

NAME

pam - portable arbitrary map file format

DESCRIPTION

The PAM image format is a lowest common denominator 2 dimensional map format.

It is designed to be used for any of myriad kinds of graphics, but can theoretically be used for any kind of data that is arranged as a two dimensional rectangular array. Actually, from another perspective it can be seen as a format for data arranged as a three dimensional array.

This format does not define the meaning of the data at any particular point in the array. It could be red, green, and blue light intensities such that the array represents a visual image, or it could be the same red, green, and blue components plus a transparency component, or it could contain annual rainfalls for places on the surface of the Earth. Any process that uses the PAM format must further define the format to specify the meanings of the data.

A PAM image describes a two dimensional grid of tuples. The tuples are arranged in rows and columns. The width of the image is the number of columns. The height of the image is the number of rows. All rows are the same width and all columns are the same height. The tuples may have any degree, but all tuples have the same degree. The degree of the tuples is called the depth of the image. Each member of a tuple is called a sample. A sample is an unsigned integer which represents a locus along a scale which starts at zero and ends at a certain maximum value greater than zero called the maxval. The maxval is the same for every sample in the image. The two dimensional array of all the Nth samples of each tuple is called the Nth plane or Nth channel of the image.

Though the format does not assign any meaning to the tuple values, it does include an optional string that describes that meaning. The contents of this string, called the tuple type, are arbitrary from the point of view of the PAM format, but users of the format may assign meaning to it by convention so they can identify their particular implementations of the PAM format.

The Layout

A PAM file consists of a sequence of one or more PAM images. There are no data, delimiters, or padding before, after, or between images.

Each PAM image consists of a header followed immediately by a raster.

Here is an example header:

```
P7  
WIDTH 227  
HEIGHT 149  
DEPTH 3  
MAXVAL 255  
TUPLTYPE RGB  
ENDHDR
```

The header begins with the ASCII characters "P7" followed by newline. This is the magic number.

The header continues with an arbitrary number of lines of ASCII text. Each line ends with and is delimited by a newline character.

Each header line consists of zero or more whitespace-delimited tokens or begins with "#". If it begins with "#" it is a comment and the rest of this specification does not apply to it.

A header line which has zero tokens is valid but has no meaning.

The type of header line is identified by its first token, which is 8 characters or less:

ENDHDR

This is the last line in the header. The header must contain exactly one of these header lines.

HEIGHT

The second token is a decimal number representing the height of the image (number of rows). The header must contain exactly one of these header lines.

WIDTH

The second token is a decimal number representing the width of the image (number of columns). The header must contain exactly one of these header lines.

DEPTH

The second token is a decimal number representing the depth of the image (number of planes or channels). The header must contain exactly one of these header lines.

MAXVAL

The second token is a decimal number representing the maxval of the image. The header must contain exactly one of these header lines.

TUPLTYPE

The header may contain any number of these header lines, including zero. The rest of the line is part of the tuple type. The rest of the line is not tokenized, but the tuple type does not include any white space immediately following **TUPLTYPE** or at the very end of the line. It does not include a newline. If there are multiple **TUPLTYPE** header lines, the tuple type is the concatenation of the values from each of them, separated by a single blank, in the order in which they appear in the header. If there are no **TUPLTYPE** header lines the tuple type is the null string.

The raster consists of each row of the image, in order from top to bottom, consecutive with no delimiter of any kind between, before, or after, rows.

Each row consists of every tuple in the row, in order from left to right, consecutive with no delimiter of any kind between, before, or after, tuples.

Each tuple consists of every sample in the tuple, in order, consecutive with no delimiter of any kind between, before, or after, samples.

Each sample consists of an unsigned integer in pure binary format, with the most significant byte first. The number of bytes is the minimum number of bytes required to represent the maxval of the image.

PAM Used For PNM (PBM, PGM, or PPM) Images

A common use of PAM images is to represent the older and more concrete PBM, PGM, and PPM images.

A PBM image is conventionally represented as a PAM image of depth 1 with maxval 1 where the one sample in each tuple is 0 to represent a black pixel and 1 to represent a white one. The height, width, and raster bear the obvious relationship to those of the PBM image. The tuple type for PBM images represented as

PAM images is conventionally "BLACKANDWHITE".

A PGM image is conventionally represented as a PAM image of depth 1. The maxval, height, width, and raster bear the obvious relationship to those of the PGM image. The tuple type for PGM images represented as PAM images is conventionally "GRAYSCALE".

A PPM image is conventionally represented as a PAM image of depth 3. The maxval, height, width, and raster bear the obvious relationship to those of the PPM image. The first plane represents red, the second green, and the third blue. The tuple type for PPM images represented as PAM images is conventionally "RGB".

The Confusing Universe of Netpbm Formats

It is easy to get confused about the relationship between the PAM format and PBM, PGM, PPM, and PNM. Here is a little enlightenment:

"PNM" is not really a format. It is a shorthand for the PBM, PGM, and PPM formats collectively. It is also the name of a group of library functions that can each handle all three of those formats.

"PAM" is in fact a fourth format. But it is so general that you can represent the same information in a PAM image as you can in a PBM, PGM, or PPM image. And in fact a program that is designed to read PBM, PGM, or PPM and does so with a recent version of the Netpbm library, will read an equivalent PAM image just fine and the program will never know the difference.

To confuse things more, there is a collection of library routines called the "pam" functions that read and write the PAM format, but also read and write the PBM, PGM, and PPM formats. They do this because the latter formats are much older and more popular, so this makes it convenient to write programs that use the newer PAM format.

SEE ALSO

pbm(5), pgm(5), ppm(5), pnm(5), libpnm(3)