

(19)



(11)

EP 2 263 191 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention of the grant of the patent:
17.01.2018 Bulletin 2018/03

(51) Int Cl.:
G06T 7/187 (2017.01)

(21) Application number: **09724041.0**

(86) International application number:
PCT/US2009/001526

(22) Date of filing: **10.03.2009**

(87) International publication number:
WO 2009/120264 (01.10.2009 Gazette 2009/40)

(54) SYSTEM AND METHOD FOR IDENTIFYING COMPLEX TOKENS IN AN IMAGE

SYSTEM UND VERFAHREN ZUR IDENTIFIZIERUNG KOMPLEXER TOKENS IN EINEM BILD

SYSTÈME ET PROCÉDÉ POUR IDENTIFIER DES JETONS COMPLEXES DANS UNE IMAGE

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO SE SI SK TR

(74) Representative: **Reichert & Lindner
Partnerschaft Patentanwälte
Bismarckplatz 8
93047 Regensburg (DE)**

(30) Priority: **28.03.2008 US 79877**

(56) References cited:
**US-A- 5 933 518 US-A1- 2006 074 835
US-A1- 2007 177 797 US-A1- 2007 177 797**

(43) Date of publication of application:
22.12.2010 Bulletin 2010/51

(73) Proprietor: **Tandent Vision Science, Inc.
San Francisco, California 94111 (US)**

- **AMADASUN M ET AL: "LOW-LEVEL SEGMENTATION OF MULTISPECTRAL IMAGES VIA AGGLOMERATIVE CLUSTERING OF UNIFORM NEIGHBOURHOODS", PATTERN RECOGNITION, ELSEVIER, GB, vol. 21, no. 3, 1 January 1988 (1988-01-01), pages 261-268, XP000021175, ISSN: 0031-3203, DOI: 10.1016/0031-3203(88)90060-X**
- **GHAZAL M ET AL: "A Real-Time Technique for Spatio-Temporal Video Noise Estimation", IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, IEEE SERVICE CENTER, PISCATAWAY, NJ, US, vol. 17, no. 12, 1 December 2007 (2007-12-01), pages 1690-1699, XP011195143, ISSN: 1051-8215, DOI: 10.1109/TCSVT.2007.903805**

(72) Inventors:
• **FRIEDHOFF, Richard, Mark
New York, NY 10019 (US)**
• **SMITH, Casey, Arthur
Grand Junction
CO 81505 (US)**
• **MAXWELL, Bruce, Allen
Benton
ME 04901-3634 (US)**

EP 2 263 191 B1

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

DescriptionBackground of the Invention

5 **[0001]** A challenge in the utilization of computers to accurately and correctly perform operations relating to images is the development of algorithms that truly reflect and represent physical phenomena occurring in the visual world. For example, the ability of a computer to correctly and accurately distinguish between a shadow and a material object edge within an image has been a persistent challenge to scientists. An early and conventional approach to object edge detection involves an analysis of brightness boundaries in an image. In the analysis it is assumed that a boundary caused by a material object will be sharp, while a boundary caused by a shadow will be soft or gradual due to the penumbra effect of shadows. While this approach can be implemented by algorithms that can be accurately executed by a computer, the results will often be incorrect. In the real world there are many instances wherein shadows form sharp boundaries, and conversely, material object edges form soft boundaries. Thus, when utilizing conventional techniques for shadow and object edge recognition, there are significant possibilities for false positives and false negatives for shadow recognition. That is, for example, a material edge that imitates a shadow and is thus identified incorrectly by a computer as a shadow or a sharp shadow boundary that is incorrectly interpreted as an object boundary. Accordingly, there is a persistent need for the development of accurate and correct techniques that can be utilized in the operation of computers relating to images, to, for example, identify material characteristics of the image.

10 **[0002]** The publication US 2007/0177797 A1 of the US patent application 11/341,753 discloses methods for identifying regions of constant reflectance in an image. With some of these methods, for each pixel in the image it is checked whether the pixel and its neighboring pixels fall within a cylinder in RGB color space from a dark end to a bright end along a positive slope. Such sets of pixels are referred to as linear tokens. Clustering and merging of similar tokens is considered, wherein similarity of tokens is defined via Euclidean distances in RGB color space.

15 **[0003]** The article "Low-Level Segmentation of Multispectral Images via Agglomerative Clustering of Uniform Neighbourhoods" by Amadasun M et al., Pattern Recognition, Elsevier, GB, vol. 21, no. 3, 1 January 1988, pages 261-268, discusses a method where an image is initially divided into a number of non-overlapping neighbourhoods. These neighbourhoods are subsequently tested for uniformity. Mean feature vectors are computed for all neighbourhoods considered to be uniform. In case the number of mean feature vectors obtained is greater than an initially set number of categories into which the image is to be partitioned, the two most similar mean vectors are merged.

Summary of the Invention

20 **[0004]** The present invention provides a method and system comprising image techniques that accurately and correctly identify regions of an image that correspond to a single material in a scene depicted in the image.

25 **[0005]** In a first exemplary embodiment of the present invention, an automated, computerized method is provided for processing an image, as defined in claim 1.

30 **[0006]** In a second exemplary embodiment of the present invention, a computer system is provided. The computer system comprises a CPU and a memory storing an image file. Pursuant to a feature of the present invention, the CPU is arranged and configured to execute a routine to perform the method of claim 1.

35 **[0007]** In accordance with yet further embodiments of the present invention, computer systems are provided, which include one or more computers configured (e.g., programmed) to perform the methods described above. In accordance with other embodiments of the present invention, computer readable media are provided which have stored thereon computer executable process steps operable to control a computer(s) to implement the embodiments described above. The automated, computerized methods can be performed by a digital computer, analog computer, optical sensor, state machine, sequencer or any device or apparatus that can be designed or programed to carry out the steps of the methods of the present invention.

Brief Description of the Drawings**[0008]**

40 Figure 1 is a block diagram of a computer system arranged and configured to perform operations related to images.

45 Figure 2 shows an $n \times m$ pixel array image file for an image stored in the computer system of figure 1.

50 Figure 3a is a flow chart for identifying Type C token regions in the image file of figure 2, according to a feature of the present invention.

Figure 3b is an original image used as an example in the identification of Type C tokens.

Figure 3c shows Type C token regions in the image of figure 3b.

5 Figure 3d shows Type B tokens, generated from the Type C tokens of figure 3c, according to a feature of the present invention.

Figure 4 is a flow chart for a routine to test Type C tokens identified by the routine of the flow chart of figure 3a, according to a feature of the present invention.

10 Figure 5 is a flow chart for constructing Type B tokens via an arbitrary boundary removal technique, according to a feature of the present invention.

Figure 6 is a flow chart for creating a token graph, according to a feature of the present invention.

Figure 7 is a flow chart for constructing Type B tokens via an adjacent planar token merging technique, according to a feature of the present invention.

20 Figure 8 is a flow chart for generating Type C tokens via a local token analysis technique, according to a feature of the present invention.

Figure 9 is a flow chart for constructing Type B tokens from Type C tokens generated via the local token analysis technique of figure 8, according to a feature of the present invention.

25 Figure 10a shows intensity profiles and corresponding first derivative and second derivative profiles for each of a material edge, a single material with a shadow over a portion of the material and a single material under constant illumination.

Figure 10b is an example of an LoG filter mask.

30 Figure 10c is an illustration of a 7 X 7 pixel patch within the n X m pixel array image file of figure 2.

Figure 11 is a flow chart for identifying regions of uniform material reflectance within the n X m pixel array image file of figure 2, according to a feature of the present invention.

35 Figure 12 is a flow chart for combining results from executions of the routine of figure 11 with differing parameters, according to a feature of the present invention.

Detailed Description of the Preferred Embodiments

40 **[0009]** Referring now to the drawings, and initially to figure 1, there is shown a block diagram of a computer system 10 arranged and configured to perform operations related to images. A CPU 12 is coupled to a device such as, for example, a digital camera 14 via, for example, a USB port. The digital camera 14 operates to download images stored locally on the camera 14, to the CPU 12. The CPU 12 stores the downloaded images in a memory 16 as image files 18. The image files 18 can be accessed by the CPU 12 for display on a monitor 20, or for print out on a printer 22.

45 **[0010]** Alternatively, the CPU 12 can be implemented as a microprocessor embedded in a device such as, for example, the digital camera 14 or a robot. The CPU 12 can also be equipped with a real time operating system for real time operations related to images, in connection with, for example, a robotic operation or an interactive operation with a user.

50 **[0011]** As shown in figure 2, each image file 18 comprises an n X m pixel array. Each pixel, p, is a picture element corresponding to a discrete portion of the overall image. All of the pixels together define the image represented by the image file 18. Each pixel comprises a digital value corresponding to a set of color bands, for example, red, green and blue color components (RGB) of the picture element. The present invention is applicable to any multi-band image, where each band corresponds to a piece of the electro-magnetic spectrum. The pixel array includes n rows of m columns each, starting with the pixel p(1, 1) and ending with the pixel p(n, m). When displaying or printing an image, the CPU 12 retrieves the corresponding image file 18 from the memory 16, and operates the monitor 20 or printer 22, as the case may be, as a function of the digital values of the pixels in the image file 18, as is generally known.

55 **[0012]** In an image operation, the CPU 12 operates to analyze the RGB values of the pixels of a stored image file 18 to achieve various objectives, such as, for example, to identify regions of an image that correspond to a single material

depicted in a scene recorded in the image file 18. A fundamental observation underlying a basic discovery of the present invention, is that an image comprises two components, material and illumination. All changes in an image are caused by one or the other of these components. A method for detecting of one of these components, for example, material, provides a mechanism for distinguishing material or object geometry, such as object edges, from illumination and shadow boundaries.

[0013] Pursuant to a feature of the present invention, a token is a connected region of an image wherein the pixels of the region are related to one another in a manner relevant to identification of image features and characteristics such as identification of materials and illumination. The pixels of a token can be related in terms of either homogeneous factors, such as, for example, close correlation of color among the pixels, or inhomogeneous factors, such as, for example, differing color values related geometrically in a color space such as RGB space, commonly referred to as a texture. Exemplary embodiments of the present invention provide methods and systems to identify various types of homogeneous or inhomogeneous tokens for improved processing of image files. The present invention utilizes spatio-spectral information relevant to contiguous pixels of an image depicted in an image file 18 to identify token regions. The spatio-spectral information includes spectral relationships among contiguous pixels, in terms of color bands, for example the RGB values of the pixels, and the spatial extent of the pixel spectral characteristics relevant to a single material.

[0014] According to one exemplary embodiment of the present invention, tokens are each classified as either a Type A token, a Type B token or a Type C token. A Type A token is a connected image region comprising contiguous pixels that represent the largest possible region of the image encompassing a single material in the scene. A Type B token is a connected image region comprising contiguous pixels that represent a region of the image encompassing a single material in the scene, though not necessarily the maximal region corresponding to that material. A Type C token comprises a connected image region of similar image properties among the contiguous pixels of the token, where similarity is defined with respect to a noise model for the imaging system used to record the image.

[0015] Referring now to figure 3a, there is shown a flow chart for identifying Type C token regions in the scene depicted in the image file 18 of figure 2, according to a feature of the present invention. Type C tokens can be readily identified in an image, utilizing the steps of figure 3a, and then analyzed and processed to construct Type B tokens.

[0016] Prior to execution of the routine of figure 3a, the CPU 12 can operate to filter the image depicted in a subject image file 18. The filters may include a texture filter, to, for example, transform patterns of differing reflectance caused by a textured material into a homogeneous representation that captures the spectral and spatial characteristics of the textured region. Identification of tokens can be difficult in a textured image. A textured image contains materials with, for example, more than one reflectance function that manifests as a defining characteristic. For example, the defining characteristic can be a pattern of colors within the texture, such that the texture displays a certain distribution of colors in any patch or region selected from anywhere within the textured region of the image.

[0017] Other textures can be defined by geometric characteristics, such as stripes or spots. The CPU 12 can execute a software module that implements any well known method, such as, for example, a Laws filter bank, wavelets or textons (see, for example, Randen, T.[Trygve], Husøy, J.H.[John Håkon], Filtering for Texture Classification: A Comparative Study, PAMI(21), No. 4, April 1999, pp. 291-310.), or convert a local area around each pixel to an histogram. Any method utilized will convert each pixel value of N color bands to a vector of T values representing the output of one or more functions applied to a local area around the pixel, for example, an 11x11 pixel array.

[0018] For example, an histogram representation for each pixel can be produced using the following algorithm:

Loop over all pixels p in an N-band (for example, RGB) input color image;

- A) Initialize N 8-bin histograms to zero, one for each color band
- B) For each pixel q within a neighborhood of p (for example, an 11x11 pixel box)

- (i) For each of the N color values C_n of q;

- (a) Increment the appropriate bins of the nth histogram;
- (b) Use interpolation so that the two bins closest to the color value get incremented proportionally;

- (ii) Concatenate the N 8-bin histogram values together into a single $8 \times N$ element vector;
- (iii) Assign the $8 \times N$ element vector to the corresponding pixel p in the output image.

After the transformation from a set of color bands to a set of filter outputs, the image is treated exactly as the original color band image with respect to identifying type C tokens.

[0019] In many instances, the texture filters may only be required on part of an input image, as much of the image may include homogeneously colored objects. Therefore, prior to application of the texture filters, it is useful to identify

EP 2 263 191 B1

and mask off regions of homogeneous color. The texture filters are then only applied to areas where there appear to be textured materials. An example algorithm for identifying textured regions is as follows:

- 5 1) Execute a type C tokenization on the N-band color values (e.g. RGB), storing the token results in a region map R, where each pixel in the region map has the tokenID of the token to which it belongs.
- 2) Execute a median filter on the region map R (e.g. each pixel P_{ij} is replaced by the median token ID of a 7x7 box around P_{ij}). Store the result in R-median.
- 10 3) Execute a filter on the original image that calculates the standard deviation of the pixels in a box around each pixel (e.g. 7x7) for each color band. Put the result in S.
- 4) For each pixel in S, divide the standard deviation calculated for each color band by an estimated noise model value. An example noise model is $S_n = A * \text{maxValue} + B * \text{pixelValue}$, where maxValue is the maximum possible color band value, pixelValue is the intensity of a particular band, and A and B are constants experimentally determined for the imaging system (e.g. A = 0.001 and B = 0.06 are typical). This step converts the standard deviation into a normalized deviation for each color band. Store the results in Sn.
- 15 5) For each pixel in Sn, sum the squares of the normalized deviations for all N color bands, take the square root of the result and divide by the number of bands N to create a deviation value D_{ij}. Compare the resulting deviation value D_{ij} to a threshold (e.g. 1.0) assign a 1 to any pixel with a deviation value higher than the threshold, otherwise assign the pixel a 0. Store the results in a texture mask image T.
- 20 6) For each pixel in T, if the texture mask value T_{ij} = 1 and the seed size of the token region with the id given in the median region map R-median_{ij} is less than a threshold (e.g. < 4), label the pixel as a textured pixel. Otherwise, label it as a homogeneous pixel. Store the result in the texture mask Tmask.
- 25

The output of the above algorithm is a mask, Tmask, which is the size of the original image. Pixels of Tmask with a 1 value should be treated as part of an image region corresponding to texture materials and pixels with a value of 0 should be treated as part of an image region corresponding to materials of homogeneous color.

30 **[0020]** A 1st order uniform, homogeneous Type C token comprises a single robust color measurement among contiguous pixels of the image. At the start of the identification routine, the CPU 12 sets up a region map in memory. In step 100, the CPU 12 clears the region map and assigns a region ID, which is initially set at 1. An iteration for the routine, corresponding to a pixel number, is set at i=0, and a number for an N x N pixel array, for use as a seed to determine the token, is set an initial value, N = N_{start}. N_{start} can be any integer > 0, for example it can be set at set at 11 or 15 pixels.

35 **[0021]** At step 102, a seed test is begun. The CPU 12 selects a first pixel, i = (1, 1) for example (see figure 2), the pixel at the upper left corner of a first N x N sample of the image file 18. The pixel is then tested in decision block 104 to determine if the selected pixel is part of a good seed. The test can comprise a comparison of the color value of the selected pixel to the color values of a preselected number of its neighboring pixels as the seed, for example, the N x N array. The color values comparison can be with respect to multiple color band values (RGB in our example) of the pixel or the filter output representation of the pixel, in the event the image was filtered, as described above. If the comparison does not result in approximately equal values (within the noise levels of the recording device) for the pixels in the seed, the CPU 12 increments the value of i (step 106), for example, i = (1, 2), for a next N x N seed sample, and then tests to determine if i = i_{max} (decision block 108).

40 **[0022]** If the pixel value is at imax, a value selected as a threshold for deciding to reduce the seed size for improved results, the seed size, N, is reduced (step 110), for example, from N = 15 to N = 12. In an exemplary embodiment of the present invention, i_{max} can be set at i = (n, m). In this manner, the routine of figure 3a parses the entire image at a first value of N before repeating the routine for a reduced value of N.

45 **[0023]** After reduction of the seed size, the routine returns to step 102, and continues to test for token seeds. An N_{stop} value (for example, N = 2) is also checked in step 110 to determine if the analysis is complete. If the value of N is at N_{stop}, the CPU 12 has completed a survey of the image pixel arrays and exits the routine.

[0024] If the value of i is less than i_{max}, and N is greater than N_{stop}, the routine returns to step 102, and continues to test for token seeds.

50 **[0025]** When a good seed (an N x N array with approximately equal pixel values) is found (block 104), the token is grown from the seed. In step 112, the CPU 12 pushes the pixels from the seed onto a queue. All of the pixels in the queue are marked with the current region ID in the region map. The CPU 12 then inquires as to whether the queue is empty (decision block 114). If the queue is not empty, the routine proceeds to step 116.

55 **[0026]** In step 116, the CPU 12 pops the front pixel off the queue and proceeds to step 118. In step 118, the CPU 12

marks "good" neighbors around the subject pixel, that is neighbors approximately equal in color value to the subject pixel, with the current region ID. All of the marked good neighbors are placed in the region map and also pushed onto the queue. The CPU 12 then returns to the decision block 114. The routine of steps 114, 116, 118 is repeated until the queue is empty. At that time, all of the pixels forming a token in the current region will have been identified and marked

in the region map as a Type C token.

[0027] When the queue is empty, the CPU 12 proceeds to step 120. At step 120, the CPU 12 increments the region ID for use with identification of a next token. The CPU 12 then returns to step 106 to repeat the routine in respect of the new current token region.

[0028] Upon arrival at $N = N_{\text{stop}}$, step 110 of the flow chart of figure 3a, or completion of a region map that coincides with the image, the routine will have completed the token building task. Figure 3b is an original image used as an example in the identification of tokens. The image shows areas of the color blue and the blue in shadow, and of the color teal and the teal in shadow. Figure 3c shows token regions corresponding to the region map, for example, as identified through execution of the routine of figure 3a (Type C tokens), in respect to the image of figure 3b. The token regions are color coded to illustrate the token makeup of the image of figure 3b, including penumbra regions between the full color blue and teal areas of the image and the shadow of the colored areas.

[0029] While each Type C token comprises a region of the image having a single robust color measurement among contiguous pixels of the image, the token may grow across material boundaries. Typically, different materials connect together in one Type C token via a neck region often located on shadow boundaries or in areas with varying illumination crossing different materials with similar hue but different intensities. A neck pixel can be identified by examining characteristics of adjacent pixels. When a pixel has two contiguous pixels on opposite sides that are not within the corresponding token, and two contiguous pixels on opposite sides that are within the corresponding token, the pixel is defined as a neck pixel.

[0030] Figure 4 shows a flow chart for a neck test for Type C tokens. In step 122, the CPU 12 examines each pixel of an identified token to determine whether any of the pixels under examination forms a neck. The routine of figure 4 can be executed as a subroutine directly after a particular token is identified during execution of the routine of figure 3a. All pixels identified as a neck are marked as "ungrowable." In decision block 124, the CPU 12 determines if any of the pixels were marked.

[0031] If no, the CPU 12 exits the routine of figure 4 and returns to the routine of figure 3a (step 126).

[0032] If yes, the CPU 12 proceeds to step 128 and operates to regrow the token from a seed location selected from among the unmarked pixels of the current token, as per the routine of figure 3a, without changing the counts for seed size and region ID. During the regrowth process, the CPU 12 does not include any pixel previously marked as ungrowable. After the token is regrown, the previously marked pixels are unmarked so that other tokens may grow into them.

[0033] Subsequent to the regrowth of the token without the previously marked pixels, the CPU 12 returns to step 122 to test the newly regrown token.

[0034] Neck testing identifies Type C tokens that cross material boundaries, and regrows the identified tokens to provide single material Type C tokens suitable for use in creating Type B tokens. Figure 3d shows Type B tokens generated from the Type C tokens of figure 3c, according to a feature of the present invention. The present invention provides several exemplary techniques of pixel characteristic analysis for constructing Type B tokens from Type C tokens. One exemplary technique involves arbitrary boundary removal. The arbitrary boundary removal technique can be applied to Type C tokens whether they were generated using N color band values (RGB in our example) of the pixel or the filter output representation of the pixel, in the event the image was filtered. Actual boundaries of any particular Type C token will be a function of the seed location used to generate the token, and are thus, to some extent arbitrary. There are typically many potential seed locations for each particular token, with each potential seed location generating a token with slightly different boundaries and spatial extent because of differences among the color values of the pixels of the various seeds, within the noise ranges of the recording equipment.

[0035] Figure 5 is a flow chart for constructing Type B tokens via an arbitrary boundary removal technique, according to a feature of the present invention. In step 200, the CPU 12 is provided with a set (T_c) of Type C tokens generated with a seed size (S) via the routine of figure 3a, with neck removal via the routine of figure 4. The seed size $S = S_{\text{max}}$, for example, $S = 4$ pixels. In step 202, for each Type C token, t_c in the set T_c the CPU 12 selects a number (for example 50) of potential seeds s_1 , to s_n . In our example, each selected seed will be a 4X4 pixel array from within the token region, the pixels of the array being of approximately equal values (within the noise levels of the recording device).

[0036] In step 204, the CPU 12 grows a new Type C token, utilizing the routines of figures 3a and 4, from each seed location, s_1 , to s_n of each token t_c in the set T_c . The newly grown tokens for each token t_c are designated as tokens r_{c1} to r_{cn} . The newly grown tokens r_{c1} to r_{cn} for each token t_c generally overlap the original Type C token t_c , as well as one another.

[0037] In step 206, the CPU 12 operates to merge the newly generated tokens r_{c1} to r_{cn} of each token t_c , respectively. The result is a new token R_t corresponding to each original token t_c in the set T_c . Each new token R_t encompasses all of the regions of the respective overlapping tokens r_{c1} to r_{cn} generated from the corresponding original token t_c . The

unions of the regions comprising the respective merged new tokens R_t are each a more extensive token than the original Type C tokens of the set. The resulting merged new tokens R_t result in regions of the image file 18, each of a much broader range of variation between the pixels of the respective token R_t than the original Type C token, yet the range of variation among the constituent pixels will still be relatively smooth. R_t is defined as a limited form of Type B token, Type B_{ab1} , to indicate a token generated by the first stage (steps 200-206) of the arbitrary boundary removal technique according to a feature of the present invention.

[0038] In step 208, the CPU 12 stores each of the Type B_{ab1} tokens generated in steps 202-206 from the set of tokens T_c , and proceeds to step 210. Type B_{ab1} tokens generated via execution of steps 202-206 may overlap significantly. In step 210, the CPU 12 operates to merge the R_t tokens stored in step 208 that overlap each other by a certain percentage of their respective sizes. For example, a 30% overlap is generally sufficient to provide few, if any, false positive merges that combine regions containing different materials. The new set of merged tokens still may have overlapping tokens, for example, previously overlapping tokens that had a less than 30% overlap. After all merges are complete, the CPU 12 proceeds to step 212.

[0039] In step 212, the CPU 12 identifies all pixels that are in more than one token (that is in an overlapping portion of two or more tokens). Each identified pixel is assigned to the token occupying the largest region of the image. Thus, all overlapping tokens are modified to eliminate all overlaps.

[0040] In step 214, the CPU 12 stores the final set of merged and modified tokens, now designated as Type B_{ab2} tokens, and then exits the routine. As noted above, the Type B_{ab2} tokens were generated from Type C tokens wether the Type C tokens were generated using N color band values (RGB in our example) of the pixel or the filter output representation of the pixel, in the event the image was filtered.

[0041] A second exemplary technique according to the present invention, for using Type C tokens to create Type B tokens, is adjacent planar token merging. The adjacent planar token merging can be implemented when an image depicts areas of uniform color, that is for non-textured regions of an image. Initially, a token graph is used to identify tokens that are near to one another. Figure 6 shows a flow chart for creating a token graph, according to a feature of the present invention. Each token t_c in the set of Type C tokens T_c , generated through execution of the routines of figures 3a and 4, is evaluated in terms of a maximum distance D_{max} between tokens defining a neighboring pair of tokens, t_c , t_n , of the set T_c , a minimum number of token perimeter pixels, P_{min} , in each token of the neighboring pair of tokens, and a minimum fraction of perimeter pixels, F_{min} , of each token of a neighboring pair of tokens, required to be within D_{max} .

[0042] In step 300, the CPU 12 selects a Type C token t_c in the set of Type C tokens T_c , and identifies the pixels of the selected token t_c forming the perimeter of the token. In a decision block 302, the CPU 12 determines whether the number of perimeter pixels is less than P_{min} , for example 10 pixels.

[0043] If yes, the CPU 12 proceeds to decision block 304 to determine whether there are any remaining tokens t_c in the set of Type C tokens T_c . If yes, the CPU 12 returns to step 300, if no, the CPU 12 exits the routine 306.

[0044] If no, the CPU 12 proceeds to step 308. In step 308, the CPU 12 generates a bounding box used as a mask to surround the selected token t_c . The bounding box is dimensioned to be at least D_{max} larger than the selected token t_c in all directions. A known distance transform (for example, as described in P. Felzenszwalb and D. Huttenlocher, Distance Transforms of Sampled Functions, Cornell Computing and Information Science Technical Report TR2004-1963, September 2004), is executed to find the distance from each perimeter pixel of the selected token t_c to all the pixels in the surrounding bounding box. The output of the distance transform comprises two maps, each of the same size as the bounding box, a distance map and a closest pixel map. The distance map includes the Euclidean distance from each pixel of the bounding box to the nearest perimeter pixel of the selected token t_c . The closest pixel map identifies, for each pixel in the distance map, which perimeter pixel is the closest to it.

[0045] In step 310, the CPU 12 scans the distance map generated in step 308 to identify tokens corresponding to pixels of the bounding box (from the region map generated via the routine of figure 3a), to identify a token from among all tokens represented by pixels in the bounding box, that has a number N_{cn} of pixels within the distance D_{max} , wherein N_{cn} is greater than P_{min} , and greater than $F_{min} * \text{perimeter pixels of the respective token}$ and the average distance between the respective token and t_c is the lowest of the tokens corresponding to the pixels in the bounding box. If these conditions are satisfied, the respective token is designated t_n of a possible token pair t_c , t_n , and a link L_{cn} is marked active.

[0046] In step 312, the CPU 12 checks to determine whether a reciprocal link L_{cn} is also marked active, and when it is marked active, the CPU 12 marks and stores in the token graph, an indication that the token pair t_c , t_n is a neighboring token pair. The reciprocal link refers to the link status in the evaluation of the token designated as t_n in the current evaluation. If that token has yet to be evaluated, the pair is not designated as a neighboring token pair until the link L_{cn} is verified as active in the subsequent evaluation of the token t_n . The CPU 12 then returns to decision block 304 to determine whether there are any further tokens in the set T_c .

[0047] Upon completion of the token graph, the CPU 12 utilizes token pair information stored in the graph in the execution of the routine of figure 7. Figure 7 shows a flow chart for constructing Type B tokens via the adjacent planar token merging technique, according to a feature of the present invention. In the adjacent planer merging technique, pairs of tokens are examined to determine wether there is a smooth and coherent change in color values, in a two dimensional

measure, between the tokens of the pair. The color change is examined in terms of a planar representation of each channel of the color, for example the RGB components of the pixels according to the exemplary embodiments of the present invention. A smooth change is defined as the condition when a set of planes (one plane per color component) is a good fit for the pixel values of two neighboring tokens. In summary, neighboring tokens are considered the same

material and a Type B token when the color change in a two-dimensional sense is approximately planar.
[0048] In step 320, the CPU 12 selects a token pair t_c, t_n from the token graph. In decision block 322, the CPU 12 determines whether the mean color in token t_c is significantly different from the mean color in the token t_n . The difference can be a function of a z-score, a known statistical measurement (see, for example, Abdi, H. (2007), Z-scores, in N.J. Salkind (Ed.), Encyclopedia of Measurement and Statistics, Thousand Oaks, CA: Sage), for example, a z-score greater than 3.0.

[0049] If the mean colors of the token pair are different, the CPU 12 proceeds to decision block 324 to determine whether there are any additional token pairs in the token graph. If yes, the CPU 12 returns to step 320. If no, the CPU 12 exits the routine (step 326).

[0050] If the mean colors are within the z-score parameter, the CPU 12 proceeds to step 328. In step 328, the CPU 12 performs a mathematical operation such as, for example, a least median of squares regression (see, for example, Peter J. Rousseeuw, Least Median of Squares Regression, Journal of the American Statistical Association, Vol. 79, No. 388 (Dec., 1984), pp. 871-880) to fit a plane to each color channel of the pixels (in our example RGB) of the token pair t_c, t_n , as a function of row n and column m (see figure 2), the planes being defined by the equations:

$$R = X_{Rn} + Y_{Rm} + Z_R \quad G = X_{Gn} + Y_{Gm} + Z_G \quad B = X_{Bn} + Y_{Bm} + Z_B$$

wherein parameter values X, Y and C are determined by the least median of squares regression operation of the CPU 12.

[0051] Upon completion of the plane fitting operation, the CPU 12 proceeds to step 330. In step 330, the CPU 12 examines each pixel of each of the tokens of the token pair t_c, t_n to calculate the z-score between each pixel of the tokens and the planar fit expressed by the equation of the least median of squares regression operation. When at least a threshold percentage of the pixels of each token of the pair (for example, 80%), are within a maximum z-score (for example, 0.75), then the neighboring token pair is marked in the token graph as indicating the same material in the image. After completion of step 330, the CPU 12 returns to decision block 324.

[0052] Upon exiting the routine of figure 7, the CPU 12 examines the token graph for all token pairs indicating the same material. The CPU 12 can achieve the examination through performance of a known technique such as, for example, a union find algorithm. (See, for example, Zvi Galil and Giuseppe F. Italiano. Data structures and algorithms for disjoint set union problems, ACM Computing Surveys, Volume 23, Issue 3 (September 1991), pages 319-344). As a simple example, assume a set of seven Type C tokens $T_1, T_2, T_3, T_4, T_5, T_6, T_7$. Assume that the result of the execution of figure 7, (performance of the adjacent planar analysis), indicates that tokens T_1 and T_2 are marked as the same material, and tokens T_1 and T_3 are also marked as the same material. Moreover, the results further indicate that tokens T_4 and T_5 are marked as the same material, and tokens T_5 and T_6 are also marked as the same material. The result of execution of the union find algorithm would therefore indicate that tokens $\{T_1, T_2, T_3\}$ form a first group within the image consisting of a single material, tokens $\{T_4, T_5, T_6\}$ form a second group within the image consisting of a single material, and token $\{T_7\}$ forms a third group within the image consisting of a single material. The groups $\{T_1, T_2, T_3\}$, $\{T_4, T_5, T_6\}$ and $\{T_7\}$ form three Type B tokens.

[0053] A third exemplary technique according to the present invention, for using Type C tokens to create Type B tokens, is a local token analysis. A local token approach generates Type C tokens using a window analysis of a scene depicted in an image file 18. Such tokens are designated as Type C_w tokens. Figure 8 is a flow chart for generating Type C_w tokens via the local token analysis technique, according to a feature of the present invention.

[0054] In step 400, the CPU 12 places a window of fixed size, for example, a 33x33 pixel array mask, over a preselected series of scan positions over the image. The window can be a shape other than a square. The scan positions are offset from one another by a fixed amount, for example $\frac{1}{2}$ window size, and are arranged, in total, to fully cover the image. The window area of pixels at each scan position generates a Type C_w token, though not every pixel within the window at the respective scan position is in the Type C_w token generated at the respective scan position.

[0055] At each scan position (step 402), the CPU 12 operates, as a function of the pixels within the window, to fit each of a set of planes, one corresponding to the intensity of each color channel (for example, RGB), and an RGB line in RGB space, characterized by a start point l_0 and an end point l_1 of the colors within the window. The planar fit provides a spatial representation of the pixel intensity within the window, and the line fit provides a spectral representation of the pixels within the window.

[0056] For the planar fit, the planes are defined by the equations:

$$R = X_{Rn} + Y_{Rm} + Z_R \quad G = X_{Gn} + Y_{Gm} + Z_G \quad B = X_{Bn} + Y_{Bm} + Z_B$$

wherein parameter values X, Y and C are determined by CPU 12 by executing a mathematical operation such as the least median of squares regression discussed above, a least-squares estimator, such as singular value decomposition, or a robust estimator such as RANSAC (see, for example, M. A. Fischler, R. C. Bolles. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. Comm. of the ACM, Vol 24, pp 381-395, 1981).

[0057] For the RGB line fit, the line is defined by: $l(r,g,b) = l_0(r,g,b) + t(l_1(r,g,b) - l_0(r,g,b))$ wherein the parameter t has a value between 0 and 1, and can be determined by the CPU 12 utilizing any of the mathematical techniques used to find the planar fit.

[0058] At each scan position, after completion of step 402, the CPU 12 operates in step 404 to examine each pixel in the window in respect of each of the planar fit representation and RGB line representation corresponding to the respective window scan position. For each pixel, the CPU 12 determines an error factor for the pixel relative to each of the established planes and RGB line. The error factor is related to the absolute distance of the pixel to its projection on either from either the planar fit or the RGB line fit. The error factor can be a function of the noise present in the recording equipment or be a percentage of the maximum RGB value within the window, for example 1%. Any pixel distance within the error factor relative to either the spatial planar fit or the spectral line fit is labeled an inlier for the Type C_w token being generated at the respective scan position. The CPU 12 also records for the Type C_w token being generated at the respective scan position, a list of all inlier pixels.

[0059] At each scan position, after completion of step 404, the CPU 12 operates in step 406 to assign a membership value to each inlier pixel in the window. The membership value can be based upon the distance of the inlier pixel from either the planar fit or the RGB line fit. In one exemplary embodiment of the present invention, the membership value is the inverse of the distance used to determine inlier status for the pixel. In a second exemplary embodiment, a zero-centered Gaussian distribution with a standard deviation is executed to calculate membership values for the inlier pixels.

[0060] After all of the scan positions are processed to generate the Type C_w tokens, one per scan position, the CPU 12 operates to compile and store a token data list (step 408). The token data list contains two lists. A first list lists all of the pixels in the image file 18, and for each pixel, an indication of each Type C_w token to which it labeled as an inlier pixel, and the corresponding membership value. A second list lists all of the generated Type C_w tokens, and for each token an indication of the inlier pixels of the respective token, and the corresponding membership value. After compiling and storing the token data list, the CPU 12 exits the routine (step 410).

[0061] Figure 9 is a flow chart for constructing Type B tokens from the Type C_w tokens generated via the local token analysis technique, according to a feature of the present invention. In step 420, the CPU 12 calculates a similarity of parameters of the spatial planer dimensions and spectral RGB lines of adjacent or overlapping Type C_w tokens generated through execution of the routine of figure 8. Overlapping and adjacent Type C_w tokens can be defined as tokens corresponding to scan positions that overlap or are contiguous. A similarity threshold can be set as a percentage of difference between each of the spatial planer dimensions and spectral RGB lines of two overlapping or adjacent Type C_w tokens being compared. The percentage can be a function of the noise of, for example, the camera 14 used to record the scene of the image file 18. All overlapping or adjacent Type C_w token pairs having a calculated similarity within the similarity threshold are placed on a list.

[0062] In step 422, the CPU 12 sorts the list of overlapping or adjacent Type C_w token pairs having a calculated similarity within the similarity threshold, in the order of most similar to least similar pairs. In step 424, the CPU 12 merges similar token pairs, in the order of the sort, and labeling pairs as per degree of similarity. Each merged token pair will be considered a Type B token. In step 426, the CPU 12 stores the list of Type B tokens, and exits the routine.

[0063] In a further exemplary technique according to the present invention, a filter response, such as, for example, a second derivative filter response, is executed by the CPU 12 to identify regions of an image having uniform material reflectance. The identified regions are then used to generate Type B tokens. A filter response can be implemented by performing a convolution of pixel values. Convolution is a mathematical technique that applies a filter mask to an array of pixels to determine a function response value for a pixel at the center of the array. The filter mask represents a mathematical function such as, for example, a Laplacian of Gaussian kernel (LoG). The LoG filter is a known filter used to determine a mathematical function output value at each pixel location of an image, an output value that is related to a derivative value at the respective pixel location, and utilized to detect material edges in an image.

[0064] Figure 10a shows intensity profiles within an image, and corresponding first derivative and second derivative profiles for each of a material edge, a single material with a shadow over a portion of the material and a single material under constant illumination. A derivative is a measure of a rate of change. In an image, a first derivative measures the rate of change in the intensity of a color band, for example, the rate of change of red, from pixel to pixel of the image. A second derivative is a measure of the rate of change of the rate of change, i.e., a measure of the rate of change of

the first derivative. The A column of figure 10a shows the intensity profile, and first and second derivative profiles at a material edge within an image. As shown in column A, the color intensity profile forms a sharp edge profile at a material boundary, resulting in a spike in the first derivative value reflecting the sharp rate of change of material intensity at the material boundary. Consequently, the second derivative value undergoes a zero crossing, first peaking in the positive direction as the first derivative peaks at the boundary, and then crossing zero and peaking in the negative direction as the first derivative peak returns to zero.

[0065] In the known LoG filter technique, high first derivative values, first derivative peaks or second derivative zero crossings among pixels, are utilized as indicia of a material edge, as reflected in column A of figure 10a. In contrast to the first derivative peaks and second derivative zero crossings of a material edge, columns B and C of figure 10a show, respectively, intensity profiles and corresponding first derivative and second derivative profiles for each of a single material with a shadow over a portion of the material and a single material under constant illumination. In column B, the intensity of the image undergoes an attenuated edge, corresponding to the penumbra of a shadow falling across pixels representing a single material. Thus, the first derivative slowly raises and falls with the attenuated change of a penumbra, and the second derivative remains at a relatively low value throughout the slow rise and fall of the first derivative remains at a relatively low value throughout the slow rise and fall of the first derivative value. As shown in column C, a constant value for the illumination of a single material results in a constant low value among pixels for each of the first derivative and second derivative values.

[0066] According to a feature of the present invention, a second derivative filter, such as, for example, a LoG filter, is used in a novel manner, to identify regions of an image, for example, patches comprising pixels having uniformly low second derivative values, as shown in each of columns B and C of figure 10a. In an exemplary embodiment, each patch is of a preselected size (for example, either a 3X3 pixel array, a 5X5 pixel array or a 7X7 pixel array). As described above, a patch of uniform material reflectance under uniform light will have a near constant color value across the pixels of the patch, resulting in a low value for the second derivative at each pixel location in the patch. Additionally, a constant material reflectance in an attenuated penumbra will also result in a low value for the second derivative among pixels of a patch.

[0067] Thus, ascertaining the location of a patch of pixels, all having a low value for the second derivative, provides a high indication of likelihood that there is not a material edge within the patch, and that all of the pixels of the patch have the same material reflectance. The present invention contemplates any function response among a patch of pixels that is relevant to a likelihood of uniform material reflectance, such as, for example, second derivative values, third derivative values or other measures of the curvature of the intensity profile of the image, where the intensity profile is the shape of the image in (x, y, intensity) space.

[0068] A Laplacian filter provides a sum of partial second derivatives at a point in the image. A Laplacian filter can therefore be implemented to ascertain function outputs proportional to the second derivative values for pixels of an image. The Laplacian of an image I is defined as:

$$\nabla^2 I = \frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2}$$

The above function is the sum of the second partial derivatives in the x and y directions (along the n rows and m columns of the array of figure 2), and is proportional to the average second derivative across all directions.

[0069] Due to the presence of noise and small-scale variations in an image, first and second derivative analysis of an image is typically performed on a blurred version of the image. A blurred version of an image smooths out any irregularities caused by noise and image variations, and can be obtained through any standard blurring function. For example, a convolution of the image with a filter mask based upon a Gaussian function. A Gaussian function can be defined by:

$$G_{\sigma}(x, y) = \frac{1}{2\pi\sigma^2} e^{-(x^2+y^2)/(2\sigma^2)}$$

The σ factor is the standard deviation and controls the degree of blurring of a convolved image. A small value for σ results in a relatively small degree of blurring, while a large value for σ results in significant blurring.

[0070] An LoG function combines the Laplacian and Gaussian functions into a single function that serves to both blur the image and ascertain the Laplacian values to provide accurate outputs proportional to the second derivative values for the pixels of an image. In place of the LoG function, a Difference of Gaussians or convolution with simple derivative masks on a blurred image can be used. An LoG function equation is as follows:

$$\nabla^2 G_\sigma(x, y) = \frac{\partial^2 G_\sigma}{\partial x^2} + \frac{\partial^2 G_\sigma}{\partial y^2} = \left(\frac{x^2 + y^2 - 2\sigma^2}{2\sigma^6\pi} \right) e^{-(x^2+y^2)/(2\sigma^2)}$$

5
 10 **[0071]** Figure 10b is an example of an LoG filter mask. The mask comprises a $2k + 1 \times 2k + 1$ array of multiplication factors, in our example, $k = 3$, for 7 rows of factors, each with 7 columns, thus, a length of 7σ relative to the center, with $\sigma = 1$ in the filter mask of figure 10b. The filter mask of figure 10b is arranged for use in connection with a convolution of the image via 7×7 pixel patch samples of the image file of figure 2. The multiplication factors express the LoG function and are applied to the pixels of the corresponding patch, and all the results are summed to provide a function response value for the pixel at the center of the 7×7 array, as will be described.

15 **[0072]** To that end, figure 10c shows an illustration of a representative pixel patch comprising an 7×7 array within the $n \times m$ pixel array of the image file of figure 2. The 7×7 patch in the example of figure 10b forms the upper left hand corner of the $n \times m$ pixel array of figure 2, and would be used to calculate a value proportional to the second derivative for the pixel $p(4,4)$, the center pixel of the representative array (shown in bold).

20 **[0073]** Each multiplication factor of the filter mask corresponds to a function response for a corresponding pixel of the 7×7 patch, and is used to calculate a multiplication result for the respective pixel of the patch. The multiplication factor applied to each pixel is determined in relation to the combined Laplacian and Gaussian functions set forth above, (an LoG filter mask) as is known in the art. The filter mask expresses the following summation equation:

$$\text{Output } p(n, m) = \sum_{i=-k}^k \sum_{j=-k}^k p(n+i, m+j) \text{LoG}(i, j).$$

25 **[0074]** Each pixel $p(n, m)$ is represented by N color bands, for example RGB, for the red, green and blue components of the color value for the pixel $p(n, m)$. The multiplication factors of the filter mask will be applied separately to each color band of the pixels of the patch, resulting in a separate second derivative value for each color band. In our example, we will describe the calculation for the red band ($p(n, m \text{ RED})$). Thus, $p(1,1 \text{ RED})$ is multiplied by the factor 0.00031426, while the next pixel of the row, $p(1,2 \text{ RED})$ is multiplied by the factor 0.002632078, and so on. The pixels of each row, 1-7 are multiplied in order, across the columns, by the multiplication factors at the corresponding rows and columns of the LoG filter mask. The results of all the multiplications are added together to provide a sum that is proportional to a second derivative value for the center pixel $p(4,4)$, as a function of the LoG function.

30 **[0075]** A similar filter response output is performed for each pixel $p(n, m)$ of the $n \times m$ array and for each color band of each pixel, to compile a list of second derivative values for the pixels $p(n, m)$ of an image file 18. At border regions of the image, for example, pixels within 3 pixels of the image boundary, for pixel values of points of the 7×7 patch that extend beyond the image boundary, estimates can be made for the respective color values, based upon various assumptions, as follows:

- 35 1. assume the values for pixels at the image boundary continue beyond the boundary, such that values for pixels in the patch that extend beyond the image boundary are set at the same color intensity as the color intensity of the nearest boundary pixel;
- 40 2. assume the image reflects at the boundary, for example, the values for pixels of the patch beyond the image boundary are a mirror image of corresponding pixels within the image; or
- 45 3. assume all values outside the image are a constant value, such as 0 or the maximum image value.

[0076] Second derivative values can be either positive or negative. In the exemplary embodiment of the present invention, an absolute value is used for the second derivative. Moreover, the values for each band, in our example, red, green and blue, are represented as a single number per pixel. The single number value can be obtained by taking the maximum second derivative filter response output from among the color bands of the pixel, taking an average for the filter response outputs of all the bands or, in hyperspectral imagery, taking a percentile value from each band, for example the 90th percentile value filter output value from among the color bands.

50 **[0077]** In one embodiment of the filter response feature, second derivative values are determined in a log space version or a gamma transform of the image. When a region of an image is very dark, and, further, in deep shadow, first and second derivatives become small, and are not useful as accurate indications of regions of uniform material reflectance. The use of logarithmic values of color intensity or a gamma transform maintains an equality of color differences, regardless of the illumination flux (or image exposure) affecting the image.

55 **[0078]** Figure 11 is a flow chart for identifying regions of uniform material reflectance within the $n \times m$ pixel array image

file of figure 2, according to the filter response feature of the present invention. In step 500, an image file 18 (see figure 2) is input to the CPU 12. In step 510, the CPU 12 calculates a filter response for each of the pixels $p(n, m)$ of the image file 18. In our exemplary embodiment of the present invention, the filter response is a value proportional to the second derivative at each pixel $p(n, m)$ of the image file 18. Such values can be calculated by the CPU 12 via a convolution of the image file 18, by the CPU 12 with a LoG filter mask, as described above in respect of figures 10b and 10c.

5 [0079] In order to execute the filter response step, appropriate filter parameters are input to the CPU 12. In the case of a LoG filter mask, the relevant parameter is the σ factor, which determines the degree of blurring. The σ factor can range, for example, between 0.5 and 3.0.

10 [0080] In step 520, the CPU 12 operates to locate patches of the pixel array with filter responses that indicate uniform material reflectance among the pixels of the patch. In the example of second derivative values, a low filter response is indicative of the absence of material edges, and thus provides a high degree of likelihood that the patch encompasses a region of the image that has a uniform material reflectance, as described above. Input parameters relevant to step 520 include patch size and shape and a definition of a low filter response. Each patch can be, for example, square, rectangular, circular, and so on. In our example, each patch can comprise a 3 X 3 square or circular patch with a radius of 9. The CPU 12 examines a sequence of patches across the entire image file 18, with each patch being centered on a different one of the pixels $p(n, m)$ of the array of the subject image file 18, across all the pixels of the array.

15 [0081] For each patch examined, the CPU 12 determines whether the filter responses of the pixels of the patch under examination, are below a threshold value. The measure can, for example, be a threshold measured against the average filter response for the pixels of the patch or a threshold measured against the maximum second derivative value among the pixels of the patch. The threshold can be determined empirically, with lower values tending to more significant segmentation (dividing the image into too many regions) and higher values leading to errors (considering two regions to be the same material reflectance when they are not). The CPU 12 lists each examined patch of the subject image file 18 that has a result below the threshold, and therefore indicates a high likelihood of uniform material reflectance within the examined patch.

25 [0082] In step 530, the CPU 12 operates to merge all sufficiently overlapping patches from among patches that are on the above compiled list of patches indicating a high likelihood of uniform material reflectance, to compose a list of merged patches as Type B tokens. A parameter relevant to step 530 is a definition of "sufficiently overlapping." For example, the CPU 12 merges all listed patches that have 5 pixels in common, or 10 pixels in common, and so on. Given that each patch has a high likelihood of uniform material reflectance, and the patches share a number of pixels in common, it can be concluded that the two sufficiently overlapping patches have the sufficiently overlapping patches share the same material reflectance. All listed patches that overlap to the extent that they share the selected number of pixels are merged, and the merging is cumulative for mutually overlapping patches to provide the maximum extent of regions of the image having uniform reflectance. The output (540) of the routine is a list of Type B tokens, a segmentation of the image into regions of uniform reflectance. Inasmuch as each merged region may include pixels having non-zero first derivatives, each identified Type B token may encompass significant illumination variation.

30 [0083] Referring now to figure 12, there is shown a flow chart for combining results from executions of the routine of figure 11 with differing parameters, according to a feature of the present invention. As noted in the description of figure 11, various variable parameters are input to the CPU 12. In our exemplary embodiment, the variable parameters include the σ factor, patch size and shape, a definition of a low filter response and a definition of sufficiently overlapping. Different combinations of parameters yield differing results. One set of parameters can capture some portions of an image well, with well defined Type B tokens, but leave other portions of the image in many small unconnected patches, while another set of parameters might capture different parts of the image well. Thus, by varying values within the set of parameters, and executing the routine of figure 11 several times, one time with each unique set of parameters, the CPU 12 can obtain several sets of segmentation results to provide a more robust overall result that is significantly improved from a result that would be obtained from any one particular set of parameters.

45 [0084] For example, a small patch size will localize material edges more effectively, while larger patch sizes avoid crossing gaps where edges occasionally disappear. The threshold should also be lowered as patch size and the σ factor increase, and so on. To that end, in step 550, an image file 18 is input to the CPU 12. In step 560, the CPU 12 executes the routine of figure 11 several times in succession, each time with a different set of input parameters.

50 [0085] In step 570, the CPU 12 operates to merge the results of the several segmentation operations for the subject image file 18. The merge operation can be similar to the merge step 530 of the routine of figure 11, for example, a merge of all identified Type B tokens that overlap by a selected number of pixels. Alternatively, the CPU 12 can implement the overlap criteria set forth below, in respect of the following description of a further exemplary embodiment of the present invention.

55 [0086] In step 580, the CPU 12 can optionally merge in Type C tokens identified by the CPU 12 during execution of the routine of figure 3a, as described above. In examining pixel patches for low second derivative values, the filter response technique inherently tends to avoid material edges, and thus, the technique provides Type B tokens that can fall short of image areas around edges. To "fill out" Type B tokens identified through the second derivative filter response

technique, the CPU 12 can merge the identified Type B tokens that overlap previously identified Type C tokens. Again the overlap criteria can be as described above. In step 590, the CPU 12 outputs a list of Type B tokens.

[0087] In a further exemplary embodiment of the present invention, the CPU 12 compiles lists of Type B tokens separately generated through each of and/or a combination of one or more of the arbitrary boundary removal, adjacent planar token merging, local token analysis and filter response techniques. The determination of the combination of techniques used depends in part on whether a particular region of the image was filtered because of texturing of the image. Since each Type B token generated through the described techniques likely represents a single material under varying illumination conditions, merging sufficiently overlapping Type B tokens generated through the use of varying and different techniques, provides a resulting, merged Type B token that represents a more extensive area of the image comprising a single material, and approaches the extent of a Type A token.

[0088] Sufficiently overlapping can be defined by satisfaction of certain pixel characteristic criteria, such as, for example:

A) The two Type B tokens have at least n of the original Type C tokens in common, for example, $n = 1$

B) The two Type B tokens have at least n pixels in common, for example, $n = 20$

C) The two Type B tokens have at least $n\%$ overlap, that is at least $n\%$ of the pixels in a first one of the two Type B tokens are also found in the second one of the two Type B tokens or vice versa, wherein, for example $n\% = 10\%$.

D) The percentage of pixels in a smaller one of the two Type B tokens, also found in the larger one of the two Type B tokens is above a preselected threshold, for example 15%.

E) A preselected combination of criteria A-D.

[0089] Merging of two sufficiently overlapping Type B tokens can be accomplished via a mathematical operation such as execution of the union find algorithm discussed above. In the case of two overlapping Type B tokens that do not satisfy the above discussed criteria, the overlapping pixels of the two tokens can be assigned to the larger one of the two Type B tokens.

[0090] As a result of execution of the token generation and merging techniques according to features of the present invention, an image can be accurately segmented into tokens representing discrete materials depicted in the scene, thus providing a basis for reduction of inaccuracies in image processing that may be caused by varying illumination conditions, for more robust computer processing of images.

[0091] In the preceding specification, the invention has been described with reference to specific exemplary embodiments and examples thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the scope of the invention as set forth in the claims that follow. The specification and drawings are accordingly to be regarded in an illustrative manner rather than a restrictive sense.

Claims

1. An automated, computerized method for processing an image, comprising the steps of:

calculating (510) a filter response for each of the pixels of the image, wherein at least one filter parameter is provided as input parameter to the calculation of the filter response, and wherein the filter response is indicative as to the presence or absence of a material edge;

identifying (520) patches of pixels in the image for which the calculated filter responses indicate uniform material reflectance among the pixels of the patch, wherein the patch size, patch shape and a criterion for the patch pixel filter responses that indicates uniform material reflectance of the patch are provided as input parameters for this step; and

merging (530) all patches identified in the previous step which exhibit a sufficient spatial overlap, wherein a definition of the sufficient overlap is provided as input parameter to this step.

2. The method of claim 1 wherein the merged patches provide single material token regions for use in processing material and illumination aspects of the image.

3. The method of claim 1 wherein the step of calculating a filter response is carried out by applying a second derivative filter to the pixels of the image.

4. The method of claim 3 wherein the second derivative filter comprises an LoG filter mask.

5. The method of claim 1 wherein calculating a filter response is carried out by applying a filter mask to pixels of the image, and patches of pixels having a predetermined filter response value are identified.

6. The method of claim 1, wherein in order to identify patches of pixels in the image for which the calculated filter responses indicate uniform material reflectance among the pixels of the patch, patches across the entire image are examined, each patch being centered on a different one of the pixels of the image, across all the pixels of the image.
- 5 7. The method of claim 1 including the further steps of selecting sets of values for the input parameters, and repeating the steps listed in claim 1 once per each one of the sets of values for the input parameters.
8. The method of claim 7 including the further step of merging results from repeated performing of the steps listed in claim 1.
- 10 9. The method of claim 2 including the further steps of identifying spatio-spectral information for the image, and utilizing the spatio-spectral information to identify single material token regions in the image.
- 15 10. The method of claim 9 including the further step of merging single material token regions provided by the merged patches of uniform material reflectance, and single material token regions identified by the step of identifying spatio-spectral information for the image.
11. A computer system which comprises:
- 20 a CPU; and
a memory storing an image file;
the CPU arranged and configured to execute a routine to perform the method of claim 1.
12. The computer system of claim 11 wherein the CPU is further arranged to execute a routine to perform the method of claim 2.
- 25

Patentansprüche

- 30 1. Automatisiertes, computergesteuertes Verfahren zum Verarbeiten eines Bildes, umfassend die Schritte:
- Berechnen (510) einer Filterantwort für jedes der Pixel des Bildes, wobei wenigstens ein Filterparameter als Eingangsparameter zur Berechnung der Filterantwort bereitgestellt wird und wobei die Filterantwort auf das Vorliegen oder Fehlen einer Materialkante hinweist;
- 35 Identifizieren (520) von Pixel-Patches in dem Bild, bei denen die berechneten Filterantworten eine einheitliche Materialreflexion unter den Pixeln des Patches anzeigen, wobei die Patch-Größe, Patch-Form und ein Kriterium für die Patch-Pixel-Filterantworten, das auf eine einheitliche Materialreflexion des Patches hinweist, als Eingangsparameter für diesen Schritt bereitgestellt werden, und
- 40 Verschmelzen (530) aller im vorhergehenden Schritt identifizierten Patches, die eine ausreichende räumliche Überlappung aufweisen, wobei eine Definition der ausreichenden Überlappung als Eingangsparameter für diesen Schritt bereitgestellt wird.
2. Verfahren nach Anspruch 1, wobei die verschmolzenen Patches einzelne Material-Token-Bereiche zur Verwendung für die Verarbeitung von Material- und Beleuchtungsaspekten des Bildes bereit stellen.
- 45 3. Verfahren nach Anspruch 1, wobei der Schritt des Berechnens einer Filterantwort durch Anwendung eines zweiten Ableitungsfilters auf die Pixel des Bildes durchgeführt wird.
- 50 4. Verfahren nach Anspruch 3, wobei der zweite Ableitungsfilter eine LoG-Filtermaske umfasst.
5. Verfahren nach Anspruch 1, wobei das Berechnen einer Filterantwort durch Anwendung einer Filtermaske auf Pixel des Bildes ausgeführt wird und Pixel-Patches, die einen vorbestimmten Filterantwortwert aufweisen, identifiziert werden.
- 55 6. Verfahren nach Anspruch 1, wobei Patches über das gesamte Bild untersucht werden, um in dem Bild Pixel-Patches zu identifizieren, bei denen die berechneten Filterantworten auf eine einheitliche Materialreflexion unter den Pixeln des Patches hinweisen, wobei jeder Patch um ein anderes der Pixel des Bildes unter allen Pixeln des Bildes zentriert

ist.

- 5 7. Verfahren nach Anspruch 1, mit den weiteren Schritten: Auswählen von Gruppen von Werten für die Eingangsparameter und Wiederholen der in Anspruch 1 aufgeführten Schritte jeweils einmal pro Wertesatz für die Eingangsparameter.
8. Verfahren nach Anspruch 7, mit dem weiteren Schritt des Verschmelzens der Ergebnisse aus der wiederholten Durchführung der in Anspruch 1 aufgeführten Schritte.
- 10 9. Verfahren nach Anspruch 2, mit den weiteren Schritten: Identifizieren räumlich-spektraler Informationen für das Bild und Nutzen der räumlich-spektralen Informationen zum Identifizieren einzelner Material-Token-Bereiche in dem Bild.
- 15 10. Verfahren nach Anspruch 9, mit dem weiteren Schritt des Verschmelzens einzelner Material-Token-Bereiche, die von den verschmolzenen Patches einheitlicher Materialreflexion bereit gestellt werden, und einzelner Material-Token-Bereiche, die durch den Schritt des Identifizierens räumlich-spektraler Informationen für das Bild identifiziert werden.
11. Computersystem, umfassend:
- 20 eine CPU und
einen Speicher zum Speichern einer Bilddatei,
wobei die CPU dazu angeordnet und konfiguriert ist, ein Programm zur Durchführung des Verfahrens nach Anspruch 1 auszuführen.
- 25 12. Computersystem nach Anspruch 11, wobei die CPU ferner dazu angeordnet ist, ein Programm zur Durchführung des Verfahrens nach Anspruch 2 auszuführen.

Revendications

- 30 1. Procédé automatisé et informatisé pour traiter une image, comprenant les étapes :
- de calcul (510) d'une réponse de filtre pour chacun des pixels de l'image, dans lequel au moins un paramètre de filtre est fourni en tant que paramètre d'entrée pour le calcul de la réponse de filtre, et dans lequel la réponse de filtre est indicative de la présence ou de l'absence d'un contour de matériau ;
- 35 d'identification (520) de parcelles de pixels dans l'image pour lesquels les réponses de filtre calculées indiquent une réflectance de matériau uniforme parmi les pixels de la parcelle, dans lequel la taille de parcelle, la forme de parcelle et un critère pour les réponses de filtre de pixel de parcelle qui indique une réflectance de matériau uniforme de la parcelle sont fournis en tant que paramètres d'entrée pour cette étape ; et
- 40 de fusion (530) de toutes les parcelles identifiées à l'étape précédente qui présentent un chevauchement spatial suffisant,
dans lequel une définition du chevauchement suffisant est fournie en tant que paramètre d'entrée à cette étape.
2. Procédé selon la revendication 1, dans lequel les parcelles fusionnées fournissent des régions de jeton de matériau unique pour une utilisation dans le traitement des aspects de matériau et d'illumination de l'image.
- 45 3. Procédé selon la revendication 1, dans lequel l'étape de calcul d'une réponse de filtre est effectuée en appliquant un filtre de dérivée seconde aux pixels de l'image.
- 50 4. Procédé selon la revendication 3, dans lequel le filtre de dérivé seconde comprend un masque de filtre LoG.
5. Procédé selon la revendication 1, dans lequel le calcul d'une réponse de filtre est effectué en appliquant un masque de filtre aux pixels de l'image, et des parcelles de pixels ayant une valeur de réponse de filtre prédéterminée sont identifiées.
- 55 6. Procédé selon la revendication 1, dans lequel, afin d'identifier des parcelles de pixels dans l'image pour lesquelles les réponses de filtre calculées indiquent une réflectance de matériau uniforme parmi les pixels de la parcelle, des parcelles dans l'image entière sont examinées, chaque parcelle étant centrée sur un pixel différent parmi les pixels

de l'image, dans tous les pixels de l'image.

- 5
7. Procédé selon la revendication 1, comportant les étapes supplémentaires de sélection d'ensembles de valeur pour les paramètres d'entrée, et de répétition des étapes énumérées à la revendication 1 une fois pour chacun des ensembles de valeur pour les paramètres d'entrée.
8. Procédé selon la revendication 7, comportant l'étape supplémentaire de fusion de résultats à partir de la réalisation répétée des étapes énumérées à la revendication 1.
- 10
9. Procédé selon la revendication 2, comportant les étapes supplémentaires d'identification d'informations spatio-spectrales pour l'image, et d'utilisation des informations spatio-spectrales pour identifier des régions de jeton de matériau unique dans l'image.
- 15
10. Procédé selon la revendication 9, comportant une étape supplémentaire de fusion de régions de jeton de matériau unique fournies par les parcelles fusionnées de réflectance de matériau uniforme, et de régions de jeton de matériau unique identifiées par l'étape d'identification d'informations spatio-spectrales dans l'image.
11. Système d'ordinateur qui comprend :
- 20
- une UC ; et
 - une mémoire stockant un fichier d'image ;
 - l'UC étant agencée et configurée pour exécuter une routine afin de réaliser le procédé de la revendication 1.
- 25
12. Système d'ordinateur selon la revendication 11, dans lequel l'UC est en outre agencée pour exécuter une routine afin de réaliser le procédé de la revendication 2.
- 30
- 35
- 40
- 45
- 50
- 55

Computer System Configured to Operate on Images




FIG. 1

Pixel Array for Storing Image Data




Image File
18

FIG. 2

Identifying Token Regions in an Image

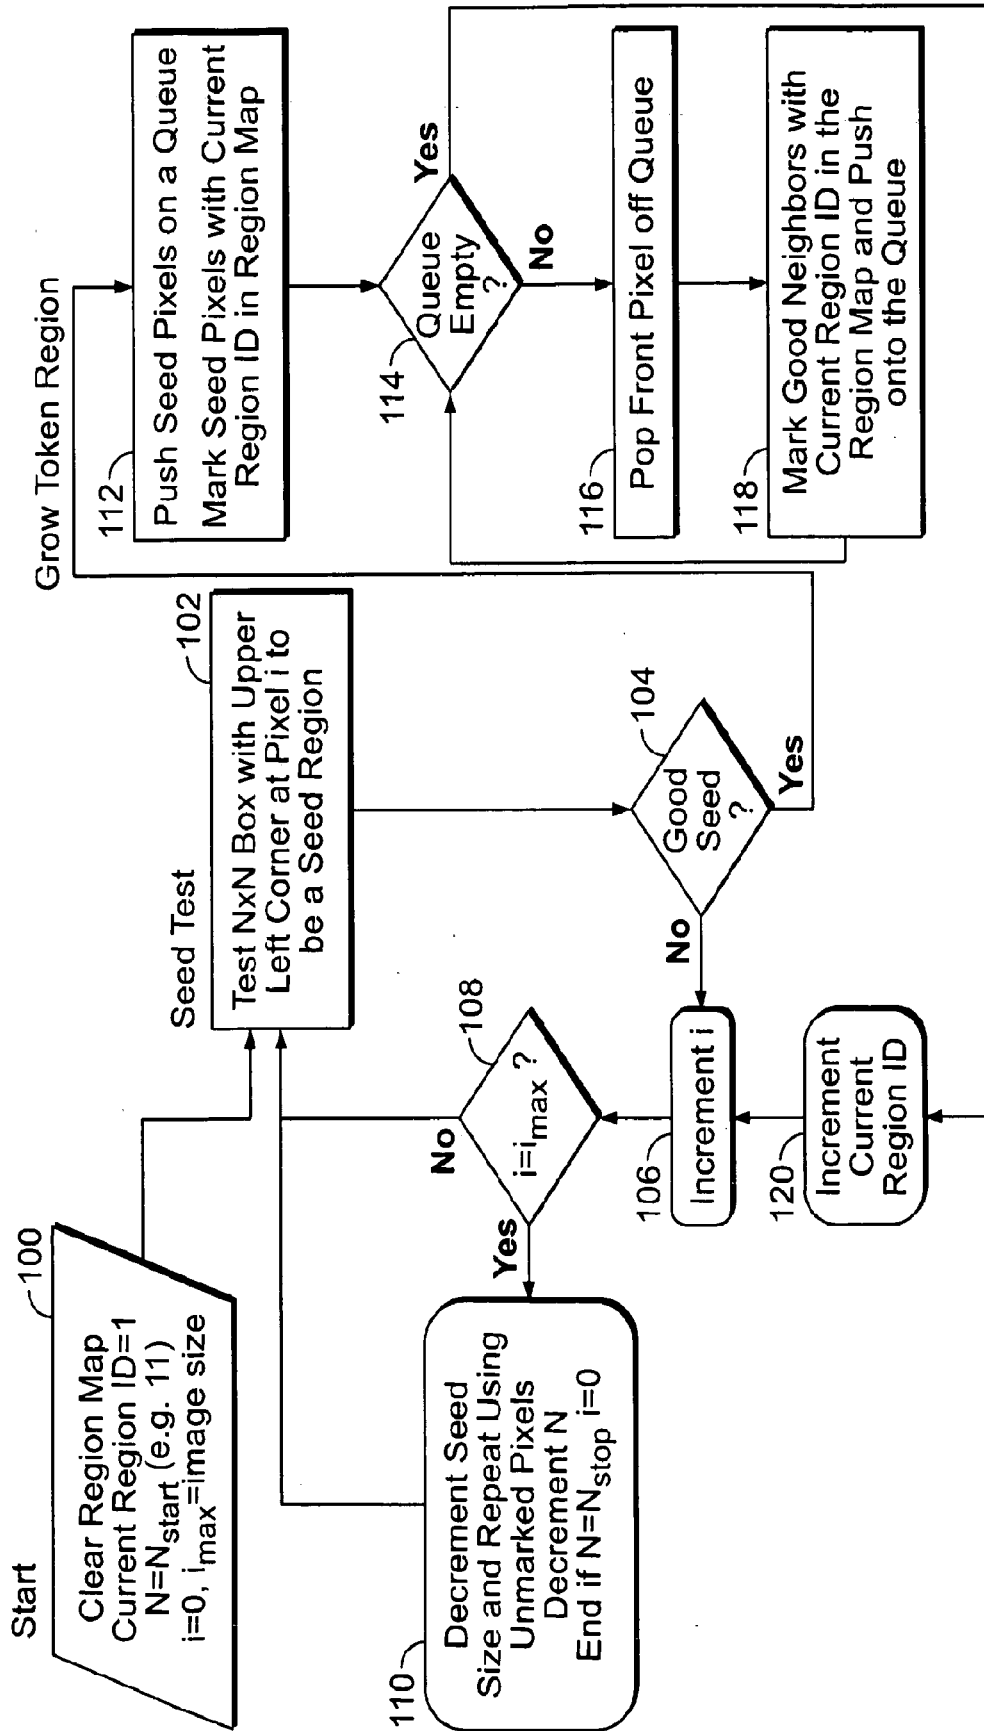



FIG. 3A

Examples of Identifying Token Regions in an Image

Original Image



- ⊗ Teal In Shadow
- ⊙ Teal
- ⊖ Blue In Shadow
- ⊘ Blue
- ⊕ Yellow

FIG.3B

Token Regions




FIG.3C

Type B Tokens




FIG.3D




FIG.4

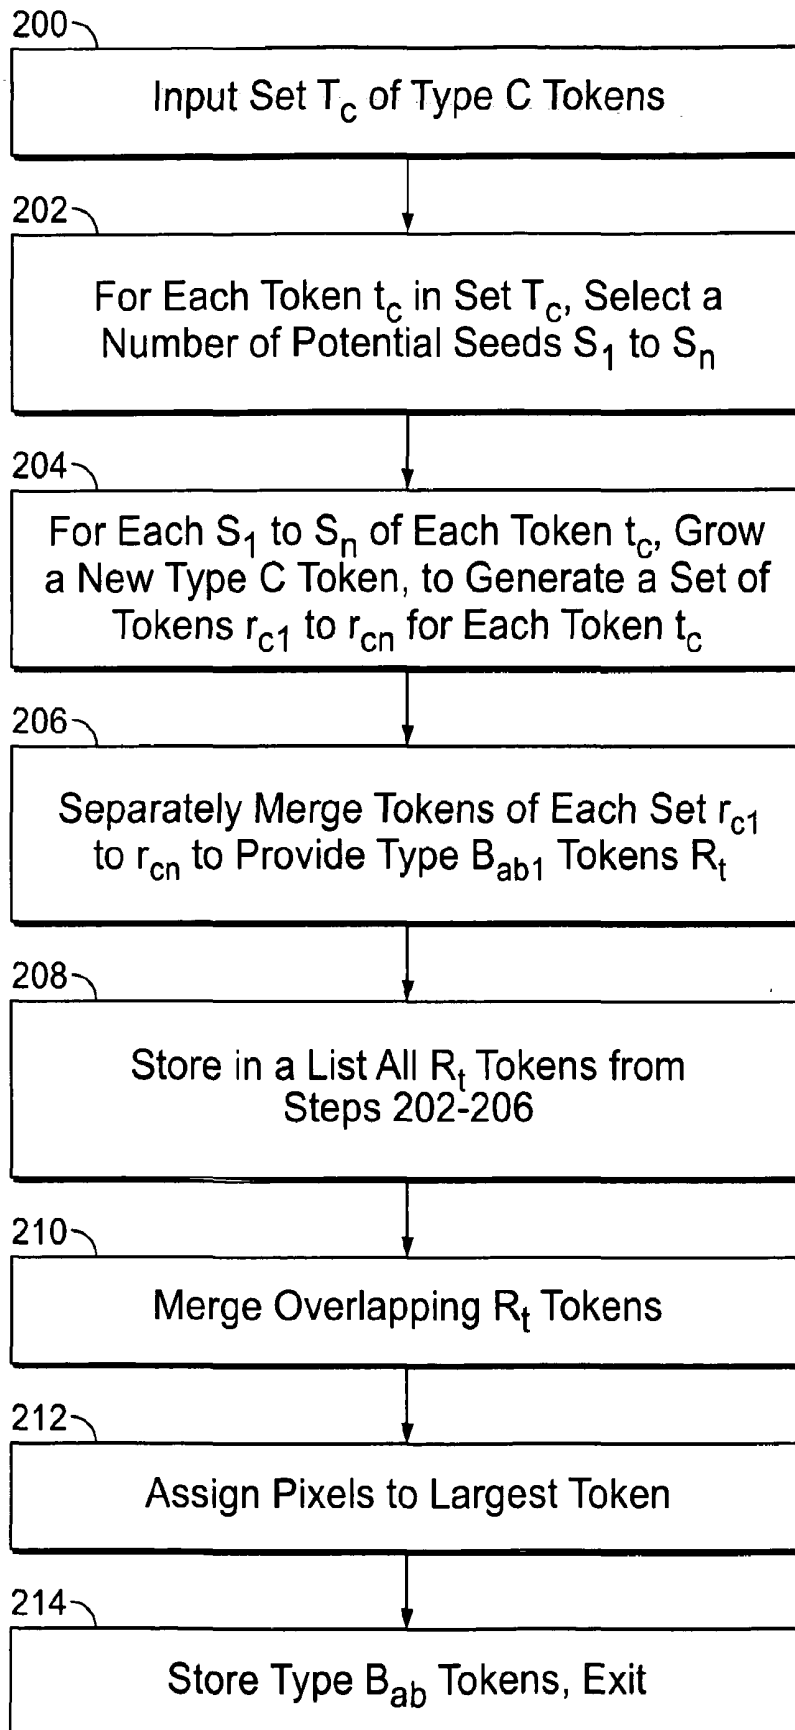


FIG. 5

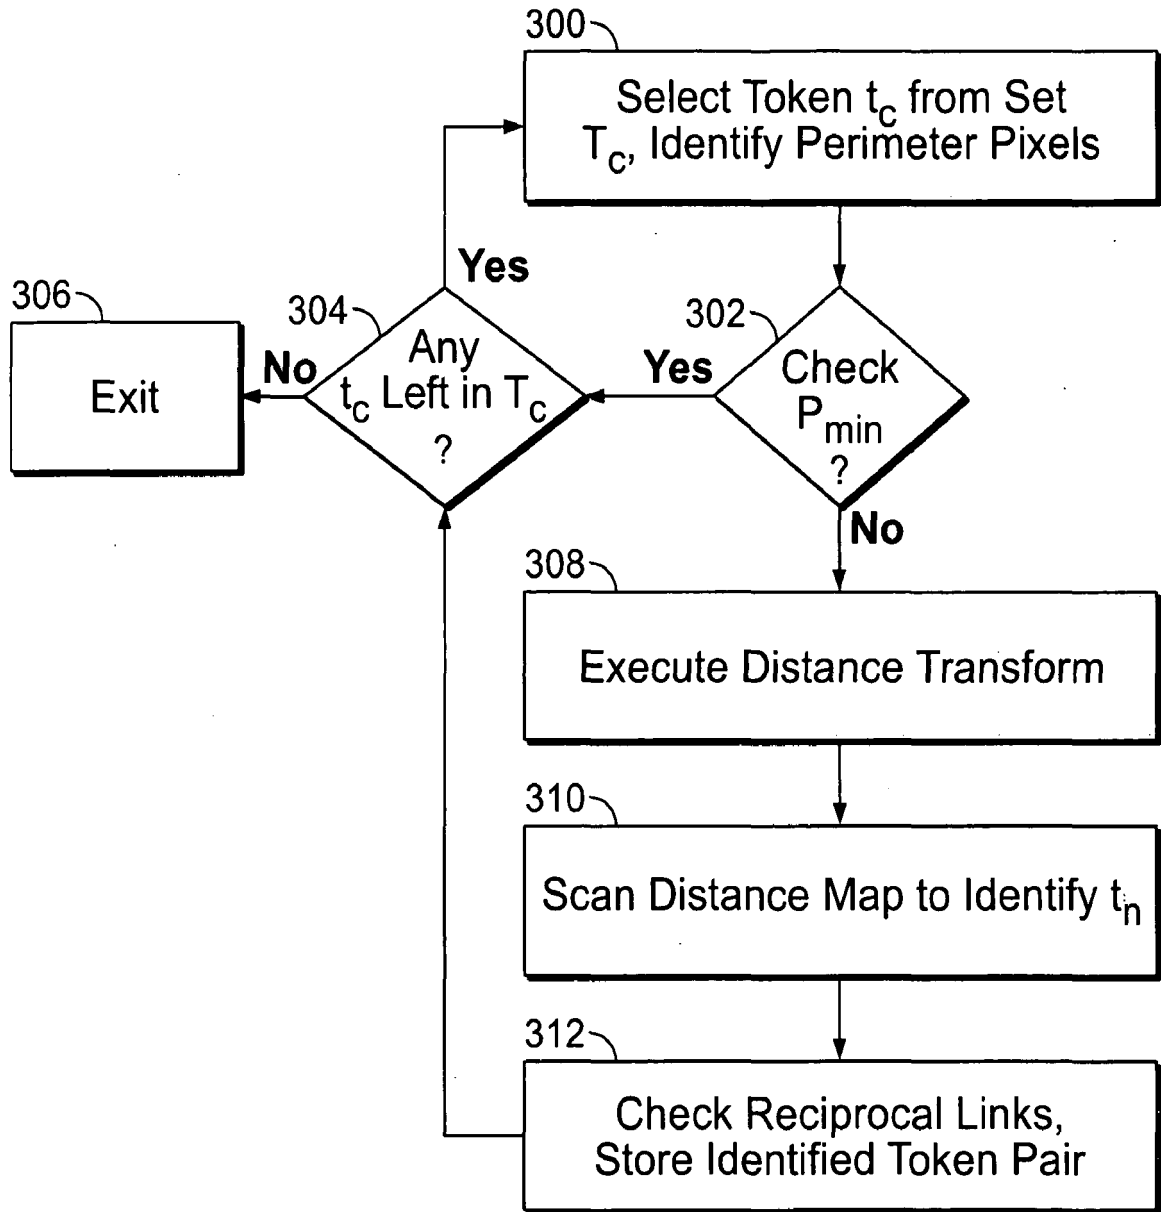


FIG. 6

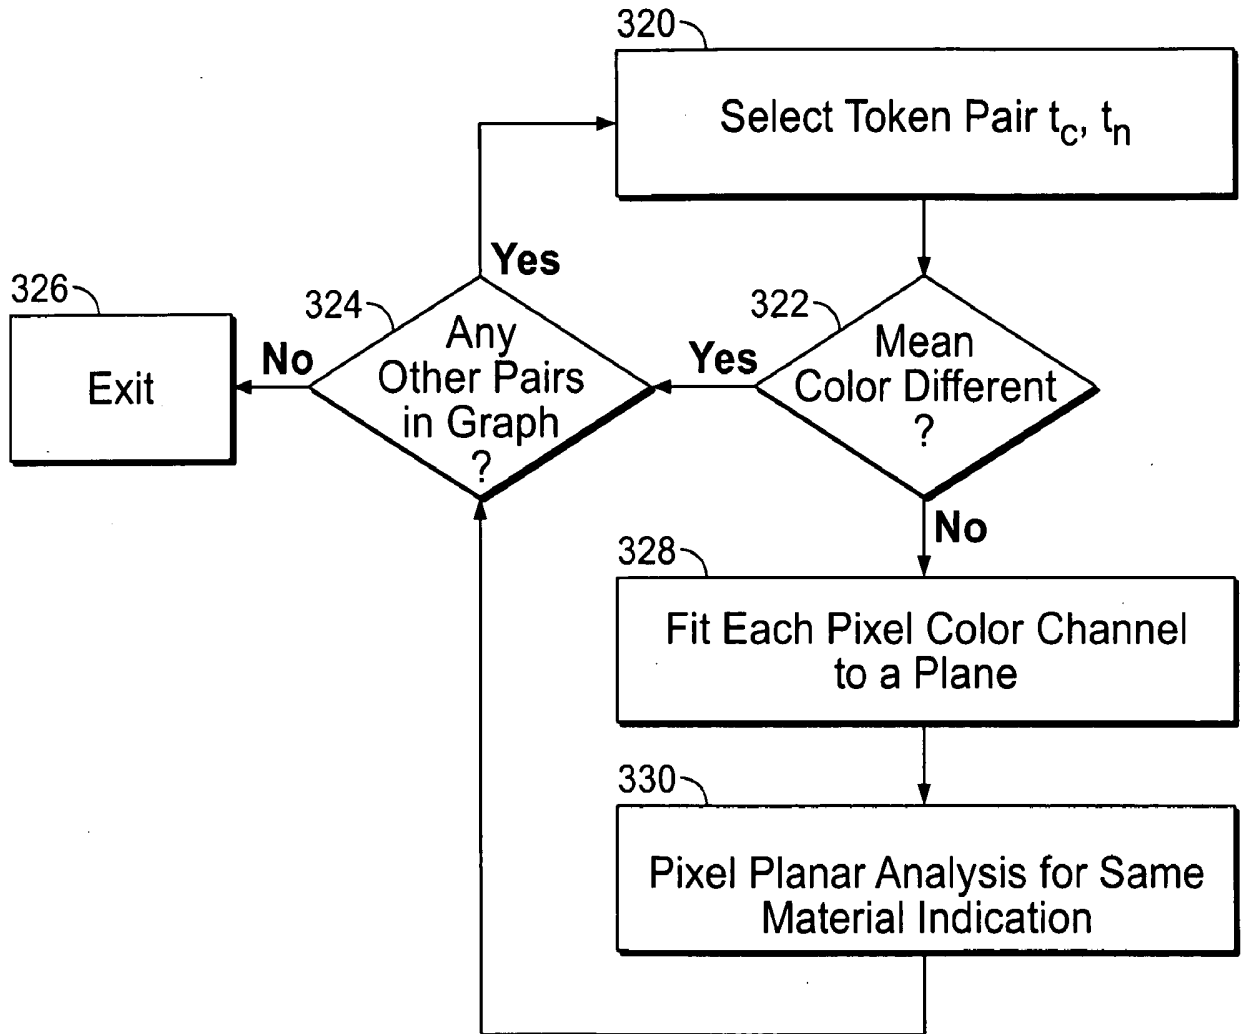


FIG. 7

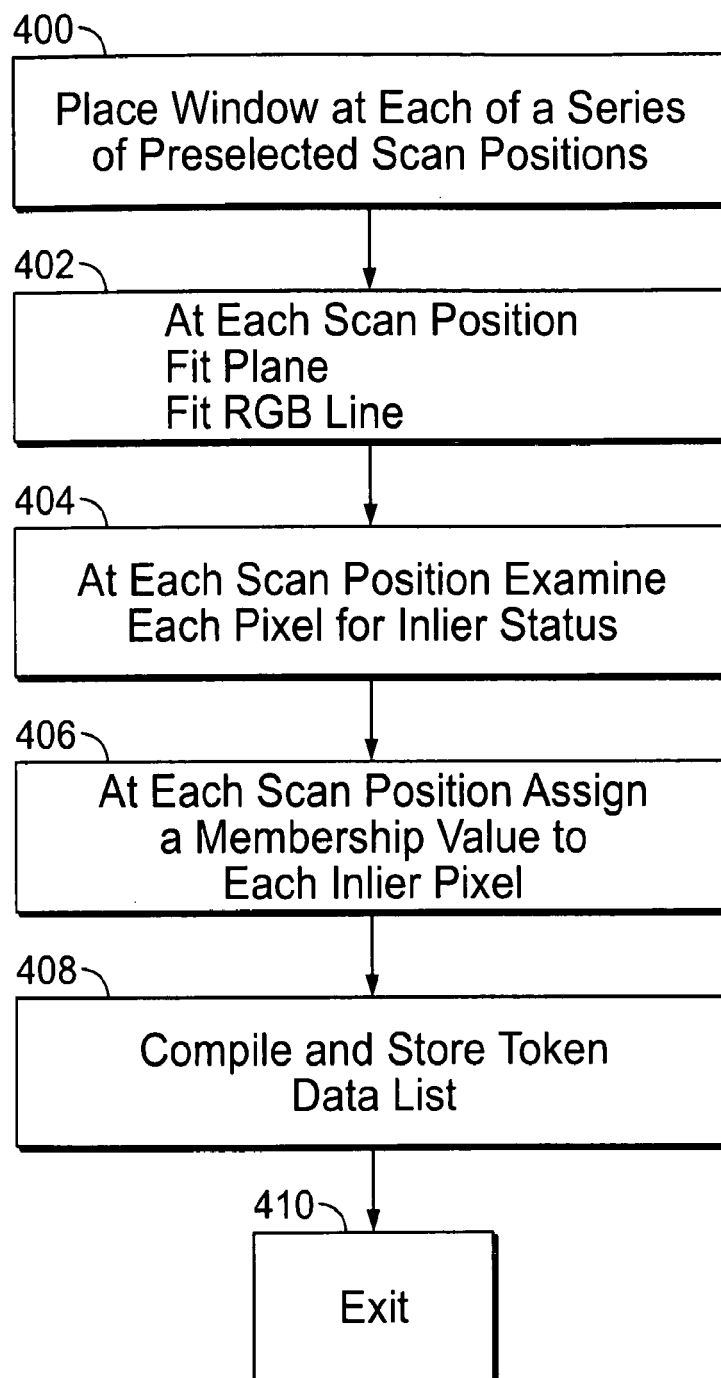


FIG. 8

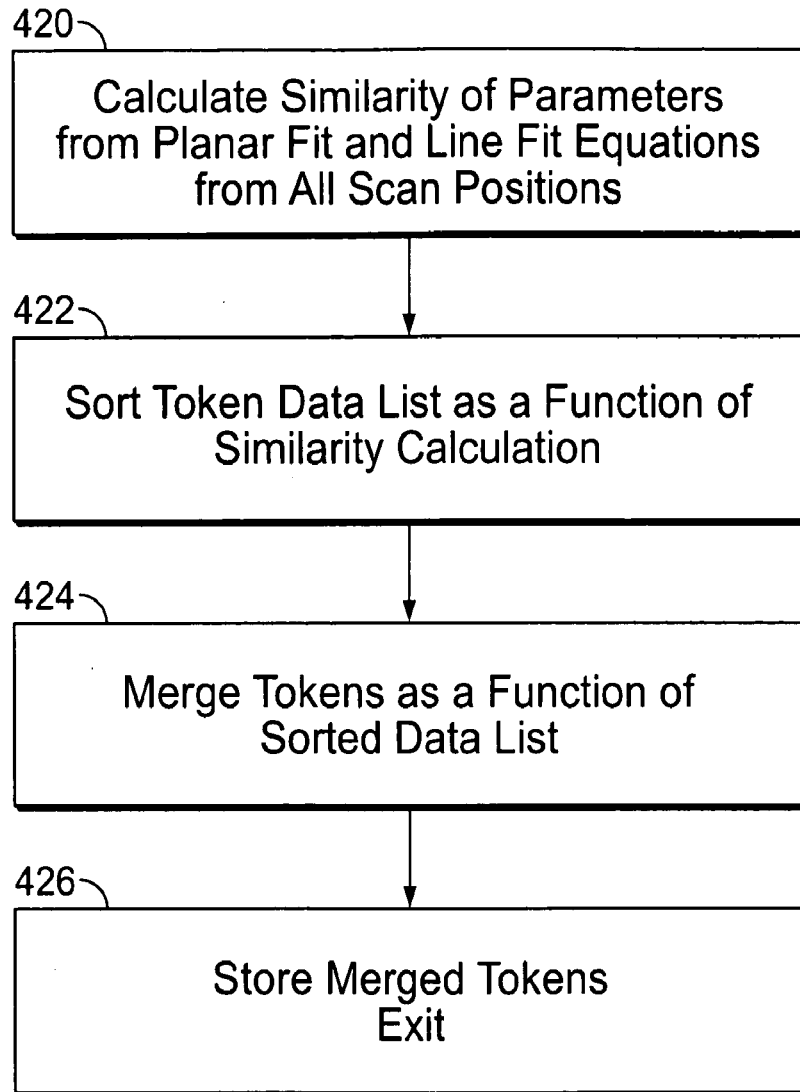


FIG. 9

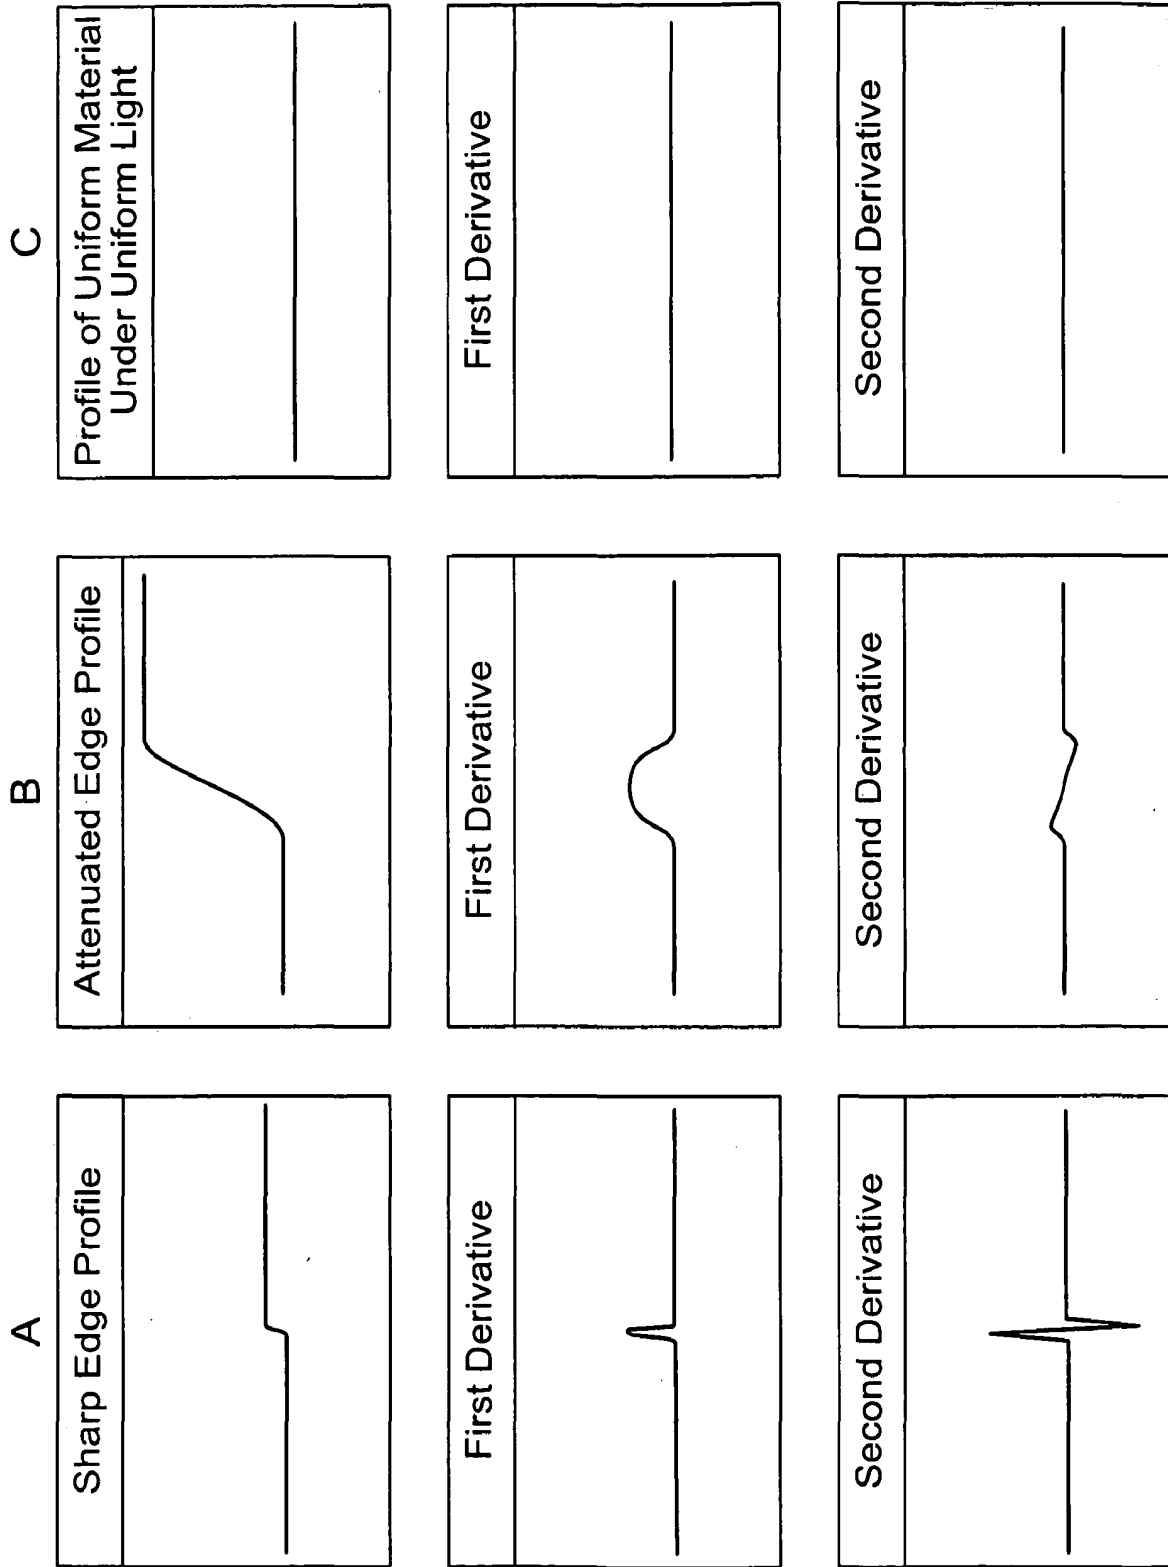


FIG.10A

0.00031426	0.002632078	0.008579021	0.012376362	0.008579021	0.002632078	0.00031426
0.002632078	0.017490147	0.039192701	0.043078559	0.039192701	0.017490147	0.002632078
0.008579021	0.039192701	0	-0.096532354	0	0.039192701	0.008579021
0.012376362	0.043078559	-0.096532354	-0.318309892	-0.096532354	0.043078559	0.012376362
0.008579021	0.039192701	0	-0.096532354	0	0.039192701	0.008579021
0.002632078	0.017490147	0.039192701	0.043078559	0.039192701	0.017490147	0.002632078
0.00031426	0.002632078	0.008579021	0.012376362	0.008579021	0.002632078	0.00031426

FIG. 10B

P(1,1)	P(1,2)	P(1,3)	P(1,4)	P(1,5)	P(1,6)	P(1,7)
P(2,1)	P(2,2)	P(2,3)	P(2,4)	P(2,5)	P(2,6)	P(2,7)
P(3,1)	P(3,2)	P(3,3)	P(3,4)	P(3,5)	P(3,6)	P(3,7)
P(4,1)	P(4,2)	P(4,3)	P(4,4)	P(4,5)	P(4,6)	P(4,7)
P(5,1)	P(5,2)	P(5,3)	P(5,4)	P(5,5)	P(5,6)	P(5,7)
P(6,1)	P(6,2)	P(6,3)	P(6,4)	P(6,5)	P(6,6)	P(6,7)
P(7,1)	P(7,2)	P(7,3)	P(7,4)	P(7,5)	P(7,6)	P(7,7)

FIG. 10C

Generalized Identification of Regions of Uniform Reflectance

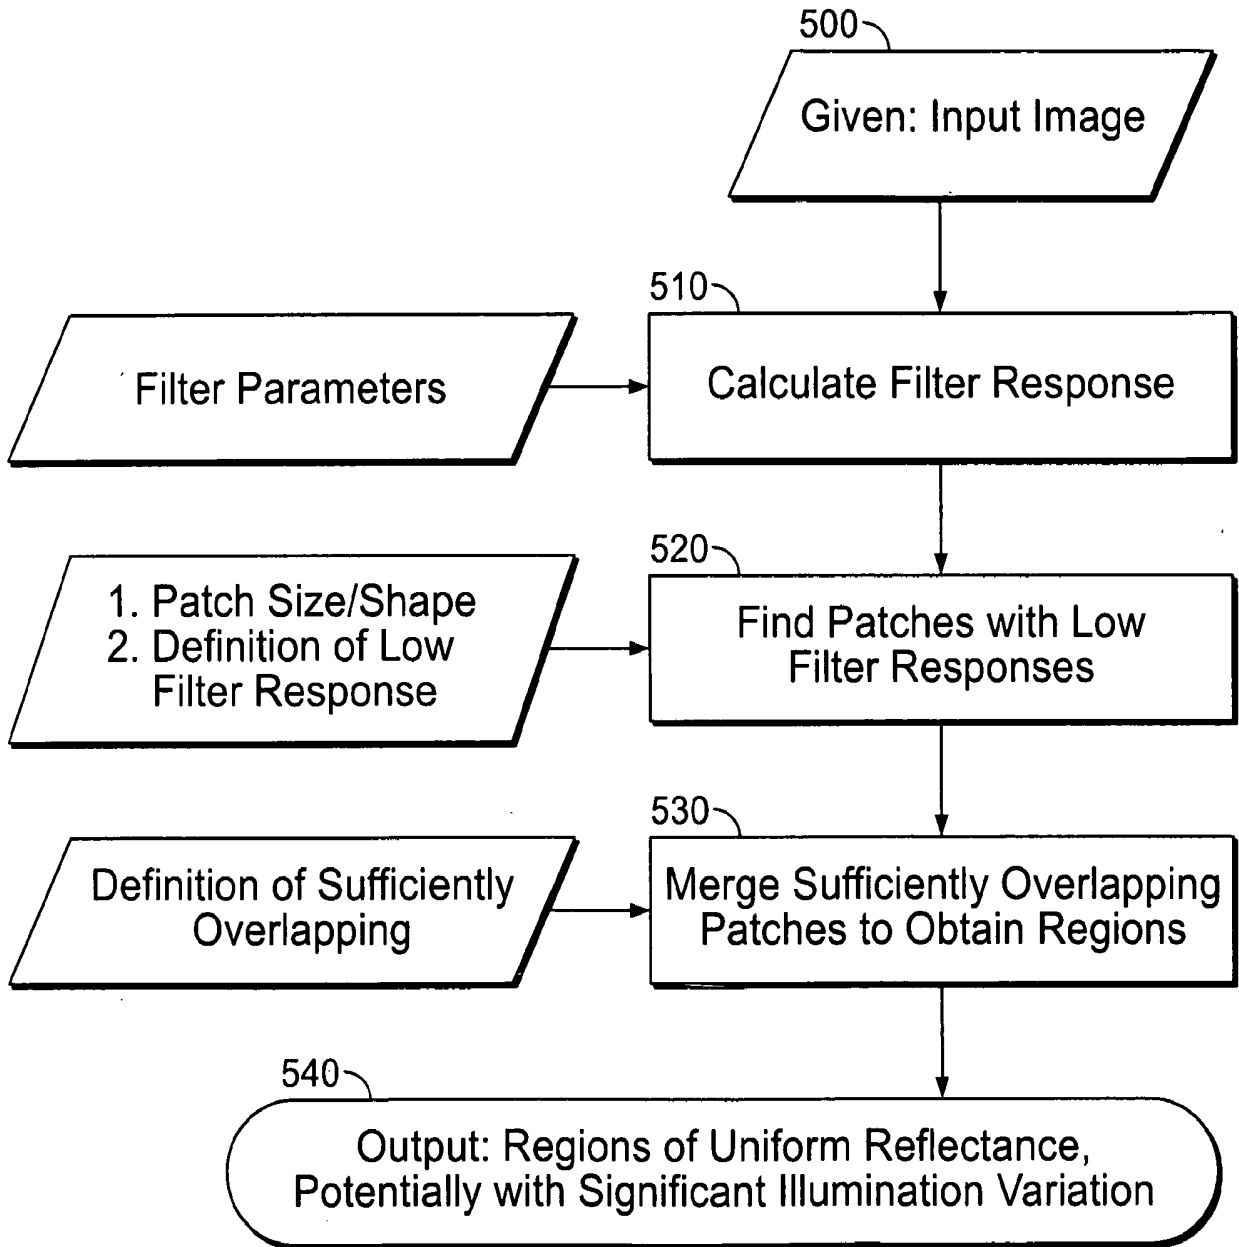


FIG. 11

Combining Results

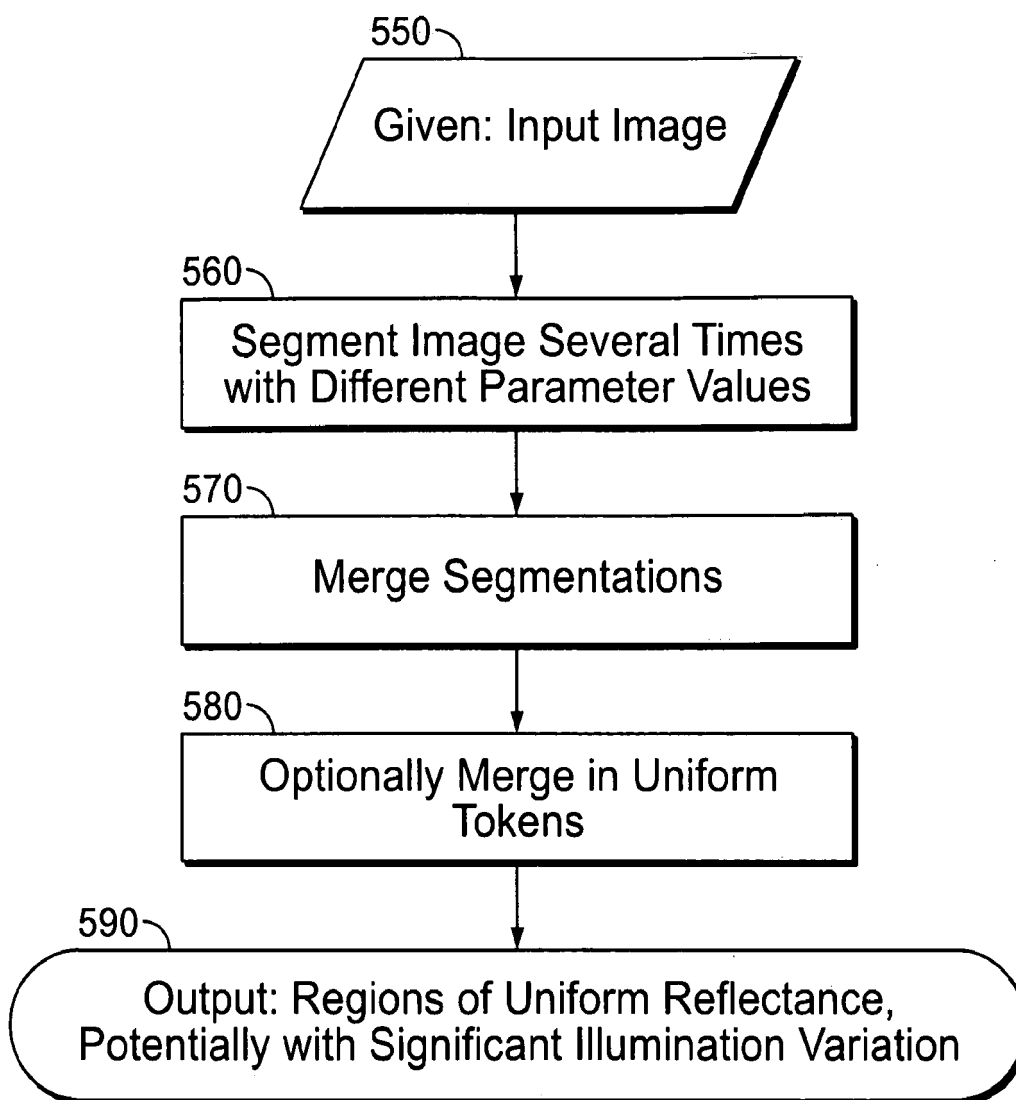


FIG. 12

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 20070177797 A1 [0002]
- US 11341753 B [0002]

Non-patent literature cited in the description

- Low-Level Segmentation of Multispectral Images via Agglomerative Clustering of Uniform Neighbourhoods. **AMADASUN M et al.** Pattern Recognition. Elsevier, 01 January 1988, vol. 21, 261-268 [0003]
- **RANDEN, T.[TRYGVE ; HUSØY, J.H.[JOHN HÅKON.** *Filtering for Texture Classification: A Comparative Study, PAMI*, April 1999, vol. 21 (4), 291-310 [0017]
- **P. FELZENSZWALB ; D. HUTTENLOCHER.** *Distance Transforms of Sampled Functions, Cornell Computing and Information Science Technical Report TR2004-1963*, September 2004 [0044]
- **ABDI, H.** *Encyclopedia of Measurement and Statistics*, Thousand Oaks. 2007 [0048]
- **PETER J. ROUSSEEUW.** Least Median of Squares Regression. *Journal of the American Statistical Association*, December 1984, vol. 79 (388), 871-880 [0050]
- **ZVI GALIL ; GIUSEPPE F. ITALIANO.** Data structures and algorithms for disjoint set union problems. *ACM Computing Surveys*, September 1991, vol. 23 (3), 319-344 [0052]
- **M. A. FISCHLER ; R. C. BOLLES.** Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Comm. of the ACM*, 1981, vol. 24, 381-395 [0056]