imgrad

August 16, 2005

Abstract

This task approximates the gradient at each pixel of an image by doing least-squares fitting of a plane to the 3x3 array of values centred on each pixel.

1 Instruments/Modes

The task is not XMM-specific.

2 Description

This task approximates the gradient at each pixel of an image by doing least-squares fitting of a plane to the 3x3 array of values centred on each pixel. Since gradient is a vector quantity, two output images are required. The user can choose between a cartesian representation (in which the output images record the x and y components of the gradient) or a polar representation (in which the output images are the magnitude and azimuth of the gradient.).

'Least-squares fitting of a plane' sounds very grand but in fact the algebra boils down to the following:

$$(\nabla I)_{x,y} = \frac{1}{6} \begin{pmatrix} \sum_{j=y-1}^{y+1} I_{x+1,j} - \sum_{j=y-1}^{y+1} I_{x-1,j} \\ \sum_{i=x-1}^{x+1} I_{i,y+1} - \sum_{i=x-1}^{x+1} I_{i,y-1} \end{pmatrix}.$$

Here $I_{x,y}$ represents the image value at the (x, y)th pixel.

No gradient value is calculated at the edge: nulls are stored in the output at these pixels. If the input image contains null-valued pixels, all 8 nearest neighbours of such pixels are set to null in the output.

The azimuth, where this is to be calculated, is $\arctan(\nabla_y / \nabla_x)$.

3 Parameters

This section documents the parameters recognized by this task (if any).

Parameter	Mand	Type	Default	Constraints

imageset	yes	dataset	dummy_default	

The input image dataset.

outputstyle	no	type	cartesian	cartesian—polar

This controls the style of the output. If `outputstyle`='cartesian', the two output images store the x and y components of the gradient; if `outputstyle`='polar', the two output images store the magnitude and azimuth of the gradient.

xgradimageset	no	dataset	xgrad_image.ds	

If `outputstyle`='cartesian', this dataset is used to store the x component of the gradient.

ygradimageset	no	dataset	ygrad_image.ds	

If `outputstyle`='cartesian', this dataset is used to store the y component of the gradient.

gradsizeimageset	no	dataset	grad_size_image.ds	

If `outputstyle`='polar', this dataset is used to store the magnitude of the gradient.

gradazimuthimageset	no	dataset	grad_azim_image.ds	

If `outputstyle`='polar', this dataset is used to store the azimuth (from 0 to π) of the gradient.

4 Errors

This section documents warnings and errors generated by this task (if any). Note that warnings and errors can also be generated in the SAS infrastructure libraries, in which case they would not be documented here. Refer to the index of all errors and warnings available in the HTML version of the SAS documentation.

badOutputStyle *(error)*
The value of the `outputstyle` parameter was not recognized.

5 Input Files

1. A FITS dataset containing a 2-dimensional image array in the primary extension. The datatype of the image may be 8-bit integer, 16-bit integer, 32-bit integer, 32-bit real or 64-bit real.

6 Output Files

1. Two FITS datasets, each containing a 2-dimensional image array in the primary extension, the same size as the input array and containing all of the keywords (except those pertaining to DSS) of the input. Both images are of datatype 32-bit real. If `outputstyle`='cartesian', the two output images store the x and y components of the gradient; if `outputstyle`='polar', the two output images store the magnitude and azimuth of the gradient.

7 Algorithm

References