus transmission over a given range increases with altitude in a non-linear manner. In order to interpolate over a linearly varying quantity, prior to interpolating over transmission, range is mapped to normalised range. Normalised range can be interpreted as the equivalent sea level path range. Equation 1 defines Normalised LOS Range; the exponent used is a fit for how transmission tends to change for some typical atmospheres [Kirk 2007/2008].

$$Rn \equiv \int_0^R e^{-0.202(A + r \sin \theta)} dr = \begin{cases} \frac{e^{-0.202A}(e^{-0.202R \sin \theta} - 1)}{-0.202 \sin \theta} & \theta \neq 0 \\ R e^{-0.202A} & \theta = 0 \end{cases} \quad (1)$$

Where R is LOS Range (km), Rn is Normalised LOS Range (km), A is Observer Altitude (km), θ is LOS Elevation

3.3 Interpolation

At rendering time the atmospheric quantities for each ray are calculated using piecewise linear interpolation from neighbouring points [Kirk, 2007/2008]. The number of neighbouring points is either 2^2 , 2^4 or 2^5 depending on the number of dimensions in the grid for that quantity. For Transmission and Path Radiance the interpolation is carried out on the Transmission Coefficient and on Sky Radiance respectively to account for their typical exponential trends with path length [Siegel, 2002]. Once interpolated they are converted back to Transmission and Path Radiance, as per Equations 2 and 3.

$$\tau = e^{\beta \cdot R_n} \quad (2)$$

$$L_{path} = L_{sky}(1 - \tau) \quad (3)$$

Some details of the interpolation scheme are unknown, such as how values are interpolated when one or more of the neighbouring points have been moved or removed because they lie outside one of the boundary conditions.

Since both the interpolation scheme and atmospheric variations are complicated it is difficult to predict the grid spacing (spectral atmosphere geometry) that is both necessary and sufficient. If the geometry vectors have a sufficient number of values in areas where atmospheric quantities are changing, then most of the atmosphere can be captured sufficiently accurately and processed in the shortest amount of time. If there are many values in unnecessary regions then the atmosphere will process unnecessarily slowly and may still be inaccurate.

3.4 Geometry Optimisation

To ensure that the errors introduced by interpolation of atmospheric quantities are negligible we have written a MatLab based testing script. The software can be used to perform a check on a given atmosphere by comparing interpolated quantities with MODTRAN quantities. This check can then be used to identify problem regions where interpolation errors are large and refine the geometry iteratively until the atmosphere exhibits no significant interpolation errors.

The script uses CameoSim to render single pixel images of either a hot or cold blackbody, a diffuse mirror, or of the sun or sky to calculate the effective quantities that were used in the ray tracer. Two sets of images are rendered, one set uses the initial atmosphere where the paths require interpolation along one or more parameters and the other set uses a more

refined atmosphere where those paths have been calculated by MODTRAN and should not contain interpolation artefacts.

The script starts by checking interpolation along just one of the parameters. This is done so that if the error is significant only one parameter has to be refined, not all of the parameters. After single parameter errors are removed multiple parameter interpolation paths are checked.

When measuring path radiance and transmission errors individually neither absolute errors nor relative errors are appropriate. Specifying absolute error criteria is too difficult, since when spectral bands are very narrow, all quantities become small, and the error criteria must change accordingly. Relative error criteria do not work when the quantities involved are close to zero, as a very large relative error still results in a quantity that is close to zero. When path ranges are very short path radiance is close to zero, and when path range is very large transmission is close to zero. So the error must be calculated on a combination of path radiance and transmission, as neither one determines the final ray radiance on its own:

$$L_i = L_{path} + \tau L_o \quad (4)$$

Where L_i is the incident radiance, L_{path} is Path Radiance, τ is Transmission and L_o is the radiance leaving the surface.

To address this problem we define minimum and maximum surface radiance for our scene, by rendering imagery with a simple atmosphere and disabling atmospheric effects on the 'eye' rays (the first ray traced). These are independent of paths between observer and the scene. For each waveband we then define minimum and maximum ray radiance and calculate the percentage errors in these two new quantities caused by interpolation.

In order to prevent the calculation of unnecessary geometry points, CameoSim allows a maximum altitude to be specified, and if a path begins or ends above this altitude then this path will not be calculated. This works poorly when the sensor is the highest object in the scene. The maximum altitude is specified as the height of the sensor, and all upward looking paths are skipped. All downward paths are calculated, even ones that end well above the terrain. A better scheme would be to specify a maximum terrain altitude, and skip any path that ends above this altitude. We have implemented this in our script so that interpolation errors are not tested above a certain altitude.

An alternate method we trialled in checking for interpolation artefacts was the image convergence approach. Here a given atmosphere was used to render an image. Then more spectral atmosphere geometry values were added to the atmosphere and the image re-rendered. If no significant differences were observed, then the original atmosphere geometry was deemed sufficient. If significant differences occurred then geometry must be refined and the process repeated. This approach had significant drawbacks:

- Definition of the initial atmosphere
- How to relate a bad pixel to geometry values
- The process is slow and awkward and must be repeated if the observer is moved

3.5 Example

Figure 10 and Figure 11 illustrate an example of the atmosphere refinement process on a crude LWIR atmosphere that was used in this report. The 7-12 μm waveband was split into 10 equal bands by equal wavelength so as to also resolve the errors spectrally. Thirteen iterations were required to remove all interpolation errors greater than 10%. The initial atmosphere geometry begins with only the maximum and minimum of each parameter, with some additional key values such as LOS Elevation = Solar Elevation. If the initial atmosphere contained more geometry, a best guess as to what was needed, the reduction in interpolation errors would have required fewer iterations to reach the desired maximum error but at the price of greater processing time per iteration. At this point we are not sure which initial condition type leads to the least total processing time.




Figure 10: The maximum interpolation error found in the spectral atmosphere is reduced by iteratively adding spectral atmosphere geometry points. Initially each parameter is treated separately but when the individual errors are all below the requirement the multiple parameter combinations are checked.




Figure 11: The number of points in the spectral atmosphere geometry increases with each iteration of the algorithm

4. Modelling BRDFs

A Bidirectional Reflectance Distribution Function (BRDF, $f(\cdots)$) takes account of the bi-directional light reflecting properties of a surface. A BRDF can be based on a detailed physics model of light interactions with the surface (e.g. Cook Torrance, He et al.) or on simple parameterisations that have less physical basis (e.g. Ward, Phong, Lafortune, Ashikhmin). Bidirectional reflectance of a surface is typically measured spectrally at multiple source-observer orientations and data is fitted to a suitable BRDF. The models are often used in physics based renderers to sensibly interpolate and extrapolate measured data to all other source-observer orientations.

4.1 Theory

BRDF models can be classified into following categories:

- Data-driven – interpolate dense sets of measured data
- Phenomenological – designed empirically to follow trends of measured data
- Theoretical – derived from physical assumptions of surface properties

Bidirectional reflectance of a surface can depend on a number of parameters, which requires the use of exceedingly complicated BRDF model types:

- Lambertian $f()$ - source-observer direction independent
- Isotropic $f(\theta_s \theta_o \Delta \phi)$ - source-observer direction dependent, rotationally invariant
- Anisotropic $f(\theta_s \phi_s \theta_o \phi_o)$ - source-observer direction dependent
- Spectral $f(\cdots \lambda)$ - wavelength dependent
- Textured $f(\cdots x_o y_o)$ - point of observation dependent
- Translucent $f(\cdots x_s y_s)$ - point of incidence dependent
- Polarised $f(\cdots \bar{E})$ - polarisation dependent

Some useful constraints on a BRDF are energy conservation which prohibits a surface from reflecting more than 100% of the light incident on it, and Helmholtz reciprocity which states that bidirectional reflectance does not change upon interchange of source and observer:

$$\int_H f(\theta_o \phi_o \cdots) \cos(\theta_o) d\omega_o \leq 100\% \quad (5)$$

$$f(\theta_s \phi_s \theta_o \phi_o) = f(\theta_o \phi_o \theta_s \phi_s) \quad (6)$$

Physics based ray tracers generally try to calculate the observed radiance leaving a surface, L_o , given its BRDF, f , and incident radiance, L_s , by evaluating the following integral over the upper hemisphere, H , of the surface:

$$L_o(\theta_o \phi_o \cdots) = \int_H f(\theta_o \phi_o \theta_s \phi_s \cdots) L_s(\theta_s \phi_s) \cos(\theta_s) d\omega_s \quad (7)$$

Usually little is known about the distribution of L_s and the ray tracer is used to sample L_s over the hemisphere in N sample directions. To achieve convergence of L_o with the least number of sample rays it is desirable to use importance sampling where ideally, if nothing is known about L_s , a normalised $f(\theta_s \phi_s) \cos(\theta_s)$ is used as the probability distribution function (PDF) for generating the sample directions $\theta_{s_i} \phi_{s_i}$.

$$L_o(\theta_o \phi_o) = \frac{n}{N} \sum_{i=1}^N L_s(\theta_{s_i} \phi_{s_i}) \quad (8)$$

$$PDF(\theta_s \phi_s) = \frac{1}{n} f(\theta_s \phi_s) \cos(\theta_s) \quad (9)$$

$$n = \int_H f(\theta_s \phi_s) \cos(\theta_s) d\omega_s \quad (10)$$

In order to generate sample directions the PDF must be integrated and inverted to solve for $\theta_{s_i} \phi_{s_i}$. This is usually impossible because the cumulant C is often not separable [Lawrence et al, 2004]:

$$C(\theta_{s_i} \phi_{s_i}) = \int_0^{\phi_{s_i}} \int_0^{\theta_{s_i}} PDF(\theta_s \phi_s) \sin(\theta_s) d\theta_s d\phi_s \quad (11)$$

where often: $C(\theta_{s_i} \phi_{s_i}) \neq a(\theta_{s_i})b(\phi_{s_i})$

This problem can be overcome by:

- sampling only the invertible part of a BRDF or;
- approximating the BRDF with a simplified invertible BRDF or;
- taking a numerical tabulated approach such as that described by [Westlund and Meyer 2002].

4.2 CameoSim BRDFs

CameoSim has a default Lambertian model and the following isotropic BRDF models: Ashikhmin Shirley (isotropic), Cook Torrance, Lorentzian Peaks model and Sandford Robertson.

Other than being isotropic the simplicity of these models limits what BRDF shapes they can model, which may lead to a significant lack of capability to fit measured data. Even the Lorentzian peaks model is limited because the peaks can only lie in the plane of reflection. In the future it may be useful to identify if additional BRDFs are need, such as iBRDF or NEF Beard-Maxwell [Westlund and Meyer 2002].

4.2.1 Ashikhmin Shirley (isotropic)

This is the isotropic simplification of the mostly empirical Ashikhmin Shirley model described in [Ashikhmin and Shirley 2000]. The three shape parameters are:

- Reflectance (spectral)
- Specular reflectance (spectral)

- Shininess exponent

Remarks:

- Obeys reciprocity
- White specular reflections
- Mirror-like at grazing angles
- The first parameters can be related to normal incidence reflectance but are not physical in the sense that actual reflectance is always below this value where the underestimate is a function of the other parameters

4.2.2 Cook Torrance (uncoated/coated)

This is predominantly a theoretical model derived using the assumptions that the surface is a collection of perfectly specular planar microfacets that obey Fresnel reflection laws [Cook and Torrance, 1982]. The two parameters used by the model are:

- Relative index of refraction (spectral)
- Root mean square of microfacets slope

Remarks:

- Obeys reciprocity
- White specular reflections
- Mirror-like at grazing angles

4.2.3 Lorentzian Peaks Model

A data driven model designed to fit data with N Lorentzian peaks. The model uses one shape parameter and N×3 piecewise linear shape functions, where each function is defined using the pairs [angle of incidence, value]:

- Reflectance (spectral)
- N-reflection angle function(s)
- N-peak width function(s)
- N-relative peak strength function(s)

Remarks:

- Number of parameters depends on number of peaks and number of points per peak
- Lobes must lie on the plane of reflection thus limiting data that can be fitted
- Model can significantly violate reciprocity if the fit or the initial data is not sensible

4.2.4 Sandford Robertson

This is largely an empirical model, initially designed in 1985 to fit data for painted metal surfaces on aircraft in the LWIR. The CameoSim implementation of the model appears to be similar but not identical to that description in the revised version by [Sandford and Robertson 1994]. The model uses one physical parameter and three empirical shape parameters:

- Hemispherical emissivity (physical parameter) (spectral)
- Diffuse reflectivity (spectral)

- Grazing angle reflectivity
- Width of specular lobe

Remarks:

- Not designed for grazing angles $>80^\circ$
- Violates reciprocity, more so at grazing source-observer elevations

4.3 Process of Verifying CameoSim Importance Sampling

Implementation of BRDFs into ray tracing with efficient and accurate importance sampling is not trivial. Thus we have developed a simple MatLab script that can be used to perform the following:

- Visualise CameoSim BRDFs used for direct light sources such as the Sun
- Visualise the effective BRDFs produced by CameoSim's importance sampling
- Visualise the BRDFs using the analytic forms described in the literature
- Perform parameter optimisation to fit a model to data in the chi squared sense

To generate a visualisation of a CameoSim BRDF the user specifies a three dimensional parameter grid using three lists of observer elevations, source elevations and delta azimuths. For each source-observer configuration the script modifies a CameoSim project to position the observer and a source in the corresponding positions. Plate BRDF properties are specified manually and the project rendered. The script then uses CameoSim predicted single pixel radiances and the known incident irradiance to calculate bidirectional reflectance values and display them as a set of 3D surface plots, one for each angle of incidence.

$$f(\theta_s, \theta_o, \Delta\phi_o) = \frac{L_o(\theta_o, \phi_o)}{dE_s(\theta_s, 0)} \quad (12)$$

where $dE_s(\theta_s, 0) = L_s \Delta\omega_s \cos \theta_s$

The script can also be used to look at the transmitted radiance in a similar manner and to estimate directional hemispherical reflectance r :

$$r(\theta_s) = \int_H f(\theta_s, \theta_o, \Delta\phi_o) \cos \theta_o d\omega_o \quad (13)$$

CameoSim BRDF plots for direct light sources are generated by using the sun as the source. In CameoSim the sun is omitted from importance sampling. Instead rays are fired directly at the sun and reflected radiance is calculated directly from the BRDF. This is efficient because the sun's position is known and it is a significant source in the VIS to MWIR. If the Sun was importance sampled with the rest of the hemisphere a great deal more radiosity rays would be required to ensure enough of them hit the sun.

To generate plots of the effective BRDF produced by CameoSim's importance sampling the sun is replaced by a blackbody disk with the same solid angle as the sun that is hot and radiating. The target surface is set for full radiosity with several million radiosity rays. The ray

tracer fires the rays into the surroundings which return zero incident radiance everywhere except for those that hit the blackbody.




Figure 12: A schematic showing the difference in how reflections from a surface are calculated for when the sun is the source, using a direct approach, or a blackbody is the source, using the importance sampling approach.

Ideally an accurate importance sampling technique and a sufficient number of radiosity rays should produce the same effective BRDF as that used in the direct approach.

4.4 Verification Results

As a result of this verification work some inconsistencies in the Cook-Torrance, Sandford-Robertson and the Lorentzian Peak models were identified. This information was communicated to the CameoSim developers and a patch was issued in September 2008 for CameoSim v5.9.0. However Figure 13 shows an example of the Ashikhmin model where some significant inconsistencies remained. The importance sampling was again changed in June 2011 for CameoSim v5.13, though at this time DSTO has not studied the changes.




Figure 13: Example of a polar plot of the Ashikhmin Shirley model with the parameters: reflectance=0.5; specular reflectance=0.5, shininess exponent=50, incidence=45° (left), 75° (right). CameoSim BRDF for direct light sources (black) and effective CameoSim BRDF used for importance sampling (red). The red exhibits some noise because of a finite number of radiosity rays used, but it can still be clearly seen to differ from the black, especially for larger angles of incidence, which indicates that an inaccurate importance sampling algorithm is being used.

Even with the patch applied the fixed models were all found to show some inconsistencies for arbitrarily chosen parameters. In the future it may be useful to evaluate if other BRDF sampling algorithms such as those described by [Lawrence et al 2004] or [Westlund and Meyer 2002] should be implemented.

We have also attempted to reproduce CameoSim's BRDF models for direct light sources in MatLab, but with limited success. At the present only the Ashikhmin Shirley model has been verified to be consistent with that described by [Ashikhmin and Shirley 2000]. The causes for the inconsistencies in Cook Torrance, Sandford Robertson and Lorentzian Peaks models are yet to be identified.

If analytic forms of the BRDFs are available our script can also optimise model parameters to find the best fit of the model to measured data, in the chi squared sense. The script uses MatLab's fitting algorithms (such as Levenberg-Marquardt) which works by iteratively varying model parameters and comparing the model with data. Fitting the Lorentzian peaks model is challenging due the model having a large number of parameters and due to problems of finding many local minima in the fit.

5. Rendering Settings

When rendering a project it is up to the user to ensure that settings are of high enough fidelity that the modelling software produces the best image it is capable of, but not so high as to require excessive rendering time. The method we use to determine the required settings is to increase them until the image changes are insignificant. This is a very imprecise metric, as what is 'insignificant' depends upon the person or algorithm viewing the image. The highest criterion of accuracy would be: 'no pixel changes', however this is too strict for almost all viewers, for instance most people will not notice a pixel noise of ± 5 grey levels in a small percentage of pixels. Also it must be noted that every sensing device introduces random noise to images, which is a possible baseline to compare against.

5.1 The Settings to be Optimised

The following settings need to be optimised:

- Minimum, maximum and the contrast setting of sample rays
- Whether to use super sampling or Poisson sample ray distribution
- Maximum number of radiosity rays, whether true radiosity is necessary, max ray depth and min ray weight
- Thermal shadowing rays and time span
- Optical shadowing rays and whether to use soft or hard shadows
- Degree of over-sampling used
- The number of factored bands to use

The first thing to note is that it is impossible to test these settings independently. For example consider sample rays and radiosity rays. If we hold the number of radiosity rays constant and increase the number of sample rays, we will have found the number of sample rays to render with that number of radiosity rays. If the number of radiosity rays was changed, then the required number of sample rays would have to be recalculated. All the settings can only be said to be high enough when improving any setting results in insignificant image changes.

The first setting that we verified is the number of sample rays. By using an all 'quick radiosity' scene the image is made independent of the number of radiosity rays. There are four sample ray settings that have to be controlled:

- Number – Used to space apart the sample rays
- Minimum – The minimum number of sample rays used per pixel
- Contrast setting – Criterion used to determine if more sample rays are necessary
- Maximum – The number of sample rays per pixel used if the contrast setting is exceeded

We always set Number to be very high as we have found this produces high fidelity images with a low performance penalty. The Minimum / Maximum scheme would be of most benefit with a mostly bland scene, and a small section which has high sub-pixel contrast. In such a scene the Minimum number of rays would be used for most pixels and the Maximum would be used where necessary. In the scenes we have rendered we have found that there is high

frequency content everywhere. It is extremely difficult to know what the Contrast should be, because if you set the minimum too low you may not discover an important region of the pixel. We decided to use the same number of rays for each pixel, i.e. Minimum = Maximum. Typically 200 rays per pixel were found to be sufficient.

Increasing the number of thermal shadowing rays, and increasing the time span made no significant difference to the image. Increasing from hard to soft optical shadowing rays and changing from quick to slow also made no significant image difference.

CameoSim has three ways to determine the quantity of the incident light on a surface: 'Quick' radiosity, 'Full but not true' radiosity and 'Full and true' radiosity. Of the three only 'Full and true' radiosity uses a physically accurate model by tracing the incident light from multiple directions. We have found that when the number of sample rays is high scenes can be rendered with 'Full and true' radiosity with a surprisingly small number of radiosity rays. On this basis there is no reason not to use the 'true radiosity' setting. Typically a maximum of 10 radiosity rays was sufficient.

Over-sampling is generating more pixels in an image than the sensor has detectors. Prior to rendering an image, it is useful to consider what post-processing will be done to the image (see Chapter 6). In most instances it will be essential to render images at a higher resolution than that of the sensor to properly account for aliasing. In a real system, as light progresses through the sensor all spatial frequencies are reduced (blurring) and when the image is formed at the detector plane spatial frequencies higher than the sensor Nyquist spatial frequency are aliased [Holst 1998]. However the order is different in synthetic image generation as aliasing occurs when an image is rendered. By generating over-sampled imagery this preliminary aliasing is reduced, blurring is done in post processing, then the image is down-sampled which is another source of aliasing. The downsides to over-sampling are that the RAM and disk space usage is higher. The degree of oversampling required can be evaluated by increasing the degree of oversampling until the resultant image is insignificantly different from a lower degree of oversampling.

CameoSim is an inband tool. That is, all spectral properties are averaged over small finite spectral wavebands. Using fewer, wider bands reduces the rendering time of images but also reduces both the spectral resolution of atmospheric quantities and the reflectance properties of surfaces. Care must be taken when using wide bands however as some sensors, i.e. hyperspectral sensors, can produce imagery at very high spectral resolutions and to ensure fidelity the model must also be solved at a very high spectral resolution. The highest fidelity rendering CameoSim can produce is defining the spectral response curve of each band to be at the resolution of MODTRAN. For this project the rendering time was dominated by the tracing of sample and radiosity rays and using very narrow bands did not significantly slow down rendering.

6. Post-Processing Imagery

Not all the physical phenomena are accounted for by CameoSim in the formation of an image. In some cases it is possible to improve the simulation fidelity by post-processing the images to account for various atmospheric and sensor related effects.

6.1 Accounting For Atmospheric Scattering and Turbulence

In the absence of an intervening medium light leaving a surface would propagate without loss towards any sensing device. However for many remote sensing applications the sensed light is significantly altered by the effects of the intervening atmosphere. CameoSim uses MODTRAN which models atmospheric scattering, absorption, emission and refraction to calculate lightshine, sky radiance, path radiance and transmission along a path, (see Chapter 3). However there are two important physical phenomena that are not modelled well: the scattering of light due to aerosols and particles, and the refraction of light due to a changing index of refraction (imaging through turbulence).

To calculate transmission along a path MODTRAN accounts for both light scattering into and out of this path, and there are significant limitations in the way MODTRAN tracks both of these quantities. To be fully consistent the light scattered into the path would be dependent upon the objects in the CameoSim model, but this is not available to MODTRAN which assumes the world consists of a featureless earth surface. Similarly the radiance scattered out of a path cannot interact with the CameoSim model or be scattered into another path. By failing to account for this CameoSim images may be inaccurate if path radiance due to scattering is significant and if nearby scene geometry significantly alters the ambient lighting conditions of those paths. We are not aware as yet of any way this can be accounted for by post-processing the CameoSim images.

To calculate the direction taken by a ray of light MODTRAN uses the index of refraction of the atmosphere which is modelled as only varying (slowly) with height. However the index of refraction can change rapidly and locally and so the image observed by the sensor can be significantly affected as each pixel in the image can be affected differently. The effects of imaging through turbulence can be complex and will manifest differently depending upon the exposure time of the sensor.

It is possible in certain circumstances to account for imaging through turbulence by post-processing the CameoSim images. When imaging through turbulence with a long exposure time (short exposure time is defined as $\ll 10$ ms [Kopeika 1998]) the effect is a blurring of the image. An optical transfer function (OTF) is a way to specify how an image is blurred and a turbulence OTF is given by [Kopeika 1998] however it must be noted that this OTF is a function of range and the only valid application would be an image with all pixels having approximately the same range to the sensor. A novel approach would be to make use of the ability of CameoSim to calculate the range of each pixel, and use a different OTF for each pixel, however the validity of this approach has not been verified.

6.2 Sensor Effects

CameoSim images are rendered as though the sensor is an ideal pinhole camera that is free from the effects of diffraction and all parts of the image are simultaneously in focus. Depending upon the application it may be necessary to modify the CameoSim image to account for how a non-ideal sensor forms an image.

There are a great number of effects that can be accounted for in post-processing as there are many sensor models that work on the assumption that the radiance at the front of the aperture is perfectly sharp and noise-free, which matches what CameoSim produces.

Imagery can be blurred by a wide range of OTFs to simulate the effects of radiance propagating through the sensor [Lloyd 1982] which include:

- Diffraction – the effects of light passing through an aperture
- Optical aberrations – the effects of light passing through a lens
- Defocus – the effect of out of focus viewing
- Random vibration – the effect of the sensor platform vibrating
- Detector – the effect of non-zero width of the detector
- Sample & hold – the effect of scanning the detector

Noise can be modelled by adding a zero mean Gaussian to the pixel values.

The analogue to digital conversion can be modelled by reducing the number of bits in the digital representation of the radiance values.

6.3 Image Processing Software Choices

All image processing performed for this project were done in MatLab as we have found this language is simple and fast to code in. The drawback to this approach is that MatLab does not come with a built in function to read the CameoSim Floating Point Image (FPI) format and we had to write it. More off-the-shelf alternatives we could have chosen are the image processing module for CameoSim or the Compass software package

The image processing module for CameoSim can read the FPI format and contains built in primitive filters, however most real filters would have to be written in C. Whilst this is the most integrated method of applying image processing a significant drawback to this approach is that the image must be processed by the computer that CameoSim is installed on, taking up valuable computer resources.

Alternatively the stand alone product Compass can be used to process images. Features are

- Can read and write the FPI format
- Includes many built in OTFs (but doesn't include turbulence or defocus OTFs)
- Extensions to the built in functions can be written in MatLab
- Can import OTFs from NVTherm

Drawbacks are that the image is processed by the single computer that Compass is installed on.

6.4 Future Work

Detector non uniformity & correction is the major source of noticeable image degradation for modern sensors as their sensitivity is so high. Whilst sensors attempt to correct for non uniformity by doing calibration in the field, the results are not perfect.

The blur resultant from atmospheric scattering will need to be included, especially for satellite images, where scattering can be more significant than turbulence [Kopeika 1998].

7. Conclusions

High fidelity results depend upon two things, having the correct data and using it well. CameoSim is a very powerful tool, but it is also a very complex one. It has taken us several years of use before we have reached the point where we can claim to understand how the model works. There are many inputs to CameoSim that rely upon the user's judgement to determine that the quality is 'high enough'. This can be assisted by having software that produces objective measures of error and methods of optimising the settings to achieve a specified level of quality. This report describes how we have done this for specifying the atmosphere geometry and the ray tracing settings.

Creating an exact geometry model of a real scene is still very difficult and time consuming, and will never be a substitute for, only complementary to, collection of imagery with real sensors. The advantages of synthetic imagery are that the elements within a scene can be controlled and set to be whatever is desired. Whilst there are limitless scenes that can be generated, the most frequently beneficial will be scenes designed to answer 'what if' questions, by constructing a scene and then changing a single feature, e.g. the range of the target, the resolution of the sensor, etc.

Having understood the models, and constructed a scene, the next task is to render an image. If a scene is too complicated this can be computationally infeasible. In this task we attempted to model a large terrain, such that it is viewable from anywhere. Perhaps this approach was a mistake, as whilst there is great flexibility in such an approach, the modelling and computational effort may make it infeasible. In future it may be best to construct a terrain with the exact observer location(s) in mind.

Having created imagery in CameoSim, the next question is 'Is it correct?'. Without a validated synthetic image generation tool, the validity depends upon faith that all the simplifications to the physics will lead to only insignificant errors. Validation is the next step for ISRD, however it is difficult to generate a test that CameoSim can fail. A common flaw in validation studies is not having enough data, and so any prediction errors can be blamed not on the software, but on the insufficient data.

7.1 Ongoing Issues

There are some elements to the NT scene that we would like to model, however it appears that, with CameoSim, these are very difficult or impossible to achieve.

CameoSim models the atmosphere as varying only in the vertical direction, thus clouds either totally fill the sky or are totally absent from the sky. Days with partial cloud cover cannot be modelled. DIRSIG seems to have solved this problem by allowing horizontal variations in atmosphere properties [Dobbs 2000].

CameoSim is limited in that there are no anisotropic or polarised BRDF models, only 4 simple isotropic models. CameoSim does have a Material API available that allows users to plug-in their own material parameterisations, however we have not investigated this API yet.

Processing the spectral atmosphere, and ensuring interpolation errors are minimal, is a very time consuming process as CameoSim will only utilize one CPU when processing, even when multiple CPUs or a cluster of computers is available.

8. Acknowledgments

Kind regards to Stephen Carr for providing direction and technical expertise in the writing of this report. Many thanks to the CameoSim team: Albert Kirk; Glenn Hutchings and Colin Stroud for ongoing technical support. The CameoSim team has read this document and taken on board the identified weak points and plan to deal with some of them within the next few releases. Finally thank you to all those involved in the review and publishing of this report.

9. References

- [Ashikhmin and Shirley, 2000] Ashikhmin, M., Shirley, P., (2000) *An Anisotropic Phong Light Reflection Model*, University of Utah
- [Boyd, 1995] Boyd, Russell (1995) *Colours of Northern Australia: Visible and Near-IR Reflectance of Natural Terrain Elements*, DSTO Technical Report DSTO-TR-01141
- [Carr and Brady, 2007] Carr, S., Brady, A. (2007) *Synthetic Image Generation for Infrared Sensor Studies*, DSTO Technical Report DSTO-TR-2006
- [Cook and Torrance, 1982] Cook, R.L., Torrance, K.E., (1982) *A Reflectance Model for Computer Graphics* ACM Transactions on Graphics, Vol. 1, No. 1, January 1982, Pages 7-24.
- [Dobbs, 2000] Dobbs, B.M (2000) *The Incorporation of Atmospheric Variation in DIRSIG* Rochester Institute of Technology, Rochester NY.
- [Eamus et al, 1999] Eamus, D., O'Grady, A., Hutley, L. (1999) Dry season conditions determine wet season water use in the wet-dry tropical savannas of northern Australia *Tree Physiology* 20, 1219-1226
- [Holst, 1998] Holst, G. (1998) *Sampling, aliasing, and data fidelity for electronic imaging systems, communications, and data acquisition* JCD Publishing, Winter Park, FL, USA
- [Kirk, 2007/2008] Kirk, Albert (2007,2008) *Private communications*, Lockheed Martin UK INSYS Ltd
- [Kopeika, 1998] Kopeika, N. (1998) *A System Engineering Approach To Imaging* SPIE Press, Bellingham, Washington, USA
- [Lawrence et al, 2004] Lawrence, J., Rusinkiewicz, S., Ramamoorthi, R., (2004) *Efficient BRDF Importance Sampling Using A Factored Representation*, Proc. ACM SIGGRAPH, vol. 23, p. 496-505.
- [Lloyd, 1982] Lloyd, J. (1982) *Thermal Imaging Systems* Plenum Press, New York, NY, USA
- [Sandford and Robertson, 1994] Sandford, B. P., Robertson, D. C., (1994) *Infrared Reflectance Properties Of Aircraft Paints*, Phillips Laboratory, Geophysics Directorate/GPOA and Spectral Sciences, Inc. Burlington, MA
- [Siegel, 2002] Siegel, R., Howell, J. (2002) *Thermal Radiation Heat Transfer - 4th Edition* Taylor & Francis, New York, NY, USA, p. 433.
- [Stokes et al, 1996] A Standard Default Color Space for the Internet – sRGB (October 2008), Stokes, M. , Anderson, M., Chandrasekar, S., & Motta, R. URL - <http://www.w3.org/Graphics/Color/sRGB>

[Westlund and Meyer, 2002] Westlund, H. B., Meyer, G. W., *A BRDF Database Employing the Beard-Maxwell Reflection Model*, Graphics Interface p. 189-201

[Winzell, 2006] Winzell, T. (2006) Basic Optical Sensor Model, Swedish Defence Research Agency Scientific Report FOI-R-2135--SE

Appendix A: Sample Imagery




Figure 14: A colour image of a shot taken at 830, showing Land Rovers (left), 2S6s (middle) and BRDM2s (right) at ranges of approximately 5.7, 8.3 and 10.5km from the sensor. The left of a pair is active, right is inactive.




Figure 15: Same shot as Figure 14 taken in MWIR (3-5 μ m full sensor response, image level=1.8 range=0.16 W/m².sr)




Figure 16: Same shot as Figure 14 taken in LWIR (7.5-10 μ m full sensor response, image level=23.5 range=0.48 W/m².sr)

Appendix B: Atmosphere Generation Specifics

Model: Winter - Tropical

Cloud: None

Volcanic: Background Moderate

Haze: Rural - 23km

Visibility: 100km

Scatter: Mie – Isaac – no azimuthal dependence

HPT profiles: Burro Of Meteorology, 8:30am, 15/9/2005, Darwin Airport

Solar Elevation: 34.83°

Low Exitance (per band, 7-12 μ m): 4.75 5.16 5.41 5.56 5.63 5.62 5.54 5.41 5.25 5.06 W/m²/sr

High Exitance (per band, 7-12 μ m): 5.93 6.29 6.35 6.36 6.32 6.29 6.23 6.18 6.04 5.82 W/m²/sr

LWIR Geometry (interpolation errors are less than 10%):

Observer Altitude (km):

0.03 0.05 0.06 0.09 0.15 0.27 0.52 1.0 1.9 2.1 2.3 2.5 2.8 3.2 3.6 3.8 4.1 4.3 4.5 5.4 5.8 6 6.1 6.3 6.5
6.7 6.9 7.1 8.0

LOS Elevation (deg):

-89 -62 -35 -26 -17 -13 -8.7 -4.4 -2.2 0.0 0.27 0.55 1.1 2.2 4.4 6.5 8.7 13 17 26 35 48 62 89

LOS range (km):

0.02 1.6 3.1 6.3 12.5 25

Solar Observer Angle (deg):

90

MWIR Geometry (interpolation errors are less than 10%):

Observer Altitude (km):

0.041 0.045 0.047 0.049 0.056 0.071 0.10 0.21 0.33 0.44 0.55 0.66 0.78 0.89 1.0 1.2 1.4 1.9 2.0 2.1
2.3 2.8 3.2 3.6 3.8 4.1 4.3 4.5 4.9 5.4 5.8 6.0 6.3 6.4 6.7 7.1 7.6 7.8 8.0

LOS Elevation (deg):

-89 -61.9 -34.8 -26.1 -17.4 -15.2 -14.2 -13.1 -12.0 -10.9 -9.8 -8.7 -7.6 -6.5 -5.4 -4.4 -3.3 -2.2 -1.9
-1.6 -1.4 -1.2 -1.1 -0.8 -0.5 -0.3 0.0 0.5 1.1 2.2 3.3 4.4 6.5 8.7 13.1 17.4 26.1 30.5 34.8 38.2 41.6
48.4 61.9 68.7 75.5 82.2 89

LOS range (km):

0.02 0.10 0.33 0.55 1.0 1.04 1.07 1.14 1.28 1.56 2.13 3.25 5.5 7.75 10 25

Solar Observer Angle (deg):

0 5.6 11 23 28 34 39 45 51 56 62 68 73 79 84 90 180

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA							
				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)			
2. TITLE Further Studies into Synthetic Image Generation using CameoSim				3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)			
4. AUTHOR(S) A. Brady and S. Kharabash				5. CORPORATE AUTHOR DSTO Defence Science and Technology Organisation PO Box 1500 Edinburgh South Australia 5111 Australia			
6a. DSTO NUMBER DSTO-TR-2589		6b. AR NUMBER AR-015-073		6c. TYPE OF REPORT Technical Report		7. DOCUMENT DATE August 2011	
8. FILE NUMBER 2009/1021594/1		9. TASK NUMBER INT 07/052		10. TASK SPONSOR DIGO		11. NO. OF PAGES 38	
						12. NO. OF REFERENCES 16	
DSTO Publications Repository http://dspace.dsto.defence.gov.au/dspace/				14. RELEASE AUTHORITY Chief, Intelligence, Surveillance and Reconnaissance Division			
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i>							
OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111							
16. DELIBERATE ANNOUNCEMENT No Limitations							
17. CITATION IN OTHER DOCUMENTS				Yes			
18. DSTO RESEARCH LIBRARY THESAURUS Computer generated imagery, Thermal models, Infrared imaging, Atmospheres, Remote sensing							
19. ABSTRACT Previous work by DSTO described the synthetic image generation process using CameoSim and RadThermIR; this report describes the progress made on the modelling of atmospheres, terrains and sensors. A major source of error in the synthetic image generation process previously identified has been resolved. The error concerns the use of interpolation in CameoSim when it uses MODTRAN to incorporate atmospheric effects in synthetic images. A software tool has been created which iteratively reduces the interpolation errors until they are insignificant. Validation studies are currently being planned in the visible to shortwave infrared. In preparation of the validation effort a study of BRDF models has been completed, which includes the physical plausibility of models, how measured data is fitted to models and how well CameoSim samples each model.							