US005467110A

United States Patent [19]

White et al.

[11] Patent Number: 5,467,110

[45] Date of Patent: Nov. 14, 1995

[54] **POPULATION ATTRIBUTE COMPRESSION**

[75] Inventors: **James M. White; Vance Faber; Jeffrey S. Saltzman**, all of Los Alamos, N.M.

[73] Assignee: **The Regents of the University of California, Office of Technology Transfer**, Alameda, Calif.

[21] Appl. No.: **589,563**

[22] Filed: **Sep. 28, 1990**

Related U.S. Application Data

[63] Continuation-in-part of Ser. No. 350,675, May 12, 1989, Pat. No. 5,130,701.

[51] **Int. Cl.**[6] **G09G 1/28**

[52] **U.S. Cl.** **345/199**; 345/202; 348/404

[58] **Field of Search** 340/701, 702, 340/703, 750; 358/80, 81, 133; 364/521; 345/199, 202, 112; 348/404, 405, 406; 395/126

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,352,105	9/1982	Harney	340/703
4,580,134	4/1986	Campbell et al. .	
4,710,806	12/1987	Iwai et al.	340/703
4,751,446	6/1988	Pineda et al. .	
4,910,589	3/1990	Nagano et al.	358/80
4,975,769	12/1990	Aizu et al.	358/80
4,994,927	2/1991	Dixit et al.	358/133
5,021,971	6/1991	Lindsay	358/133

OTHER PUBLICATIONS

P. Heckbert, "Color Image Quantization for Frame Buffer Display," 16 Computer Graphics No. 3. pp. 297–306 (Jul. 1982).

Y. Linde et al., "An Algorithm for Vector Quantizer Design," COM–28 IEEE Trans. Comm. No. 1, pp. 84–95 (Jan. 1980).

J. H. Friedman et al., "An Algorithm for Finding Best Matches in Logarithmic Expected Time," 3 ACM Trans. Math. Software, No. 3, pp. 209–226 (Sep. 1977).

Yamaguchi et al; "Octree–related data structures and algorithms"; IEEE CG&A; Jan. 1984; pp. 53–59.

Doctor et al; "Display techniques for octree–encoded objects"; IEEE CG&A; Jul. 1981; pp. 29–38.

Primary Examiner—Ulysses Weldon
Assistant Examiner—M. Fatahiyan
Attorney, Agent, or Firm—Ray G. Wilson

[57] **ABSTRACT**

An image population having a large number of attributes is processed to form a display population with a predetermined smaller number of attributes that represent the larger number of attributes. In a particular application, the color values in an image are compressed for storage in a discrete look-up table (LUT). Color space containing the LUT color values is successively subdivided into smaller volumes until a plurality of volumes are formed, each having no more than a preselected maximum number of color values. Image pixel color values can then be rapidly placed in a volume with only a relatively few LUT values from which a nearest neighbor is selected. Image color values are assigned 8 bit pointers to their closest LUT value whereby data processing requires only the 8 bit pointer value to provide 24 bit color values from the LUT.

6 Claims, 4 Drawing Sheets

Microfiche Appendix Included
(1 Microfiche, 14 Pages)

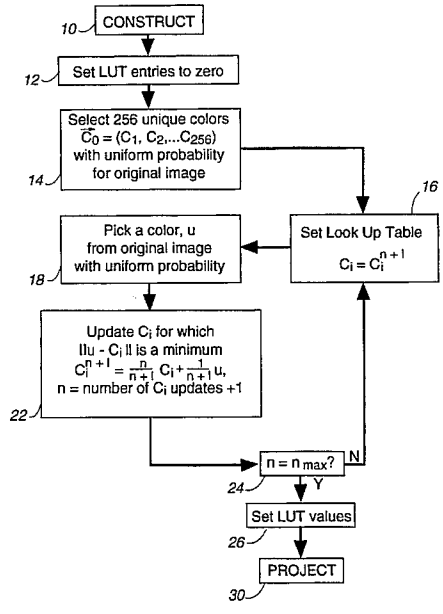

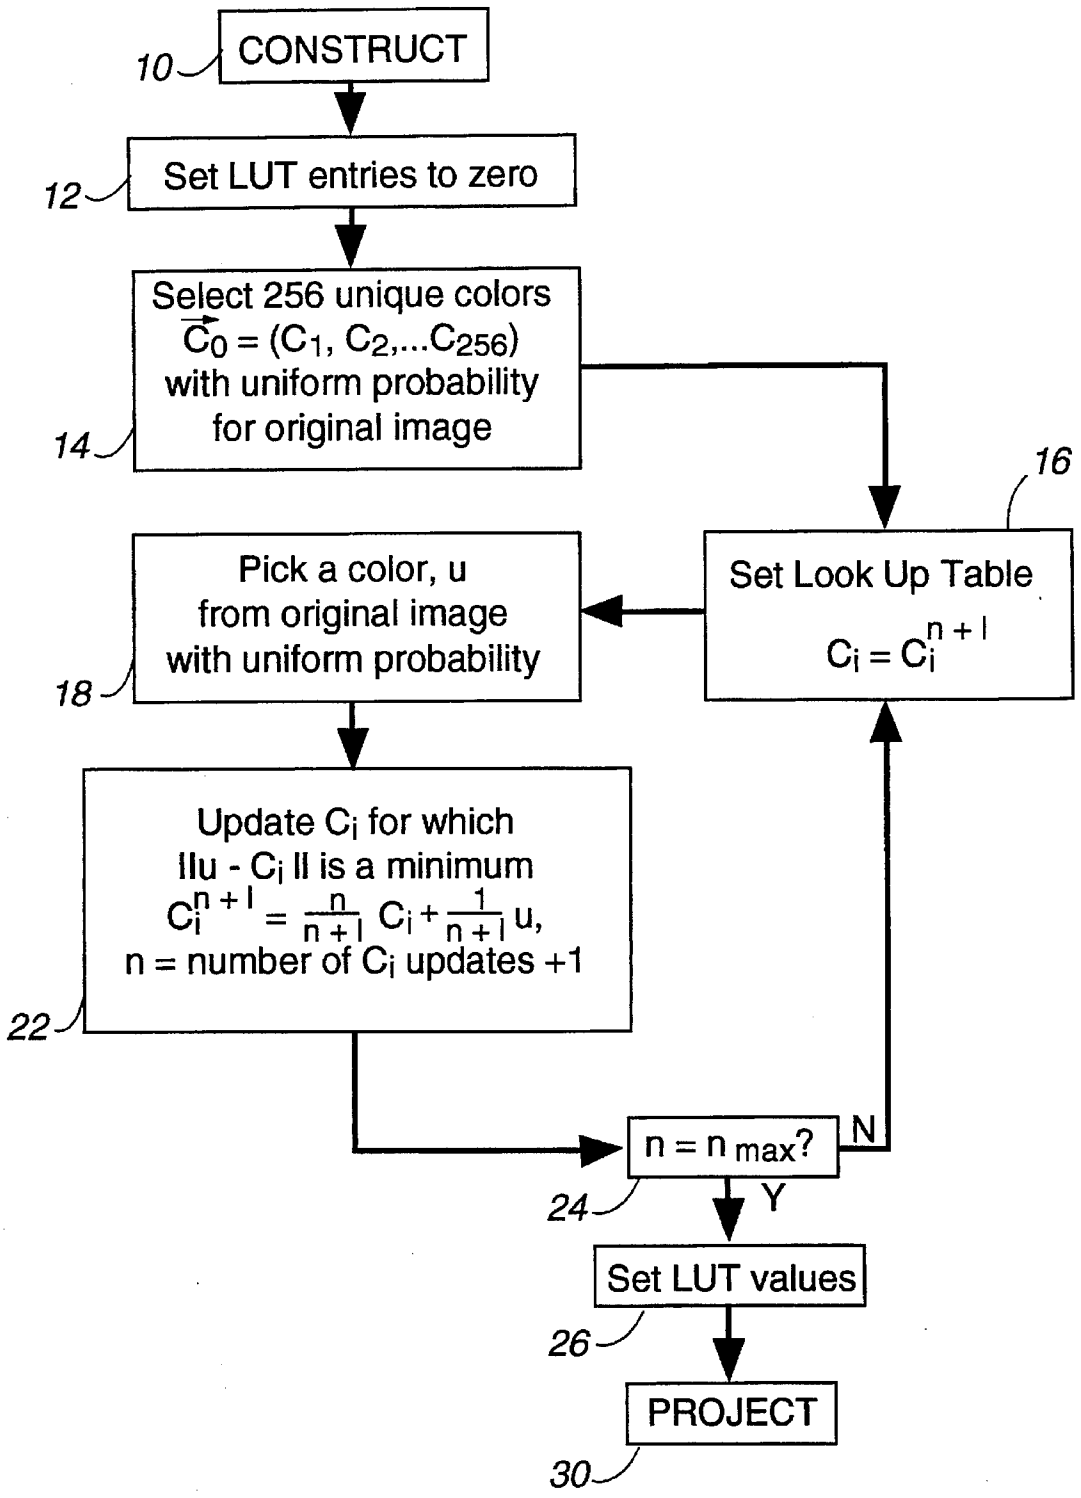
10
CONSTRUCT
12
Set LUT entries to zero
14
Select 256 unique colors
$\vec{C}_0 = (C_1, C_2, \ldots C_{256})$
with uniform probability
for original image
16
Set Look Up Table
$C_i = C_i^{n+1}$
18
Pick a color, u
from original image
with uniform probability
22
Update C_i for which
$\|u - C_i\|$ is a minimum
$C_i^{n+1} = \frac{n}{n+1} C_i + \frac{1}{n+1} u$,
n = number of C_i updates +1
24
$n = n_{max}$?
N
Y
26
Set LUT values
30
PROJECT

Fig. 1A

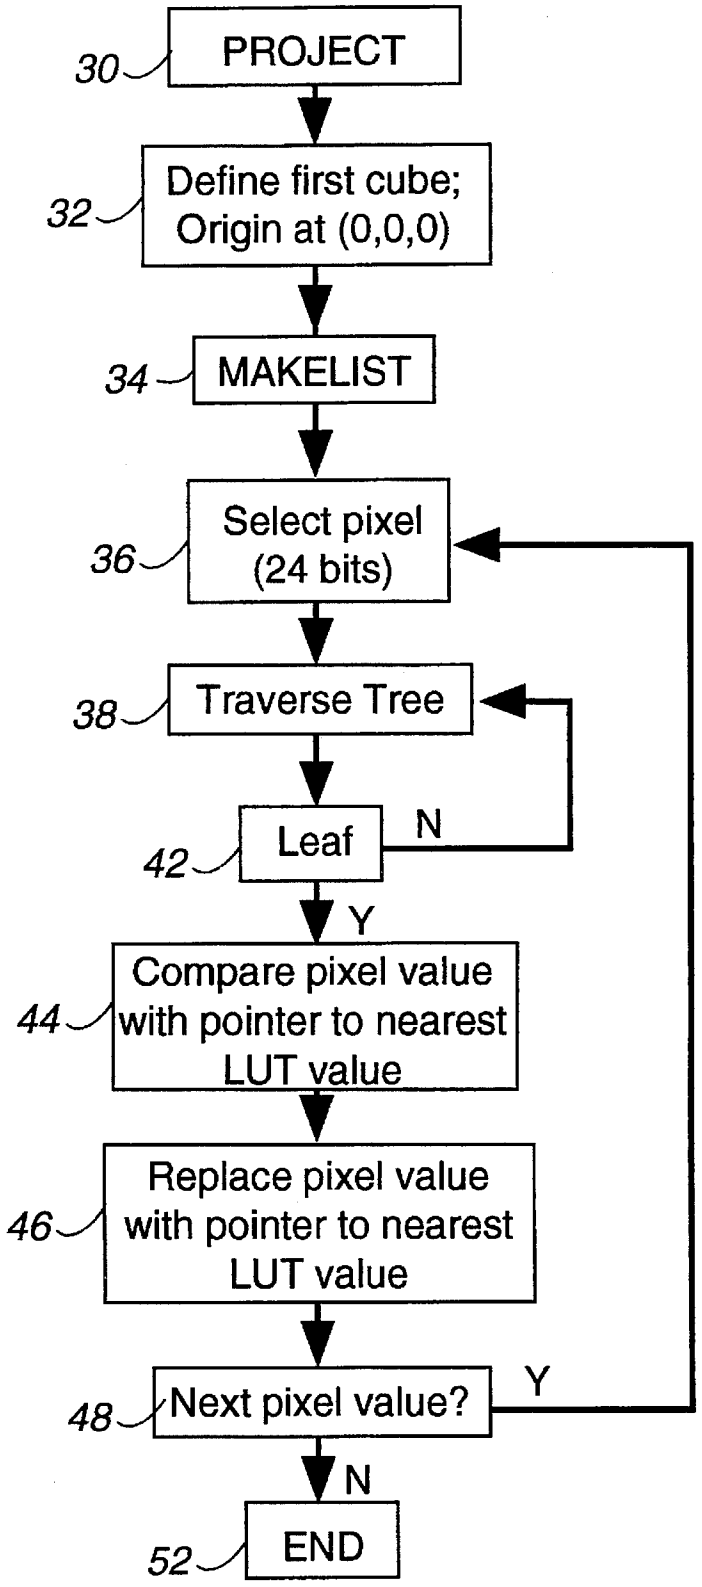
30
PROJECT
32
Define first cube; Origin at (0,0,0)
34
MAKELIST
36
Select pixel (24 bits)
38
Traverse Tree
42
Leaf
N
Y
44
Compare pixel value with pointer to nearest LUT value
46
Replace pixel value with pointer to nearest LUT value
48
Next pixel value?
Y
N
52
END

Fig. 1B

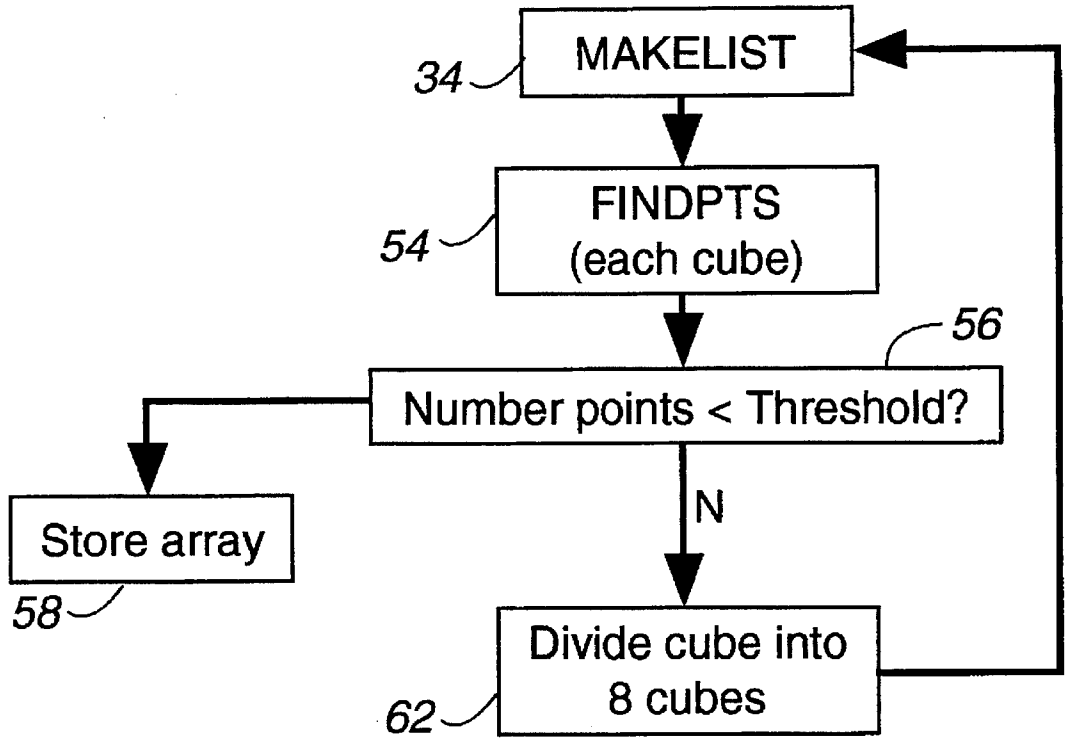
MAKELIST
34
FINDPTS
(each cube)
54
56
Number points < Threshold?
N
Store array
58
Divide cube into
8 cubes
62

Fig. 1C

(1 Microfiche, 14 Pages)

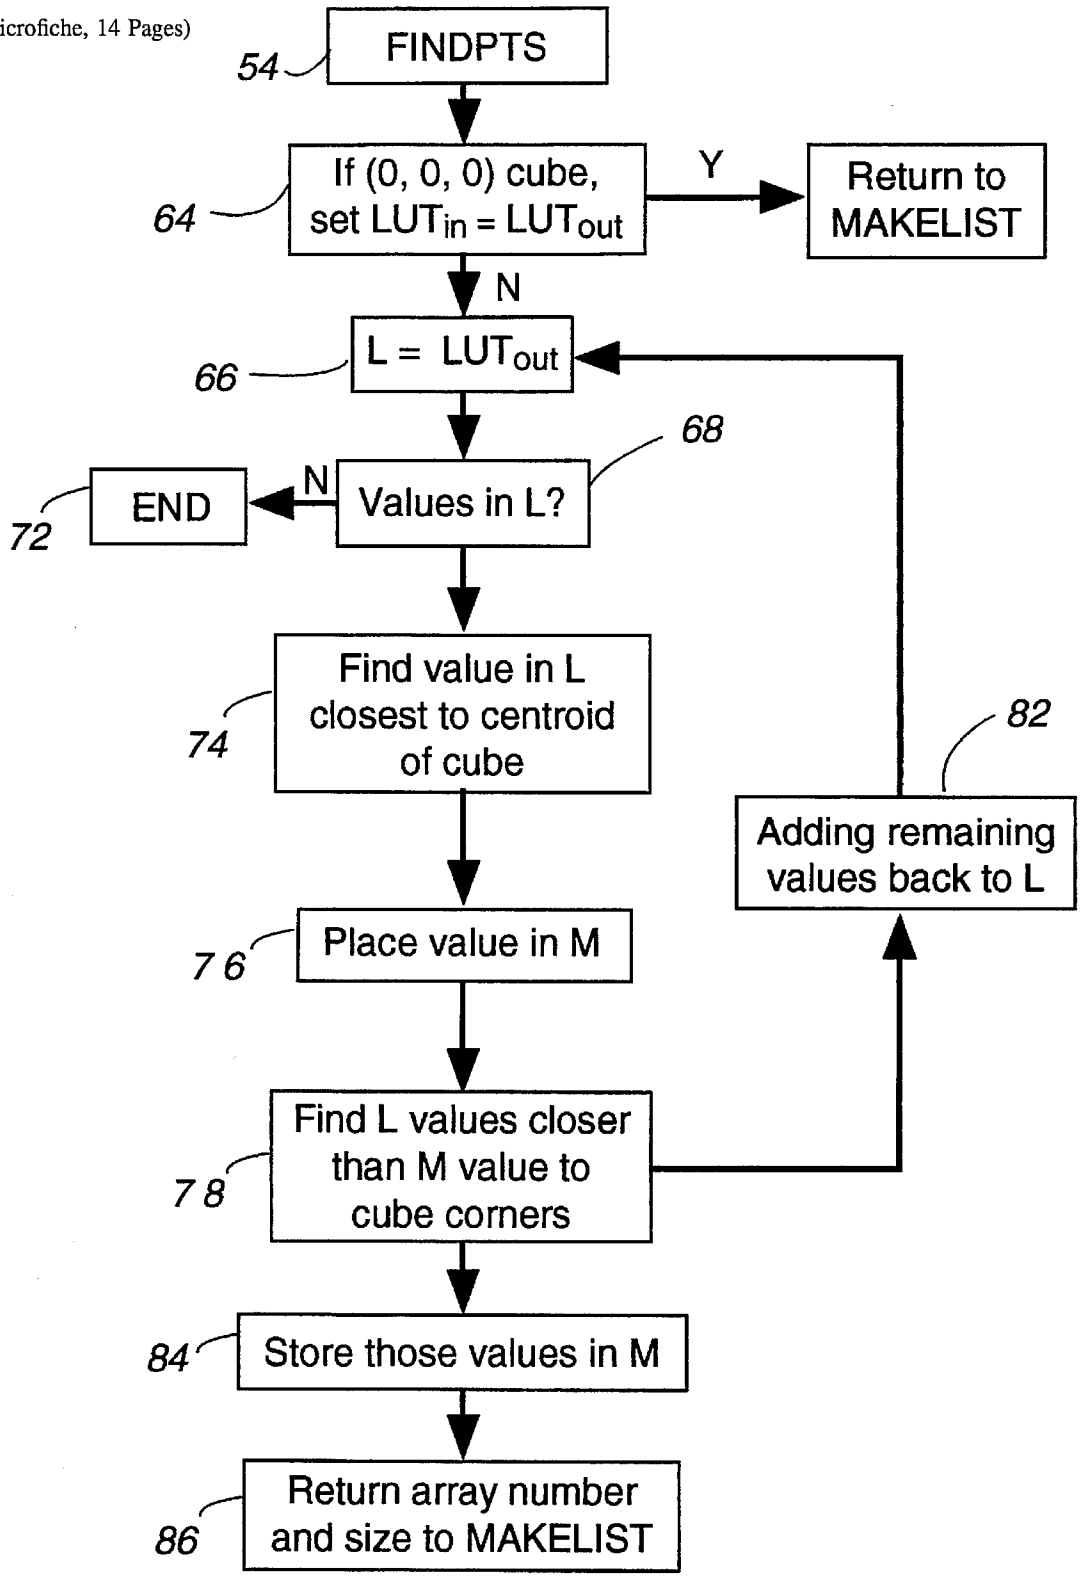

Fig. 1D

POPULATION ATTRIBUTE COMPRESSION

RELATED CASES

This application is a continuation-in-part application under 35 U.S.C. 120 entitled to the benefit of the filing date for patent application Ser. No. 07/350,675, filed May 12, 1989, now U.S. Pat. No. 5,130,701.

MICROFICHE APPENDIX

A microfiche appendix forms a part of the following description, having 1 microfiche with 15 frames. The microfiche appendix contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

BACKGROUND OF INVENTION

This invention relates to vector quantization and, more particularly, to methods for forming a representation of a population from a look-up table having a predetermined number of attributes representing a larger number of attributes of the defined population. This invention is the result of a contract with the Department of Energy (Contract No. W-7405-ENG-36).

The field of vector quantization generally concerns the representation of a large number of specified attributes of a given population with a smaller number of attributes which are distributed over the population to approximate the distribution of the large number of attributes. In the following discussion, the population is pixels of a video display system and the attributes are colors to be associated with the pixels. It will be understood that the processes described herein are applicable to any population which can be described by specified attributes such that the term pixel means any specified population and the term color means any selected set of attributes.

A display system for a video system, e.g. a data processing system, displays discrete colors at individual pixels. The color represented at each pixel is typically formed from a plurality of phosphors, each generating a specific color amplitude response to an activating electron beam. A typical set of phosphors may represent the primary colors red, blue, and green, from which a complete color spectrum may be formed. High resolution digital representations of color images present twenty-four bit words, eight bits each for the three colors, to encode the color to be presented. This representation provides over 16 million discrete colors (2^{24}).

A twenty-four bit color representation, however, is beyond the capability of most display systems, particularly where a real-time video capability is desired. A conventional display system uses only an eight bit word for color representation, which enables 256 colors (2^8) to be selected to represent an image. The eight bit word does not, however, directly represent a color, but an address in a look-up table (LUT). The LUT then contains twenty-four bit representations at each of the 256 addresses. The display system assigns each actual color to one of the stored colors in the LUT and the stored color is actually used to generate the color displayed for each pixel during the raster generation of a color display. U.S. Pat. No. 4,751,446, issued Jun. 14, 1988, to Pineda et al., incorporated herein by reference, describes one embodiment of a LUT for providing color data to a video display.

To present a high resolution color image, an optimum set of colors must be selected to represent the image. A method that will compress a color image for LUT decompression will have the following features:

1. It must produce from a list of colors in the original image a second smaller list of colors (the representative LUT colors); the representative colors need not be present in the original image.

2. It must replace each of the original colors with an index into the LUT.

In one approach to color image quantization, a fixed set of color representations are stored in the LUT. The fixed set of colors may be uniformly spaced or may be based on a statistical distribution of the input colors. A uniform quantization is computationally fast, but provides poor color representations. In a statistical distribution representation, an algorithm must be selected to map the image colors onto the LUT to adequately represent the distribution of actual image colors.

A number of algorithms have been applied for this color mapping, some of which are discussed in P. Heckbert, "Color Image Quantization for Frame Buffer Display," 16 Computer Graphics, No. 3, pp. 297–302 (July 1982), incorporated herein by reference. In one algorithm, the densest regions in the image color distribution are selected to form the LUT. This algorithm apparently does not perform well on images with a wide variety of colors or with a small number of colors. Other algorithms attempt to define volumes in the color space, i.e., Voronoi regions, and select colors from those volumes, either as volume averages or as centroids of the volumes. Substantial computing time is required to define and iterate the Voronoi volumes.

Once a LUT is formed to represent the image colors, the image colors must then be mapped onto the LUT, i.e., each image color pixel must be replaced with the address of the representative color value in the LUT. This operation is a "nearest neighbor" problem. A straightforward approach is to test all LUT values for each image color and choose the closest LUT value. This approach requires substantial computation time.

Heckbert teaches a process for a locally sorted search. Color space is divided into equal volumes each containing a sorted list of LUT values. The list includes all LUT values that are the nearest neighbors of some point in that volume. Then any image color in the volume needs to be compared only with the LUT values in the sorted list.

However Heckbert provides only uniform volume sizes that may contain none, or a large number, of LUT values to be searched against image pixel values. Further, the list of values associated with each volume is not minimized.

Yet another search technique is described in J. H. Friedman et al., "An Algorithm for Finding Best Matches in Logarithmic Expected Time," 3 ACM Trans. Math. Software, pp. 209–226 (September 1977). A file (LUT) containing N records, each described by K real valued keys (color values) is searched for the closest match or nearest neighbor to a given query record (image color). The K-D tree structure described by Friedman does not adequately sort the data when forming the tree and additional sorting must be done when accessing the record space with a given key. Further, the K-D structure may not reach a terminal value on a single traverse of the tree and several traverses of the tree may be required, with a concomitant large storage space overhead required to maintain priority queues during the search.

These and other problems of the prior art are addressed by the present invention wherein an image reconstruction process sorts LUT values to a tree structure that allows a rapid traverse and a minimum nearest neighbor query at each terminal.

Accordingly, one object of the present invention is to generate a high resolution color representation of an image by rapidly replacing image pixel values with the addresses of assigned color values in a LUT.

It is another object of the present invention to provide a method for quantizing image colors from a LUT which operates on rapid sequential displays in real time.

Additional objects, advantages and novel features of the invention will be set forth in part in the description which follows, and in part will become apparent to those skilled in the art upon examination of the following or may be learned by practice of the invention. The objects and advantages of the invention may be realized and attained by means of the instrumentalities and combinations particularly pointed out in the appended claims.

SUMMARY OF INVENTION

To achieve the foregoing and other objects, and in accordance with the purposes of the present invention, as embodied and broadly described herein, the method of this invention may comprise associating selected attributes in a given population with a smaller number of attributes in a look-up table (LUT) to form a reconstructed population using the LUT attributes. The attribute space containing the LUT attributes is refined into a plurality of adaptive volumes in the attribute space wherein each of the adaptive volumes has a predetermined maximum number of associated LUT attributes that are closest to interior points of the volume. A tree structure is then formed from the adaptive volumes in the attribute space. A node of the tree corresponds to a volume in the attribute space and each node either points to smaller volumes formed from the node volume or is a leaf of the tree, i.e., one of the adaptive volumes. Each population attribute is then traversed along the tree to arrive at a leaf containing said population attribute and the associated LUT attributes. The distances between the population attribute and only a small number of LUT attributes need to be determined to locate the closest LUT attribute to represent the population attribute. The tree structure is formed only once for each population and the association of each population attribute with a LUT attribute requires only a relatively small computation time.

In a particular embodiment of the invention the population is the pixels forming an image with color attributes, e.g., red, green, and blue. The image colors are compressed into the LUT color values for subsequent image reconstruction. Image colors are associated with LUT colors by forming a tree structure and traversing the tree to find a leaf with a maximum number of the closed LUT values to the population color value being examined. The leaf color values are then examined to determine the closest LUT color value to represent the population color value. A twenty-four bit color value is then replaced with the eight bit address of the closest LUT color value.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and form a part of the specification, illustrate an embodiment of the present invention and, together with the description, serve to explain the principles of the invention. In the drawing:

FIG. **1A** is a block flow diagram for forming a LUT suitable to display a selected color image.

FIG. **1B** is a block flow diagram for reconstructing a color image from the LUT formed in FIG. **1A**.

FIG. **1C** is a subroutine for FIG. **1B** for defining cubes in color space.

FIG. **1D** is a subroutine for FIG. **1B** for determining the nearest LUT values to the interior of the cubes defined in FIG. **1C**.

DETAILED DESCRIPTION OF THE INVENTION

In accordance with one embodiment of the present invention, a rapid association is performed of image pixel colors with a limited color palette stored in a look-up table (LUT). The color space containing the LUT color palette is divided into a plurality of adaptive volumes, each containing at most a predetermined maximum number of LUT color values. The divided color space is then organized in a tree structure with nodes corresponding to the successively divided volumes. Each node either defines a terminal volume, a leaf, with a limited number of LUT values to be examined, or contains the addresses of smaller included volumes. At each node the image pixel color value is either associated with a closest LUT color value or is referred to the next closest included node.

In a preferred method for forming a LUT, Monte Carlo sampling techniques are used to directly form the LUT for image color representation without any need to directly compute Voronoi regions, i.e. regions enclosing image colors which are nearer to a given LUT color representation than any other LUT color representation. The following steps are required to form the LUT:

1. Pick 256 unique colors $c_0=(c_1,c_2,\ldots c_{256})$ with uniform probability from the original image.

2. Pick one color, u, from the original image with uniform probability.

3. Update the color c_i for which the distance between u and c_i, i.e., $|u-c_i|$, is minimum by computing a weighted average, e.g. $c_i1=(n/n+1)c_i+(1/n+1)u$, where n is the number of times that c_i has been previously updated plus one.

4. Return to 2. The image sampling and LUT updating are continued until some selected criterion is obtained, e.g. a predetermined number of samples have been taken or the LUT color representations remain within predetermined limits for successive updates.

Once the LUT is formed, each color in the image is indexed to its closest c_i. It is believed that these assignments may be equivalent to forming the Voronoi regions of the prior art but without explicitly computing the regions, which is a very time consuming computational step.

In order to index the image colors into the LUT color representations, a nearest LUT color must be found for each actual image color. The problem can be expressed as follows: Given a set of color arrays

r_i,g_i,b_i $0\leqq i<N_p$

called the pixel set, and

r'_j,g'_j,b'_j $0\leqq j<N_t>256$

called the look-up table, find a new color array

c_i, $0\leqq i<N_p$

such that for each i

$c_i=j_m$,

where r'_{jm},g'_{jm},b'_{jm} satisfies

$\min D(r'_j,g'_j,b'_j,r_i,g_i,b_i)$ for $0 \leqq j < N_t$

and

$D(r'_j,g'_j,b'_j,r_i,g_i,b_i) = (r'_j-r_i)^3+(g'_j-g_i)^2+(b'_j-b_i)^2$

For the video application described, $r'_j,g'_j,b'_j,r_i,g_i,b_i$ and c_i are all integers in the interval [0,255], but the technique does not depend on this restriction.

In accordance with the present invention, this replacement of actual image values with an index to the closest LUT values is done using an adaptive algorithm rather than a one-by-one computation. It will be appreciated that both the look-up table and the pixel set are contained in a cube with each side having length **256**. If this cube is subdivided into eight equal cubes, the LUT values nearest every interior point in each of the refined cubes can now be determined.

The cube volumes are adaptively formed rather than uniformly sized. The cubes are adaptively refined to define a number of cubes each having no more than a predetermined maximum number of LUT values associated with the cube. The pixel values in the interior of each cube are now compared with only the reduced number of LUT values most closely associated with that cube rather than the entire LUT. A tree structure that is easily traversed is formed from the cubes.

As hereinafter described, the LUT values associated with each cube volume are the smallest subset of points in the LUT closer to every point of the cube than the complement of this subset. First, the LUT is copied into a temporary list, L. Then a value in the list, L, is found closest to the centroid of a selected subcube, $p_j=(r'_j,g'_j,b'_j)$, and stored in a list, M, as one member of a list of points nearest to the subcube. It is then determined for each of the remaining points, P_k, in the list, L, whether the selected subcube all lies within the half plane formed by a plane perpendicular to the midpoint of a line joining p_j and p_k and nearest to p_j, i.e., to determine those points, p_k, nearer the cube corners than point p_j. If so, p_k is added to the list, M, for further consideration. The list M is the desired list of points for the selected subcube, i.e. a new, smaller LUT for the image colors located in the subcube.

Referring now to FIG. **1A**, there is shown a flow diagram for a function labeled CONSTRUCT **10** which takes a data structure describing a twenty-four bit image and returns a look-up table containing twenty-four bit color representation values effective to form a high resolution color image from an eight bit index representation. CONSTRUCT **10** takes a data structure describing a twenty-four bit image and returns a look-up table containing color values that can best be used to reconstruct a color image and which are addressable by an eight bit pointer. The twenty-four bit image is described with the first four arguments of the function data structure. The first three variables point to arrays that are the intensities for the red, green, and blue components of the twenty-four bit image. These arrays are unsigned characters with a range from 0 to 255 inclusive. The fourth argument of the data structure is an integer giving the number of pixels in the image.

The next four arguments of the data structure constitute the description of the look-up table. The first three arguments point to arrays that are unsigned characters, where the arrays contain red, green, and blue intensities in the range of 0 to 255. The fourth argument is an integer giving the number of desired entries in the look-up table. The second to last entry is a pointer to an array of integers having the same size as the look-up table arrays. The last entry in the data structure is an integer specifying the number of iterations the algorithm should use in finding the look-up table. Upon completion of the function call, the function value returned is the number of distinct elements in the look-up table. This number is generally the number specified in the function call, but may be a smaller number if fewer colors are found after the specified number of samples.

When CONSTRUCT **10** is called, the existing LUT entries are set **12** to zero. The desired color image is digitized and is sampled **14** with uniform probability to select an initial set of 256 unique colors which form **16** the LUT. An iteration loop is now established to better represent the actual image colors in the LUT. The image colors are sampled **18** with equal probability to obtain a series of color samples, u. For each color, u, the closest LUT representation is found and updated **22** to form a replacement LUT representation. The updated value is an average of the LUT value and the sampled image value with the sample value weighted as one divided by the number of samples plus one. In one embodiment, the LUT value is weighted as one minus the sample weight. The iteration loop continues until a selected number of samples **24** have been taken. It will be appreciated that later samples will have progressively less effect on the stored LUT values.

Once the selected number of samples **24** has been taken, the LUT values are set **26** and PROJECT **30** is called for replacing the image colors with pointers into the LUT. Referring now to FIG. **1B**, there is shown a flow chart for function call PROJECT **30**, which builds an eight bit image from a twenty-four bit image and a look-up table. For each pixel in the twenty-four bit image, PROJECT **30** finds the closest value in the look-up table and stores the index of that value in the output array. The twenty-four bit image is described with the first four arguments in the data structure of PROJECT **30**. The first three variables point to arrays that are the intensities for the red, green, and blue components of the twenty-four bit image. These arrays are unsigned characters with a range from 0 to 255 inclusive. The fourth argument of the function data structure is an integer giving the number of pixels in the image.

The next four arguments of the function constitute a data structure for the description of the look-up table. The first three arguments point to arrays that are unsigned characters and contain the red, green, and blue intensities in the range of 0 to 255. The fourth argument is an integer giving the number of entries in the look-up table. The last argument of the function call points to an array used to describe the generated eight bit image. This array is a collection of unsigned characters used as pointers from a pixel location to an entry in the look-up table. To find the eight bit representation of a pixel at the n'th location in the pixel image, the number in the n'th location of the pointer array is found and that pointer number is used as the color index to the LUT.

In order to efficiently construct the eight bit pixel image, PROJECT **30** creates a tree structure from the LUT to quickly find the nearest entry in the LUT for a given pixel value from the twenty-four bit image. Each node of the structure corresponds to a cube in red-green-blue space. A data structure is associated with each cube and contains four short integers, an array of eight pointers pointing to other cubes and another pointer pointing to an array of indices if the cube is also a leaf of the tree. The first three integers specify the absolute origin of the cube and the fourth integer specifies whether the cube is a leaf of the tree or provides pointers to eight other cubes.

PROJECT **30** first defines **32** a cube in color space having its origin at (0,0,0). The function MAKELIST **34** (see FIG. **2C**) is called to recursively build the tree structure of cubes which is later traversed by PROJECT **30** to find the entry in

the LUT nearest to a given pixel. MAKELIST **30**. examines each cube to determine whether the cube should be further refined, i.e. subdivided, by calling itself, or whether the LUT values within the cube should be stored and another cube examined.

Once MAKELIST **34** has defined a tree, PROJECT **30** selects **36** a pixel from a location in the pixel array and the tree is traversed **38** until a leaf is reached **42**, i.e. a cube which contains the pixel and a predetermined maximum number of LUT color values. The pixel color value is then compared **44** only with the leaf array to find the closest array value. The pixel image color value is replaced **46** by the pointer argument associated with the array value. It will be appreciated that this comparison between pixel image colors and LUT colors involves only a few LUT colors which have been associated with the leaf cube. The pixel selection loop is repeated **48** until all of the twenty-four bit image values have been replaced with eight bit pointers to the LUT.

The function MAKELIST **34** refines the cube structure in the manner shown in FIG. **1C** so that each defined cube has no more than a predetermined maximum number of LUT values nearest to interior points of the cube. As each cube is presented to MAKELIST **34**, the function FINDPTS **54** (see FIG. **1D**) is called to determine the number of LUT values nearest to interior points of the cube. If the number of LUT values is determined **56** to be less than the selected maximum the array of color values found by FINDPTS **54** is stored for that cube. If the number of LUT values is greater than the maximum, the cube is refined, i.e. subdivided, **62** into eight smaller cubes and MAKELIST **34** is called recursively for each cube until the tree structure is defined.

Referring now to FIG. **1D**, the function FINDPTS **54** is called to find the LUT values which are closer to the defined cube than any other values and returns an array with those points. If the cube to be evaluated is the initial cube **64**, the program returns to MAKELIST **34** for the cube to be refined. Otherwise, an initial list L is established **66** having the values from the original LUT. For each cube from MAKELIST **34** a value in L is found **74** which is nearest the centroid of the cube and this value is placed **76** in a list M. The list L is then examined **78** for values which are closer to the cube corners than the centroid approximation stored in M. All of the closer values which are found in L are placed **84** in array M for storage. The remaining values are placed **82** back in list L for examination. The array M number and size is returned to MAKELIST **34** to determine whether array M is stored in the tree or is further refined.

Thus, function call CONSTRUCT **10** defines a set of LUT values to represent the actual image colors. Using function calls FINDPTS **54** and MAKELIST **34**, a tree structure is defined to locate the LUT values in the image color space. Then, function call PROJECT **30** traverses the tree structure with each image pixel value to find the closest LUT value. Only a single traverse of the tree structure is needed for each pixel to find its leaf. PROJECT **30** assigns a pointer **46** described by an 8 bit number to relate each image pixel value to a LUT value. A display of the image is then created using the LUT values by converting each image pixel color value in a conventional manner to a display pixel color value through the assigned 8 bit pointers.

For further analysis of the resulting display, no further pixel manipulation is needed, since the above sequence has obtained an accurate compression of the image pixel colors into the LUT colors. However, the resulting visual representation can have a contoured appearance arising from the assignment of some range of color values.

If it is desired to improve the aesthetic appearance of the display, a spatial integration algorithm, or dithering, can be introduced in PROJECT **30**. A true spatial integration can provide two desirable improvements in the resulting display. First, the technique globally averages the image attributes whereby the display has overall average attributes which match the image. Secondly, local anomalies are reduced. The integration provides several LUT color values in a local area whose average more closely approximates the image color value than a single LUT value. Further, in a region where the LUT value changes, i.e., a contour line, the integration mixes the adjacent LUT values to smooth out the apparent contour line. In one embodiment, the pointer array (FIG. **1B**, PROJECT **30**, steps **36–48**) is formed by including additional steps to find the error between an image pixel value and its corresponding LUT value, and incrementing the next image pixel value by the error before finding the corresponding LUT value.

A program listing to accomplish the flow diagram shown in FIGS. **1A–1D** is depicted in the attached microfiche appendix. The program includes a subroutine for dithering during the process for forming the pointer array to compress the image attributes to the LUT attributes.

While the above process has generally been described in terms of the color attributes (red, green, and blue) of a pixel population, it has obvious application to any set of attributes selected to represent a certain population. For example, a high resolution black and white image can be simulated on low resolution video monitors. Further, the construction of the look-up table is sufficiently rapid to support a real time video display as the speed of the image correlation algorithm is improved.

The foregoing description of the preferred embodiment of the invention has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed, and obviously many modifications and variations are possible in light of the above teaching. The embodiment was chosen and described in order to best explain the principles of the invention and its practical application to thereby enable others skilled in the art to best utilize the invention in various embodiments and with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the claims appended hereto.

What is claimed is:

1. A method for associating selected attributes in a given population with a smaller number of attributes in a look-up table (LUT) to form a reconstructed population using said LUT attributes, comprising the steps of:

refining an attribute space containing said LUT attributes into a plurality of adaptive cube volumes in said attribute space wherein each said adaptive cube volume has a predetermined maximum number of associated ones of said LUT attributes that are closest to interior points of said cube volume;

forming a tree structure defining a plurality of nodes in said attribute space where each node of said tree is a corresponding cube volume in said attribute space and either points to smaller cube volumes formed from said corresponding cube volume or to said associated LUT attributes when said corresponding cube volume is a leaf of said tree consisting of one of said adaptive cube volumes;

traversing said tree with each of said selected population attributes to arrive at one said leaf containing said selected population attribute and said associated LUT attributes; and

determining the closest one of said associated LUT attributes in said leaf to said population attribute to

replace said population attribute.

2. A method according to claim 1 wherein the step of refining said attribute space further comprises the steps of:

dividing said attribute space into first cube volumes;

determining for each said first cube volume one of said LUT attributes closest to the centroid of said each first cube volume;

determining the set of remaining LUT attributes closer to corners of said first cube volume than said LUT attribute closest to said centroid; and

further dividing said first volume into second volumes when the number of attributes in said set exceeds said predetermined maximum of attributes.

3. A method according to claim 2, wherein the step of traversing said tree comprises the steps of:

locating a said first volume closest to a selected population attribute;

determining whether said first volume is a leaf for said selected population attribute;

if said first volume is not a leaf, locating a said second volume closest to said selected population attribute; and

repeatedly selecting closest successive volumes until a leaf is reached.

4. A method for associating image pixel color values with a selected number of color values in a look-up table (LUT) to form a reconstructed color image using said LUT color values, comprising the steps of:

refining a color space containing said LUT color values into a plurality of adaptive volumes in said color space wherein each said adaptive volume has a predetermined maximum number of associated ones of said LUT color values that are closest to interior points of said volume;

forming a tree structure defining a plurality of nodes in said color space where each node of said tree is a corresponding volume in said color space and either points to smaller volumes formed from said corresponding volume or to said associated LUT values when said corresponding volume is a leaf of said tree consisting of one of said adaptive volumes;

traversing said tree with each said image pixel color value to arrive at one said leaf containing said pixel value and said associated LUT color values; and

determining the closest one of said associated LUT color values in said leaf to said image pixel color contained in said leaf to replace said image pixel color.

5. A method according to claim 4 wherein the step of refining said color space further comprises the steps of:

dividing said color space into first volumes;

determining for each said first volume one of said LUT color values closest to the centroid of said each first volume;

determininq the set of remaining LUT attributes closer to corners of said first cube volume than said LUT attribute closest to said centroid; and

further dividing said first volume into second volumes when the number of color values in said list exceeds said predetermined maximum number of color values.

6. A method according to claim 5, wherein the step of traversing said tree comprises the steps of:

locating a said first volume closest to a selected image pixel color;

determining whether said first volume is a leaf for said selected image pixel color;

if said first volume is not a leaf, locating a said second volume closest to said selected image pixel color; and

repeatedly selecting closest successive volumes until a leaf is reached.

* * * * *