

(19)



(11)

EP 2 435 956 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention of the grant of the patent:

06.03.2019 Bulletin 2019/10

(51) Int Cl.:

G06T 7/11 ^(2017.01)

G06T 7/90 ^(2017.01)

(86) International application number:

PCT/US2010/001472

(21) Application number: **10780914.7**

(22) Date of filing: **19.05.2010**

(87) International publication number:

WO 2010/138163 (02.12.2010 Gazette 2010/48)

(54) MULTI-RESOLUTION ANALYSIS IN AN IMAGE SEGREGATION

ANALYSE MIT MEHRFACHER AUFLÖSUNG IN EINER BILDAUFSPALTUNG

ANALYSE À RÉOLUTIONS MULTIPLES DANS UNE SÉGRÉGATION D'IMAGE

(84) Designated Contracting States:

**AL 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 SM TR**

(30) Priority: **25.05.2009 US 454926**

(43) Date of publication of application:

04.04.2012 Bulletin 2012/14

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

(72) Inventors:

- **TOLLIVER, David, Allen
Pittsburgh
PA 15201 (US)**
- **DANA, Kristin, Jean
Spring Lake
NJ 07762 (US)**
- **STEIN, Andrew, Neil
Pittsburgh
PA 15207 (US)**
- **ALLDRIN, Neil
Pittsburgh
PA 15217 (US)**

(74) Representative: **Reichert & Lindner**

**Partnerschaft Patentanwälte
Bismarckplatz 8
93047 Regensburg (DE)**

(56) References cited:

**US-A1- 2003 161 516 US-A1- 2007 009 176
US-B1- 6 424 732**

- **BRENDAN MOORE ET AL: "Learning Face Appearance under Different Lighting Conditions", BIOMETRICS: THEORY, APPLICATIONS AND SYSTEMS, 2008. BTAS 2008. 2ND IEEE INTERNATIONAL CONFERENCE ON, IEEE, PISCATAWAY, NJ, USA, 29 September 2008 (2008-09-29), pages 1-8, XP031371238, ISBN: 978-1-4244-2729-1**
- **TAPPEN M F ET AL: "Estimating Intrinsic Component Images using Non-Linear Regression", CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION, 2006 IEEE COMPUTER SOCIETY, NEW YORK, NY, USA 17-22 JUNE 2006, IEEE, PISCATAWAY, NJ, USA, vol. 2, 17 June 2006 (2006-06-17), pages 1992-1999, XP010923098, DOI: 10.1109/CVPR.2006.114 ISBN: 978-0-7695-2597-6**
- **MAKROGIANNIS S ET AL: "Scale space segmentation of color images using watersheds and fuzzy region merging", PROCEEDINGS 2001 INTERNATIONAL CONFERENCE ON IMAGE PROCESSING. ICIP 2001 - THESSALONIKI, GREECE, OCT. 7 - 10, 2001; [INTERNATIONAL CONFERENCE ON IMAGE PROCESSING], INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, NEW YORK, NY, vol. 1, 7 October 2001 (2001-10-07), pages 734-737, XP010564964, DOI: 10.1109/ICIP.2001.959150 ISBN: 978-0-7803-6725-8**

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).

EP 2 435 956 B1

DescriptionBackground of the Invention

[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 and illumination characteristics of the image.

[0002] The article "Learning Face Appearance under Different Lighting Conditions" by Brendan Moore et al. in *Bio-metrics: Theory, Applications and Systems*, 2008, BTAS 2008, IEEE, Piscataway, NJ, USA, discusses a machine learning approach for estimating intrinsic faces and hence de-illuminating and re-illuminating faces directly in the image domain. An observed image is decomposed into its intrinsic components using multi-scale patches of the observed image from a three-level Laplacian pyramid.

[0003] The article "Estimating Intrinsic Component Images using Non-Linear Regression" by Marshall F Tappen et al, *Conference on Computer Vision and Pattern Recognition 2006*, vol 2, pp. 1992 - 1999, IEEE, Piscataway, NJ, USA presents a method for estimating intrinsic component images from a single image based on estimators predicting filtered versions of the image. The estimators are used to obtain a set of local linear constraints from local image data.

Summary of the Invention

[0004] The present invention provides a method and system comprising image processing techniques that utilize spatio-spectral information relevant to an image, derived from at least one representation from a set of selectively varied representations of the image, for example, multi-resolutions such as a scale-spaced pyramid of representations, to accurately and correctly identify illumination and material aspects of the image.

[0005] In a first exemplary embodiment of the present invention, an automated, computerized method according to claim 1 is provided for processing an image. According to a feature of the present invention, the method comprises the steps of providing an image file depicting the image, in a computer memory, forming a set of selectively varied representations of the image file and performing an image segregation operation on at least one preselected representation of the image of the image file, to generate intrinsic images corresponding to the image. According to a feature of the exemplary embodiment of the present invention, the selectively varied representations comprise multi-resolution representations such as a scale-spaced pyramid of representations. In a further feature of the exemplary embodiment of the present invention, the intrinsic images comprise a material image and an illumination image.

[0006] In a second exemplary embodiment of the present invention, a computer system according to claim 10 is provided. The computer system comprises a CPU and a memory storing an image file containing an image. According to a feature of the present invention, the CPU is arranged and configured to execute a routine to form a set of selectively varied representations of the image and to perform an image segregation operation on at least one preselected representation of the image of the image file, to generate intrinsic images corresponding to the image.

[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]

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

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

5 Figure 2b is an 11×11 pixel array formed from the upper left hand corner of the image file of figure 2a, for use in the generation of an histogram, according to a feature of the present invention.

Figure 2c shows a diagram of the histogram generated from the 11×11 pixel array of figure 2b.

10 Figure 3a is a data flow block diagram for an image segregation operation performed according to an exemplary embodiment of the present invention including a set of selectively varied representations of the image of the image file of figure 2a.

15 Figure 3b shows a graphic representation of a scale-spaced Gaussian pyramids the set of selectively varied representations of figure 3a.

Figure 4 is a functional block diagram for an image segregation system architecture for performing the data flow of figure 3a, implemented in the computer system of figure 1, according to a feature of the present invention.

20 Figure 5 shows a graphical user interface for use in connection with an implementation of the image segregation system architecture feature of the present invention.

Figure 6a is a flow chart for identifying Type C token regions in the image file of figure 2a, according to a feature of the present invention.

25

Figure 6b is an original image used as an example in the identification of Type C tokens.

Figure 6c shows Type C token regions in the image of figure 6b.

30 Figure 6d shows Type B tokens, generated from the Type C tokens of figure 6c, according to a feature of the present invention.

Figure 7 is a flow chart for a routine to test Type C tokens identified by the routine of the flow chart of figure 6a, according to a feature of the present invention.

35

Figure 8 is a flow chart for constructing Type B tokens via an arbitrary boundary removal technique, according to a feature of the present invention.

40 Figure 9 is a flow chart for creating a token graph, containing token map information, according to a feature of the present invention.

Figure 10 is a flow chart for constructing Type B tokens via an adjacent planar token merging technique, according to a feature of the present invention.

45 Figure 11 is a flow chart for generating Type C tokens via a local token analysis technique, according to a feature of the present invention.

Figure 12 is a flow chart for constructing Type B tokens from Type C tokens generated via the local token analysis technique of figure 11, according to a feature of the present invention.

50

Figure 13 is a representation of an $[A] [x] = [b]$ matrix relationship according to a feature of the present invention.

Figure 14 is a functional block diagram for a service provider component for use in the image segregation system architecture of figure 4.

55

Detailed Description of the Preferred Embodiments

[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.

5 **[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.

10 **[0011]** As shown in figure 2a, each image file 18 comprises an $n \times m$ 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,

15 as a function of the digital values of the pixels in the image file 18, as is generally known.
[0012] According to a feature of the present invention, in an image process, the CPU 12 operates to analyze and process information, for example, the RGB values of the pixels of an image stored in an image file 18, to achieve various objectives, such as, for example, a segregation of illumination and material aspects of a scene depicted in the image of an image file 18, to generate corresponding intrinsic images. The intrinsic images can include, for example, a material image and an illumination image.

20 **[0013]** To that end, figure 3a shows a data flow diagram for an image segregation operation performed according to an exemplary embodiment of the present invention. In block 1000, an image file 18 is selected and input to the CPU 12. In block 1002, the CPU 12 operates to form a set of selectively varied representations of the image, depicted in image files 18a, 18b, ... 18n, from the selected image file 18. The variations among the representations can be based upon any preselected spacial, spectral and/or temporal groupings of the original image. The variations are selected so as to provide image representations most suitable for efficient and accurate analysis of illumination and material aspects of the image.

25 **[0014]** In one preferred embodiment of the present invention, computer operation for illumination and material image segregation can be performed at a relatively coarse scale representation of the image, and then used to constrain a solution at a relatively finer scale resolution. By enforcing a segregation that is consistent across different scales of resolution, performance of a segregation operation is not only more speedy and efficient, but also more accurate.

30 **[0015]** In another preferred embodiment, a segregation of illumination and material aspects of the image is selectively divided among representations at differing scales of image resolution. Different features of a scene depicted in an image can have different intrinsic scales. For example, edges in an image have an intrinsic scale most effectively analyzed at a fine scale of resolution, while non-edged regions, such as regions of uniform material, can be accurately analyzed at relatively coarse scales of resolution. Thus, an image is divided by edge and non-edge regions, segregating the edge regions at a fine scale of resolution, and the remaining non-edge regions at a relatively coarse scale of resolution. More generally, pixels of the image (or other preselected regions of the image) can each be divided according to the intrinsic scale of the respective pixel, and processed at a representation of the image at an appropriate scale. The division of segregation operations in this manner provides a sequence of simple solution systems, rather than one relatively large complex system to solve.

35 **[0016]** Accordingly, for example, in one exemplary embodiment of the present invention, the CPU 12 can form multi-resolution representations such as a scale-spaced pyramid of representations of the image. As shown in figure 3b, in a scale-spaced pyramid, the set of image files 18a, b, c ... n, are generated, each at a different scale of resolution relative to the resolution of the original image, from a finest resolution (defined as the resolution of the original image in image file 18) to relatively coarser resolutions, at each upper level of the pyramid. Each of image files 18, 18a, ... 18n of figure 3b is an array of pixels, as generally shown in figure 2a. Image file 18a is a coarser resolution version of the original image file 18, and image file 18b is a coarser resolution version of image file 18a, and so on. For example, as illustrated in figure 3b, pixels $P(5, 3)$, $P(6, 3)$, $P(5, 4)$ and $P(6, 4)$ of fine resolution image file 18 map to, and therefor correspond to, a single pixel, $P_a(3, 2)$ in the relatively coarse resolution image file 18a.

40 **[0017]** A scale-spaced pyramid(s) is formed via well known Gaussian and/or Laplacian pyramid formation techniques. In a Gaussian pyramid, the original image file 18 is convolved with a Gaussian kernel g , for a low pass filtered, subsampled version of the original image (image file 18a), having a cut off frequency set by the parameter σ . A full Gaussian pyramid is formed when each coarse scale (e.g. image file 18a) is convolved with the Gaussian kernel to form a next level (image file 18b), and so on (to image file 18n). A Laplacian pyramid provides a series of bandpass filtered versions of the original image, wherein each scale of the pyramid corresponds to a different band of image frequencies.

45 **[0018]** If a series of selectively varied images represented by image files 18, 18a, ... 18n is stated as a Gaussian pyramid of low passed filtered versions of the original image, and expressed by original image i_0 , and scales i_1, i_2, \dots, i_n , respectively, then any level of the pyramid is related to a downsampled next lower (finer) level by: $i_n = \downarrow(g * i_{n-1})$, where

↓ denotes a downsampling, g represents a low pass filter, such as, for example, the Gaussian low pass filter and $*$ denotes a spatial convolution.

[0019] For a Laplacian pyramid, in the series of images represented by image files 18, 18a, ...18n, each level of the pyramid I_n can be expressed in terms of a difference between levels of the Gaussian pyramid: $I_n = I_n - \uparrow I_{n+1}$, where \uparrow denotes upsampling.

[0020] Moreover, starting with a multiplicative model for the image: $I = M \times L$, where I is the original image, M is the material component of the image and L is the illumination component, it can be shown that each level of the pyramid can be decomposed into a material pyramid and an illumination pyramid. For an additive model, the log color values are $\log I = \log M + \log L$. In terms of the Gaussian pyramid, the decomposition into material and illumination components can be expressed as: $g_\sigma * \log I = g_\sigma * \log M + g_\sigma * \log L$.

[0021] In block 1004, the CPU 12 operates to perform an image segregation operation on at least one of the representations of the pyramid, for example, original image file 18 and image file 18a, to generate intrinsic images including, for example, each of an illumination image and a material image for each selected scale of the image. Each of the intrinsic images corresponds to the image depicted at the resolution scales of image file 18 and image file 18a, respectively. According to a feature of the present invention, the image segregation operation performed by the CPU 12 at each scale of resolution, is implemented according to a spatio-spectral operator/constraint/solver model to identify the illumination and material components of the original image or any representation of the original image.

[0022] 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. Spatio-spectral information is information relevant to contiguous pixels of an image depicted in an image file 18, such as spectral relationships among contiguous pixels, in terms of color bands, for example RGB values of the pixels, and the spatial extent of the pixel spectral characteristics relevant to a characteristic of the image, such as, for example, a single material depicted in the image or illumination flux effecting the image. When one of material and illumination is known in an image, the other can be readily deduced.

[0023] Moreover, the illumination flux includes an incident illuminant and an ambient illuminant. The spectra for the incident illuminant and the ambient illuminant can be different from one another. Thus, a spectral shift is caused by a shadow, i.e., a decrease of the intensity of the incident illuminant. The spectral shift can cause a variance in color of material depicted in the scene, from full shadow, through the shadow penumbra, to fully lit. Pursuant to a feature of the present invention, spectral shift phenomena is captured in spatio-spectral information. The spatio-spectral information includes a spectral ratio: a ratio based upon a difference in color or intensities between two areas of a scene depicted in an image, which may be caused by different materials (an object edge), an illumination change (illumination boundary) or both.

[0024] According to a further feature of the present invention, spatio-spectral operators are generated to facilitate the image segregation process performed in step 1004, for the segregation of illumination and material aspects of a scene depicted in the retrieved image file 18, and any other representations of the image, such as image file 18a. Spatio-spectral operators comprise representations or characteristics of an image that encompass spatio-spectral information usable in the processing of material and illumination aspects of an image. The spatio-spectral operators are subject to constraints that define constraining spatio-spectral relationships between the operators, for input to a solver. The operators can be identified at any one of the selectively varied representations of the original image, and the constraints defined in terms of operators within a single representation or between operators in different representations.

[0025] A constraint can be imposed among the different levels of the pyramid, for example, to improve and expedite segregation at the finer scale (image file 18). To that end, as described above, a segregation into illumination and material is first quickly and efficiently performed by the CPU 12 (according to the spatio-spectral operator/constraint/solver model, described in detail below) for the relatively coarse image (for example, image file 18a). Thereafter, for more consistent and accurate results, during segregation at the finer scale, material color at each local area of the relatively finer scale image (for the material image generated for image file 18) is constrained such that a weighted average of the color at a local area of the fine scale equals the color of a corresponding local area of the coarse scale.

[0026] Such a weighted average constraint can be expressed, in a simplified analysis for a three pixel local area of the fine scale representation, as follows: if the log intensities are related as: $i_2(x, y) = 0.5 * i_1(x, y) + 0.25 * i_1(x-1, y) + 0.25 * i_1(x+1, y)$, wherein i_2 (corresponding, for example, to the log of image file 18a) is a blurred version of i_1 (corresponding to the log of image file 18); then the log material values are related as: $m_2(x, y) = 0.5 * m_1(x, y) + 0.25 * m_1(x-1, y) + 0.25 * m_1(x+1, y)$, wherein m_1 and m_2 are the material components of i_1 and i_2 , respectively. Typically, the analysis is performed for a local area of the fine scale that is larger than the area corresponding to the pixel of the coarse scale representation, to avoid aliasing. For example, in pyramid wherein the coarse scale has half the pixels of the fine scale, a 5 X 5 pixel square would be used to calculate the weighted average for the color at the fine scale, and then constrained to be the color of a single pixel in the coarse scale representation.

[0027] In the exemplary embodiment of the present invention, the solver includes a mathematical processing engine that operates to obtain an optimized solution for the generation of an intrinsic image, such as a material image and/or

an illumination image derived from the original image stored in the retrieved image file 18 or any other representation of the image, as a function of the constraining relationships between the spatio-spectral operators.

[0028] Spatio-spectral operators include, for example, tokens, token map information, log chromaticity representation values, X-junctions, BIDR model representations, a boundary representation, and a texon histogram based pixel representation. Each of the spatio-spectral operators can be computed and applied at the image file 18, and/or at any of the image files 18a, ...18n, and any solution at any level can be related to any solution at any other level each via the upsampling and downsampling relationships described above.

[0029] 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 use of tokens recognizes the fact that a particular set of material/illumination/geometric characteristics of an image extends beyond a single pixel, and therefore, while the image processing described herein can be done on a pixel level, tokens expedite a more efficient processing of image properties. The pixels of a token can be related in terms of either homogeneous factors, such as, for example, close correlation of color values among the pixels, or nonhomogeneous factors, such as, for example, differing color values related geometrically in a color space such as RGB space, commonly referred to as a texture.

[0030] Exemplary embodiments of the present invention provide methods and systems to identify various types of homogeneous or nonhomogeneous 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, or contiguous pixels of any selectively varied representation (image files 18a, ...18n) to identify token regions within the original image and/or within any of the variations thereof. Such tokens identified in different selectively varied representations of the image can be used in, for example, the weighted average color constraint discussed above.

[0031] According to one exemplary embodiment of the present invention, homogeneous 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, for example, similar color and intensity, where similarity is defined with respect to a noise model for the imaging system used to record the image. The weighted average constraint discussed above, can be implemented at a Type C token level, constraining the weighted average color of a token(s) at the fine scale to the color of a corresponding token(s) at a coarse scale.

[0032] A linear token is a nonhomogeneous token comprising a connected region of the image wherein adjacent pixels of the region have differing color measurement values that fall within a cylinder in RGB space, from a dark end (in shadow) to a bright end (lit end), along a positive slope. The cylinder configuration is predicted by a bi-illuminant dichromatic reflection model (BIDR model), according to a feature of the present invention, when the color change is due to an illumination change forming a shadow (i.e. a decrease in the intensity of the incident illuminant as the interplay between the incident or direct illuminant and the ambient illuminant in the illumination field) over a single material of a scene depicted in the image.

[0033] For purposes of describing, identifying and using linear tokens, the BIDR model can be stated as: $I_{(x,y,z,\theta,\phi,\lambda)} = c_b(\lambda) I_d(\lambda) \gamma_b + M_a(\lambda) c_b(\lambda)$, where: $I_{(x,y,z,\theta,\phi,\lambda)}$ is the radiance of a surface point at (x, y, z) in the direction θ, ϕ for the wavelength λ , $c_b(\lambda)$ is the geometry independent body reflectance of a surface for the wavelength λ , $I_d(\lambda)$ is the incident illuminant for the wavelength λ , γ_b is the product of a shadow factor $s_{x,y,z}$ and a geometric factor $m_b(\theta_i)$, and $M_a(\lambda)$ is the integral of the ambient illuminant and geometric body reflectance over a hemisphere, excluding the incident illuminant. For more detailed information on the BIDR model, reference should be made to U. S. Application Serial No. 11/341,751, filed January 27, 2006, entitled: "Bi-illuminant Dichromatic Reflection Model For Image Manipulation," published as US 2007/0176940 on August 2, 2007.

[0034] Token map information indicates locations of tokens within an image, relative to one another. The map information is used to identify neighboring tokens for performing an analysis of token neighbor relationships relevant to constraining spatio-spectral relationships between tokens, for input to the solver. A token map can be generated for each selectively varied representation of the original image in respect of tokens identified at each respective variation.

[0035] Log chromaticity representation values provide illumination invariant values for pixels of the image. Logarithmic values of the color band values of the image pixels are plotted on a log-color space graph. The logarithmic values are then projected to a log-chromaticity projection plane oriented as a function of the BIDR model. The chromaticity plane values are substituted for the color band values (for example, RGB values) of each pixel. For more detailed information on log chromaticity representation values, reference should be made to U. S. Application Serial No. 11/403,719, filed April 13, 2006, entitled: "Method And System For Separating Illumination And Reflectance Using a Log Color Space," published as US 2007/0242878 on October 18, 2007.

[0036] An X-junction is an area of an image where a material edge and an illumination boundary cross one another.

An X-junction is an optimal location for an accurate analysis of material and illumination aspects of an image.

[0037] A boundary representation is an arrangement of pixels, on each side of a boundary, formed by, for example, adjacent Type B tokens. The arrangement is used to facilitate an analysis of the boundary to classify the boundary as a material boundary on a smooth surface (as opposed to another type of boundary, for example, an illumination edge, depth boundary or simultaneous illumination and material change). The pixel representation is configured to provide samples of pixels within each of the Type B tokens forming the boundary. The pixels of the samples are subject to spatio-spectral analysis, and the results are compared to determine the likelihood that the respective boundary corresponds to a material change.

[0038] A texton is a homogeneous representation for a region of an image that comprises a texture. Image texture can be defined as a function of spatial variation in pixel intensities. Image texture patterns are frequently the result of physical or reflective properties of the image surface. Commonly, an image texture is associated with spatial homogeneity and typically includes repeated structures, often with some random variation (e.g., random positions, orientations or colors). Image textures are also often characterized by certain visual properties such as regularity, coarseness, contrast and directionality. An example of image texture is the image of a zebra skin surface as it appears to be spatially homogeneous and seems to contain variations of color intensities which form certain repeated patterns. Some image textures can be defined by geometric characteristics, such as stripes or spots. A texton based operator, such as a texture token, as will be described, transforms patterns of differing reflectance caused by a textured material into a homogeneous representation that captures the spectral and spatial characteristics of the textured region in the image.

[0039] Constraints between spatio-spectral operators comprise, for example, an anchor constraint, a same illumination constraint, a smooth illumination constraint, a Type B token constraint, a Linear token constraint, a BIDR enforcement constraint, a same texton histogram constraint, a log chromaticity similarity constraint, an X junction constraint, and a boundary representation constraint. Other constraints can be formed as between spatio-spectral operators, or other indicia, from different ones of the selectively varied representations of the image, such as, for example, the weighted average color constraint described above. Each constraint is configured as a constraint generator software module that defines the spatio-spectral operators utilized by the respective constraint and provides an expression of the constraining relationship imposed upon the constituent operators.

[0040] An anchor constraint utilizes a number of brightest/largest Type C tokens in an image. The constraining relationship is that the material of the selected brightest/largest Type C tokens is constrained to be an absolute value for the color/brightness observed in the image. The constraint anchors a material map for the image at an absolute brightness to avoid relative brightness constraints.

[0041] A same illumination constraint utilizes Type C tokens and Type B tokens identified in an image and token map information. The constraining relationship is that adjacent Type C tokens, as indicated by the token map information, are at the same illumination, unless the adjacent Type C tokens are part of the same Type B token. The term "same" in connection with the term "illumination" is used to mean an average value with respect to a noise model for the imaging system used to record the image. This constrains any observed differences in appearance between adjacent Type C tokens, that are not part of the same Type B token, to be a material change, as will appear.

[0042] A smooth illumination constraint is similar to the same illumination constraint. However, rather than constraining all pixels of adjacent Type C tokens to be of the same illumination, as in the same illumination constraint, in the smooth illumination constraint, the constraint is based upon the average illumination of the pixels near a shared boundary between adjacent Type C tokens. This constrains the illumination field to be somewhat smooth, as opposed to piecewise constant (the same, as defined above) throughout a token.

[0043] A Type B token constraint also utilizes Type C tokens and Type B tokens. However, the constraining relationship is that all Type C tokens that are part of the same Type B token are constrained to be of the same material. This constraint enforces the definition of a Type B token, that 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. Thus, all Type C tokens that lie within the same Type B token are by the definition imposed upon Type B tokens, of the same material, though not necessarily of the same illumination. The Type C tokens are therefore constrained to correspond to observed differences in appearance that are caused by varying illumination.

[0044] Accordingly, the Type B token constraint is complementary to the same and smooth illumination constraints, which, as opposed to illumination change, constrain observed differences to correspond to material change, as described above. This is due to the fact that in each of the same and smooth illumination constraints, Type C tokens that are adjacent and not part of the same Type B token, are constrained to the same illumination. These Type C tokens should comprise different materials, since by the constraint, they are not in the same Type B token and therefore, by the definition of Type B tokens enforced by the constraint, do not encompass a single material, so illumination should be a constant, and any observed difference is considered as attributable to a material change.

[0045] To summarize, pursuant to a feature of the present invention, the Type C and Type B token spatio-spectral operators are defined to provide characteristics of an image that enable segregation of illumination and material. Type C tokens each comprise a connected image region of similar image properties, for example similar color, as recorded

and stored in an image file 18. Thus, adjacent Type C tokens indicate some form of change in the image or else they would form the same Type C token. Type B tokens encompass a single material. The complementary constraints of the same/smooth illumination constraints and the Type B token constraint enforce relationships between the tokens that indicate either a material change or an illumination change.

[0046] If the adjacent Type C tokens are within the same Type B token, as in the Type B token constraint, the differences between them should correspond to illumination change due to the same material property of the common Type B token. If the adjacent Type C tokens are not within the same Type B token, as in the same/smooth illumination constraints, the difference between them should then correspond to a material change since they are not both defined by a common, single material Type B token.

[0047] A Linear token constraint utilizes Type C tokens and Linear tokens. The constraining relationship is that a difference between two Type C tokens, spaced by a Linear token, approximately equals a characteristic illuminant spectral ratio for the image. As defined, a Linear token follows a cylinder configuration along a positive slope, through color space. The BIDR model predicts that the positive slope equals a characteristic illuminant spectral ratio for the image. Thus, the color difference between two Type C tokens, one at each of the dark end and bright end of a Linear token, should reflect the value of the respective characteristic illuminant spectral ratio for the image.

[0048] A BIDR enforcement constraint utilizes Type C tokens and a BIDR model defined normal vector for the log-chromaticity projection plane. The constraining relationship is that the illumination for all Type C tokens in a local patch of the image forms a set of parallel lines in log-color space, the orientation of the parallel lines being defined by the BIDR model defined normal vector. The constraint therefore enforces the illumination field present in the image to explicitly fit the BIDR model prediction for the illumination.

[0049] Thus, each of the Linear token constraint and the BIDR enforcement constraint utilize BIDR model predictions as a basis to segregate illumination and material aspects of an image. The BIDR model predicts a color change in an image when the color change is due to an illumination change forming a shadow (i.e. a decrease in the intensity of the incident illuminant as the interplay between the incident or direct illuminant and the ambient illuminant in the illumination field) over a single material of a scene depicted in the image. The color change prediction of the BIDR model accurately constrains all color band variations among Type C tokens to illumination field effects occurring in an image by operating as a function of the interplay between the spectral variations occurring between incident illuminant and ambient illuminant components of the illumination field. Thus, BIDR model based constraints couple all color band variations into one integral constraining relationship.

[0050] A same texon histogram constraint utilizes Type C tokens and texon histogram operators identified for texture regions, such as texture tokens, within an image. A texon analysis is utilized wherein each pixel of the image (or pixels of those regions of an image identified as comprising a texture) from the recorded color band representation of the respective image file 18, such as, for example, RGB color band values, is converted to a two band representation wherein the two bands comprise a texon label and a texon histogram label. The two band representations are then used to identify texture tokens, as will be described below. A constraint can be imposed that all Type C tokens within the same texture token are of the same mean material.

[0051] A log chromaticity similarity constraint utilizes Type C tokens and log chromaticity representation values. The constraining relationship is that those Type C tokens having pixels with similar log chromaticity representation values are constrained to a same color value, with observed differences being attributed to variations in the illumination field.

[0052] An X-junction constraint utilizes Type C tokens and X-junction operators. As noted above, an X-junction is an area of an image where a material edge and an illumination boundary cross one another. X-junctions are typically identified by four Type C tokens, two pairs of same material Type C tokens forming the material edge, with each same material pair including an illumination boundary dividing the respective same material into lit and shadowed pairs of Type C tokens. The constraining relationship: 1) a Type B token constraint is imposed between each same material pair of Type C tokens forming the X-junction (those with an illumination boundary between them), and 2) a same illumination constraint is imposed between each pair of Type C tokens forming the material edge of the X-junction. For a more detailed description of X-junctions and the relationships of constituent tokens, reference should be made to U. S. Application Serial No. 11/341,742, filed January 27, 2006, entitled: "Method And System For Identifying Illumination Flux In An Image," published as US 2006/0177149 on August 10, 2006.

[0053] A boundary representation constraint is defined by a standard ratio constraint. An analysis performed on a boundary representation, when indicating a material change, provides an estimate of the ratio of colors between two adjacent regions defined by the boundary, for example, the adjacent Type B tokens, even when the illumination varies over the regions. The constraint states that the ratio of the colors of two adjacent regions is X. The boundary representation analysis is executed at the level of Type B tokens, to classify a boundary as being caused by a material change, then propagated down to the level of the constituent Type C tokens. For a more detailed description of a boundary analysis, at the Type B token level, reference should be made to U. S. Application Serial No. 12/079,878, filed March 28, 2008, entitled "System and Method For Illumination Invariant Image Segmentation."

[0054] According to a feature of the present invention, the boundary representation constraint states that all adjacent

pairs of Type C tokens along the boundary, (one Type C token on each side of the boundary, and all of the Type C tokens being within the Type B tokens forming the respective boundary), have colors that satisfy the ratio X, as indicated by the boundary representation analysis.

[0055] According to a preferred embodiment of the present invention, each of the above described constraints can be classified into one of three basic types of constraints, an absolute material color constraint, a same material constraint and a relative reflectance constraint. The absolute material constraint constrains the material at a particular location of an image to be a certain color, as implemented in, for example, the anchor constraint. The same material constraint constrains operators relevant to an image (for example, two pixels or Type C tokens) to be of the same material. The same material type of constraint can be implemented in, for example, Type B, X-junction, log chromaticity similarity, same texon histogram and linear token constraints. The relative reflectance constraint constrains operators relevant to an image (for example, two pixels or Type C tokens) to have a similarity of reflectance characteristics, such as defined by smooth illumination and same illumination constraints, and which can be specified by X-junction, and boundary representation constraints.

[0056] In another example of the selectively varied representation embodiment of the present invention, such as the multi-resolution embodiment, groupings provided by the spatio-spectral operators, such as Type C tokens, Type B tokens and texture tokens, are identified at one scale, such as a fine scale (for example image file 18) and constraining relationships among those groupings are projected to a coarse scale (for example image file 18a). In this exemplary embodiment of the present invention, a constraint implemented at a fine scale representation, such as, for example, the same material constraint of a Type B token constraint, is arranged to constrain pixels or Type C tokens to be within the same Type B token, and further, to be within a relatively limited neighborhood of the fine scale representation. The neighborhood is limited to an extent somewhat less than the extent of a typical relevant Type B token, for example, a neighborhood comprising a 50 x 50 pixel array. The limited neighborhood restriction has the advantage of assuring a sparse matrix for an accurate solution, with operating efficiencies for speedy processing, resulting from the limited extent of the image subject to analysis according to the constraint. However, the implementation across multiple scales also assures a solution according to a long range constraint since a short range within a coarse scale representation corresponds to a relatively long range in the fine scale representation.

[0057] A projection of the constraining relationship imposed on a grouping, from one scale, to another scale, is implemented, for example, by upsampling the grouping identified and used in the constraint at the fine scale representation, to a next coarse scale representation. The performance of the constraint is repeated at the coarse scale, again limiting the neighborhood processed at the coarse scale, and subject to the inter-scale relationship defined by the filtering/down-sampling operation, and so on throughout the various scales of the pyramid. The limit at the coarse scale is a multiple of the same limit of the relatively fine scale. For example, as shown in figure 3b, a single pixel, $P_a(3, 2)$ in the relatively coarse resolution image file 18a corresponds to pixels $P(5, 3)$, $P(6, 3)$, $P(5, 4)$ and $P(6, 4)$ of fine resolution image file 18. The grouping can be identified at any of the scales of resolution, and projected via upsampling or downsampling to the other scales.

[0058] An exemplary solver according to a feature of the present invention comprises a mathematical processing engine for executing an optimizing function, for example, optimization of results in an equation expressed by : $[A] [x] = [b]$, where $[A]$ is a matrix of values that are to be satisfied by (and therefore, taken as solved for by) the definitions of the operator(s) and the constraining relationship(s) for the operator(s), as indicated by selected constraint(s), $[x]$ is a matrix of variables for which the equation is finding an optimal solution, for example, one of an illumination or material component of an image component, for example, a pixel or token, and $[b]$ is a matrix of values observed in an image selected for processing, for example, the recorded values for the RGB color bands of each pixel of an image file 18 or color values of pixels within a selectively varied representation of the original image depicted in image file 18 (e.g. image files 18a ... 18n). The optimizing equation can be implemented in a mathematical optimizing function selected from a set of known optimization solvers such as, for example, known convex optimization operations such as a least squares solver, or a preconditioned conjugate gradient solver.

[0059] According to the selectively varied representation embodiment of the present invention, a solution according to the $[A] [x] = [b]$ matrix equation can be accomplished at each representation, for example, at each scale of the scale-spaced pyramid, either simultaneously or sequentially. The inter-scale relationship constrains the multi-scale solution, while introducing efficiencies in operation execution for increased speed. By dividing processing across different scales, solutions are found for a set of smaller systems rather than one large complex system. Thus, the solutions are more accurate and the CPU 12 can execute in faster computation times while utilizing substantially less memory.

[0060] According to a further feature of the present invention, factors including bounds, are introduced in a solver operation, in addition to constraining relationships, as a function of real world illumination and material phenomena, to keep material/illumination values within physically plausible ranges, such as a limit 1, limit infinity solver (L_1 , L_∞), a bounded least squares solver, or a bounded L_1 , L_∞ solver, as will be described below.

[0061] Figure 4 shows a functional block diagram of an image segregation system architecture, implemented in, for example, the computer system of figure 1, according to a feature of the present invention. Alternatively, the functional

blocks of figure 4 can be implemented in a dedicated hardware circuit arranged to perform the functionality of the blocks of figure 4. An image 32 (as depicted in an image file 18, or one of the selectively varied representations in image files 18a, ... 18n) is input to a preprocessing block 33. The preprocessing block 33 can perform such functions as correction of chromatic aberration in the image 32, combining multiple images to provide a high dynamic range image, linearize pixel data for the image, and so on, for an image optimized for processing. The pre-processed image is then input to a Type C tokenization block 35 which operates to identify Type C tokens in the pre-processed image, in the manner described below. Type C tokens are common to many of the constraints utilized in exemplary embodiments of the present invention, thus, an initial identification of Type C tokens for an input image 32 expedites further processing.

[0062] In an exemplary embodiment of the present invention, the CPU 12 executes code to implement both the preprocessing block 33 and the Type C tokenization block 35, as well as a service provider 24, that functions as a central agent and caching structure (configured in the memory 16), to handle an image for processing according to the teachings of the present invention. The service provider 24 receives and stores the pre-processed image and related Type C token information from the Type C tokenization block 35, and is coupled to an operators block 28 (executed by the CPU 12) arranged to generate any other operators for the image required by selected constraints, as will appear. The service provider 24 is also coupled to a global features extraction input 29. The global features extraction input 29 can be used to provide the system with information relevant to an image being processed, such as an indication of light source when the image was taken (sunlight, fluorescent light, incandescent light), time of day, location, domain knowledge, such as information relevant to the nature of the image, such as interior, exterior, buildings, lawns with green grass, trees with leaves in bloom, etc., and any other parameters relevant to image processing. The service provider 24 stores the global features extraction input 29 with a related input image 32.

[0063] A constraint builder 26 is coupled to the service provider 24. The constraint builder 26 uses a constraint generator library (configured within the memory 16) that stores the constraint generator software modules for the various constraints described above. The service provider 24 and constraint builder 26 operate to arrange spatio-spectral operators relevant to the pre-processed image, according to selected ones of the constraint generator software modules, in for example, the $[A] [x] = [b]$ matrix equation.

[0064] A solver 30 (executed by the CPU 12) is coupled to the constraint builder 26, and implements an optimization operation, as described above, for an optimal solution for the $[A] [x] = [b]$ matrix equation, for use in generating intrinsic images from the pre-processed image. The solver 30 is also coupled to a post-processing block 36 (executed by the CPU 12) for certain post-processing operations. The post-processing operations can include, for example, monotonicity maintenance. In monotonicity maintenance, if two large regions exhibit a linear transition in the input image 32, the transition should remain a linear transition in the output intrinsic image 34. Post-processing can also include illumination propagation, that serves to fill in holes left by the solver 30, illumination-map based white balancing and other filtering, smoothing processes. The post-processing block 36 outputs intrinsic images 34.

[0065] Referring now to figure 5, there is shown a graphical user interface (GUI) for use in connection with an exemplary implementation of the image segregation system architecture feature of the present invention. The GUI of figure 5 is displayed on the monitor 20 of the computer system 10 by the service provider 24 for a user to select a desired image segregation operation. The upper left hand corner of the GUI indicates Open Image, Crop Image, Show Parameters, Segregate and Multi-Resolution selection indicators. A user can move and click a cursor on a desired selector indicator. The Open Image indicator lists all image files 18 currently stored in the memory 16 and enables the user to select an image for processing. The selected image is input 32 (see figure 4) to the service provider 24 (via the preprocessing block 33 and the Type C tokenization block 35) which operates to display the selected image at the upper center of the monitor 20 (figure 5).

[0066] A material image derived by operation of the exemplary segregation system from the selected image is output 34 (see figure 4) after execution of the image segregation processing by the solver 30 and displayed at the lower right hand of the monitor 20 (figure 5). The derived illumination image is displayed at the lower right hand of the monitor 20 (figure 5).

[0067] According to a feature of the present invention, the Crop Image selector permits a user to crop a selected image so as to process a portion of the overall image. The Show Parameter selector displays parameters related to the selected image file 18. Parameters for each image file 18 can be stored in a parameter data file associated with a corresponding image file 18, and include any parameters relevant to the processing of the image depicted in the associated image file 18, for example the global features extraction input 29. Parameters can include any data relevant to image processing such as, for example, any variable for pixel analysis by the CPU 12, as for example, in the generation of spatio-spectral operators, and domain knowledge, such as information relevant to the nature of the image, such as interior, exterior, buildings, lawns with green grass, trees with leaves in bloom, etc.

[0068] A multi-resolution selector can display a drop down menu that provides options such as either a Gaussian or Laplacian pyramid generation, with specification of the number of scales, and the resolution of each scale. The drop down menu can also include selections for a computer operation for illumination and material image segregation, to be performed at a relatively coarse scale representation of the image, and then used to constrain a solution at a relatively

finer scale resolution. In another selection option, a segregation of illumination and material aspects of the image is selectively divided among representations at differing scales of image resolution, as described above.

[0069] Below the selection indicators is a list of each of the optimizing functions that can be used as the solver 30, and a further list of each of the constraint generators contained in the constraint generator library of the constraint builder 26. A user selects a desired mathematical operation and one or more of the constraints to be imposed upon the selected image. After selection of the image to be processed, the constraints to be imposed and the mathematical operation to be executed, the user can click on the Segregate indicator to commence image segregation processing.

[0070] Upon commencement of the image segregation processing, the service provider 24 generates a scale-spaced pyramid, if a multi-resolution operation is selected, and retrieves the constraint generator software modules for the selected constraints to identify the spatio-spectral operators utilized by the selected constraints. Any spatio-spectral operators not already stored by the service provider 24 are generated by the operators block 28, for the image being segregated, and the service provider 24 caches the results. The cached results can be reused in any subsequent operation for a selected image, with the same set of associated parameters.

[0071] In a multi-resolution operation, the service provider 24 can sequence operations, such as perform the segregation at the relatively coarse representation of image file 18a, store the results, then proceed to a segregation of image file 18, enforcing the color constraint described above. The service provider 24 (operating via the CPU 12) can also be operated to execute a standard edge detector, then define regions of the relatively fine resolution image of image file 18 around each identified edge. An image segregation is then performed in respect of the identified edge regions of the image file 18. A subsequent image segregation is performed on relatively coarse image file 18a for regions corresponding to all non-edge regions of image file 18.

[0072] In an example of an implementation of the spatio-spectral operator/constraint/solver model of the present invention, when the selected constraint is a same illumination constraint, the service provider 24 identifies Type C tokens, Type B tokens and a token map for the selected image. The Type C tokens were generated by the Type C tokenization block 35. The service provider 24 operates the operators block 28 to generate the remaining operators specified by the same illumination constraint.

[0073] Referring now to figure 6a, there is shown a flow chart for generating Type C token regions in the image file of figure 2a, (as depicted in one or more of each of image file 18, and image files 18a, ...n), according to a feature of the present invention. Type C tokens can be readily identified in an image by the Type C tokenization block 35, utilizing the steps of figure 6a. The operators block 28 can then analyze and process the Type C tokens to construct Type B tokens when specified by a selected constraint, as will appear.

[0074] Prior to execution of the routine of figure 6a, the CPU 12 can operate to filter the image depicted in a subject image file 18. The filters may include an image 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 in the image. Identification of Type B tokens can be difficult in an image texture. 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.

[0075] 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 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:

- 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 token ID of the token to which it belongs (see description of figure 5a).
- 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.
- 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 $S_n = A \cdot \text{maxValue} + B \cdot \text{pixelValue}$, 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 S_n .
- 5) For each pixel in S_n , 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.
- 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 $R\text{-median}_{ij}$ 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.

[0076] The output of the above algorithm is a mask, Tmask, which is an array of the size of the original image file 18 (n, m). Pixels (p(n, m)) having a corresponding Tmask location in the array with a 1 value should be treated as part of an image region corresponding to texture materials and pixels with a corresponding Tmask array location having a value of 0 should be treated as part of an image region corresponding to materials of homogeneous color.

[0077] The CPU 12 can execute a software module that implements any well known method handling textures, such as, for example, a Laws filter bank, or wavelets (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). See also: Are Filter Banks Necessary? Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, June 2003, volume 2, pages 691-698.

[0078] In one exemplary embodiment of the present invention, Type C tokens are generated using an intensity histogram representation for each pixel of an image (or, preferably, for pixels of the image in regions identified as comprising a texture). The Type C tokens generated using intensity histograms are then used to generate Type B tokens. An intensity histogram representation for each pixel in a texture filter output image can be produced using the following algorithm:

A) Loop through all memory addresses/indexes in the Tmask array (n, m);

B) At each memory address/index of the Tmask array determine whether the value at each location is 1 or 0;

C) If a memory address/index of the Tmask has a 1 value indicating texture:

a. Open the corresponding memory address/index in the corresponding image file 18 p(n, m) and access that memory address/index;

b. Open that memory address/index for writing data;

c. Initialize N 8-bin histograms to zero, one for each color band;

d. For each pixel q within a neighborhood of p (for example, an 11x11 pixel array):

For each of the N color values C_n of q (e.g. RGB wherein each color value is set between 0-255);

(i) If N color value is between 0 and 31, then increment the first bin in 8-bin histogram of color N;

(ii) If N color value is between 32 and 64, then increment the second bin in 8-bin histogram of color N;

(iii) If N color value is between 65 and 97, then increment the third bin in 8-bin histogram of color N;

(iv) If N color value is between 98 and 130, then increment the fourth bin in 8-bin histogram of color N;

(v) If N color value is between 131 and 163, then increment the fifth bin in 8-bin histogram of color N;

(vi) If N color value is between 164 and 196, then increment the sixth bin in 8-bin histogram of color N;

(vii) If N color value is between 197 and 229, then increment the seventh bin in 8-bin histogram of color N;

(viii) If N color value is between 230 and 255, then increment the eighth bin in 8-bin histogram of color N;

e. Concatenate the N 8-bin histogram values together into a single intensity histogram comprising an 8xN element vector;

f. Assign the 8xN element vector to the corresponding pixel p in the image file array p(n, m);

g. If current or next memory address/index is not last/Tmask (n, m), then move on to the next memory address/index in Tmask (n, m) array and execute step B;

h. If current memory address/index is last/equal to Tmask (n, m), then cease looping through Tmask (n, m) array.

D) If a memory address/index has a 0 value:

i. If current address/index is not a last element in Tmask (n, m) array, then move on to the next memory address/index in Tmask (n, m) array and execute step B;

ii. If current memory address/index location is a last element in Tmask (n, m) array, then cease looping through Tmask (n, m) array.

[0079] After the transformation for each pixel in the textured region of the image from a set of color bands, for example RGB, to an intensity histogram vector comprising a set of filter outputs, for example the 8 X N concatenation, the image is treated exactly as the original color band image with respect to identifying type C tokens with the 8 X N element intensity histogram vector being used in place of the RGB color band values.

[0080] Fig. 2b shows an 11 X 11 pixel array at the upper left hand corner of an image file 18. The 11 X 11 pixel array of figure 2b is used as an illustration for a method to produce an intensity histogram representation, as for example, for pixel p(6,6) at the center Z of the array, according to the exemplary algorithm described above. In our example, it is

assumed that the 11 X 11 array is within a checkerboard pattern within the image depicted in the image file 18. The checkerboard pattern comprises a pattern of alternating red, green and blue boxes as shown. Thus, pixel p(6,6) of our example will have a 1 value in the corresponding Tmask generated by the CPU 12.

[0081] Pursuant to the exemplary algorithm, three 8 bin histograms are initialized by the CPU 12 (see figure 2c), one for each of the red, green and blue color values of the subject pixel p(6,6). The 11 X 11 pixel array is then used to determine the values within the bins of the three histograms. In our example, the upper left hand section of the 11 X 11 array is a red box of the checkerboard pattern, and each pixel within that section (p(1,1) to p(5,5)) has an RGB value of (123,0,0). The upper right hand section is green, thus the pixels within the section (p(1,7) to p(5,11)) each have an RGB value of (0,123,0). The lower left hand section is a blue box, with each of the pixels p(7,1) to p(11,5) having RGB values of (0,0,123). Finally the lower right hand section of our example is again a red section, with each of the pixels p(7,7) to p(11,11) having RGB values of (123,0,0), in the repeating red, green and blue checkerboard pattern.

[0082] As shown in Fig. 2b, the one (1) pixel thick cross-shaped border lines between red quadrant (p(1,1) to p(5,5)), green quadrant (p(1,7) to p(5,11)), blue quadrant (p(7,1) to p(11,5)) and red quadrant (p(7,7) to p(11,11)) or any pixel between p(6,1) to p(6,m) and any pixel between p(n,6) to p(11,6) may have an RGB value of (123,123,123). Pixel p(6,6) at the center Z may also have the RGB value of (123,123,123).

[0083] Next, an 11 X 11 pixel box 301, with pixel Z at its center, is used to determine the value of pixel Z via an histogram representation of pixel Z's RGB color band values using the steps exemplarily described below. As shown in Fig. 2b, p(1,1) has RGB color value of (123,0,0) indicating some presence of red color and absence of green color and blue color. Since the red value of p(1,1) is 123, the corresponding bin for that number value in red 8 bin histogram, as shown in Fig. 2c, would be incremented i.e. fourth bin within the red 8-bin histogram would be incremented. This process would repeat itself until p(n,m) in pixel box 301 is reached. By this time, three 8-bin histograms, as shown in Fig. 2c, may be filled with different values. These values may then be concatenated together into a single 8 X 3 element vector and assigned to pixel Z in 11 X 11 pixel array depicted in the image file 18 of figure 2b.

[0084] Consequently, according to the exemplary algorithm described above and a checkerboard pattern within 11 X 11 pixel box 301 shown within image 18 depicted in Fig. 2b, the 3 8-bin histogram would be filled as follows:

[0085] In red 8-bin histogram, the fourth bin would have a value of 71 (the amount of repetitions of red values between 98 and 130 in 25 pixels of 2 red quadrants plus the amount of repetitions of red values between 98 and 130 in 11 pixels between p(6,1) to p(6,m) plus the amount of repetitions of red values between 98 and 130 in 10 pixels between p(n,6) to p(11,6) resulting in $25+25+11+10=71$). In a similar calculation, the first bin would have a value of 50 to indicate all of the pixels with a zero value for red. All the other bins in red 8-bin histogram would have the values of 0. Thus, the red 8-bin histogram would be [50, 0, 0, 71, 0, 0, 0, 0].

[0086] In green 8-bin histogram, the fourth bin would have a value of 46 (the amount of repetitions of green values between 98 and 130 in 25 pixels of the green quadrant plus the amount of repetitions of green values between 98 and 130 in 11 pixels between p(6,1) to p(6,m) plus the amount of repetitions of green values between 98 and 130 in 10 pixels between p(n,6) to p(11,6) resulting in $25+11+10=46$). The first bin has a value of 75 (for zero values of green). All the other bins in green 8-bin histogram would have the values of 0. Thus, the green 8-bin histogram would be [75, 0, 0, 46, 0, 0, 0, 0].

[0087] In blue 8-bin histogram, the fourth bin would have a value of 46 (the amount of repetitions of blue values between 98 and 130 in 25 pixels of the blue quadrant plus the amount of repetitions of blue values between 98 and 130 in 11 pixels between p(6,1) to p(6,m) plus the amount of repetitions of blue values between 98 and 130 in 10 pixels between p(n,6) to p(11,6) resulting in $25+11+10=46$). The first bin, again, has a value of 75. All the other bins in blue 8-bin histogram would have the values of 0. Thus, the blue 8-bin histogram would be [75, 0, 0, 46, 0, 0, 0, 0].

[0088] As a result of the above described algorithm, each 8-bin histogram, three in our example, would be concatenated together into a single 8 X 3 element vector. Accordingly, this vector would be [50, 0, 0, 71, 0, 0, 0, 0, 75, 0, 0, 46, 0, 0, 0, 0, 75, 0, 0, 46, 0, 0, 0, 0].

[0089] Next, the above described vector would be assigned at p(6,6) at the center Z of the 11 X 11 array and therefore replace the RGB value of pixel p(6,6), which originally was (123, 123, 123), with an intensity histogram set forth as [50, 0, 0, 71, 0, 0, 0, 0, 75, 0, 0, 46, 0, 0, 0, 0, 75, 0, 0, 46, 0, 0, 0, 0].

[0090] 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 of figure 6a, the CPU 12 (executing as the Type C tokenization block 35) 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 11 or 15 pixels.

[0091] At step 102, a seed test is begun. The CPU 12 selects a first pixel, $i = (1, 1)$ for example (see figure 2a), 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 intensity histogram representation of the pixel, in the event the image was filtered for texture regions, as described above. If the comparison does not result in approximately equal values (for example, within the noise levels of the recording device for RGB values) for the pixels in the seed, the CPU 12 increments the value of i (step 106), for example, $i = (1, 2)$, for a next $N \times N$ seed sample, and then tests to determine if $i = i_{\max}$ (decision block 108).

[0092] If the pixel value is at i_{\max} , 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 5a parses the entire image at a first value of N before repeating the routine for a reduced value of N .

[0093] 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.

[0094] 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.

[0095] When a good seed (an $N \times N$ 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.

[0096] 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. In the event the pixels comprise intensity histogram representations, the token can be marked as Type C_T .

[0097] 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.

[0098] Upon arrival at $N = N_{\text{stop}}$, step 110 of the flow chart of figure 5a, or completion of a region map that coincides with the image, the routine will have completed the token building task. Figure 6b 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 6c shows token regions corresponding to the region map, for example, as identified through execution of the routine of figure 6a (Type C tokens), in respect to the image of figure 6b. The token regions are color coded to illustrate the token makeup of the image of figure 6b, including penumbra regions between the full color blue and teal areas of the image and the shadow of the colored areas.

[0099] Upon completion of the routine of figure 6a by the Type C tokenization block 35, the service provider 24 stores the Type C token region information for the selected image. Prior to commencing any process to generate Type B tokens from the identified Type C tokens, the operators block 28 tests each identified Type C token to make certain that each Type C token encompasses a single material. 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.

[0100] 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.

[0101] Figure 7 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 6 can be executed as a subroutine directly after a particular token is identified during execution of the routine of figure 6a. 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.

[0102] If no, the CPU 12 exits the routine of figure 7 and returns to the routine of figure 6a (step 126).

[0103] 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 6a, 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.

[0104] 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.

[0105] 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 6d shows Type B tokens generated from the Type C tokens of figure 6c, 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.

[0106] Figure 8 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 6a, with neck removal via the routine of figure 7. The seed size $S = S_{max}$, 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).

[0107] In step 204, the CPU 12 grows a new Type C token, utilizing the routines of figures 6a and 7, 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.

[0108] 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.

[0109] 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.

[0110] 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.

[0111] In step 214, the CPU 12 (as the Type C tokenization block 35 or the operators block 28) 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 whether 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.

[0112] 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 9 shows a flow chart for creating a token graph, containing token map information, 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 6a and 7, 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} .

[0113] 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.

[0114] 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.

[0115] 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.

[0116] 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 6a), 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.

[0117] 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 .

[0118] 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 10. Figure 10 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 planar merging technique, pairs of tokens are examined to determine whether 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.

[0119] 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.

[0120] 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).

[0121] 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.

[0122] 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.

[0123] Upon exiting the routine of figure 10, 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 9, (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.

[0124] 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 11 is a flow chart for generating Type C_w tokens via the local token analysis technique, according to a feature of the present invention.

[0125] 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.

[0126] 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 I_0 and an end point I_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.

[0127] 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).

[0128] For the RGB line fit, the line is defined by:

$l(r,g,b) = I_0(r,g,b) + t(I_1(r,g,b) - I_0(r,g,b))$ 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.

[0129] 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.

[0130] 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.

[0131] 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 is 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).

[0132] Figure 12 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 108. 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.

[0133] 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_p token. In step 426, the CPU 12 stores the list of Type B_p tokens, and exits the routine.

[0134] In a further exemplary embodiment of the present invention, the CPU 12 (executing as the operators block 28) 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, and local token analysis 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.

[0135] 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.

[0136] 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.

[0137] As a result of execution by the Type C tokenization block 35 and/or the operators block 28 (via the CPU 12) 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 (Type B tokens) and tokens representing regions of robust similar color (Type C tokens), thus providing a basis for computational efficiencies, as the token representations capture spatio-spectral information of a significant number of constituent pixels. The service provider 24 stores all of the Type C and Type B tokens generated through execution of the above described token generation techniques, along with the relevant token map information, for example, as determined during execution of the adjacent planar token merging technique, and cross-references the stored operator results to the associated selected image file 18, for use in any segregation processing of the selected image. The cross reference information also includes information on the various selectively varied representations (image files 18a ... 18n) that may have been subject to a tokenization operation.

[0138] In our example of a same illumination constraint, the service provider 24 identifies Type C and Type B tokens as the operators required by the selected constraint. The Type C tokenization block 35 generated the Type C tokens. The service provider 24 operates the operators block 28 to execute the above described techniques, to generate the relevant Type B tokens for the image 32, as well as a token map. The constraint builder 26 organizes the generated token operators according to the exemplary matrix equation, $[A][x] = [b]$, for input to the solver 30. In the same illumination constraint, the constraining relationship of the relevant constraint generator software module is that adjacent Type C tokens, as indicated by the token map information, are lit by the same illumination, unless the adjacent Type C tokens are part of the same Type B token.

[0139] According to the multi-resolution exemplary embodiment of the present invention, the tokenization processes can be executed by the CPU 12 in respect to the original image file 18 or any of the selectively varied representations, image files 18a ... 18n, and constraints applied at any level.

[0140] Each Type C token stored by the service provider 24 is identified by a region ID, and includes a listing of each constituent pixel by row and column number (and a cross reference to the associated image file wherein the token was identified (image files 18, 18a ... 18n)). Each pixel of a Type C token will be of approximately the same color value, for

example, in terms of RGB values, as all the other constituent pixels of the same Type C token, within the noise level of the equipment used to record the image. An average of the color values for the constituent pixels of each particular Type C token can be used to represent the color value for the respective Type C token. Each Type B token is identified by constituent Type C tokens, and thus can be processed to identify all of its constituent pixels via the respective constituent Type C tokens.

[0141] Pursuant to a feature of the present invention, a model for image formation reflects the basic concept of an image as comprising two components, material and illumination. This relationship can be expressed as: $I = ML$, where I is the image color, as recorded and stored in the respective image file 18 and/or convolved in representations 18a ... 18n, M the material component of the recorded image color and L the illumination component of the recorded image color. The I value for each Type C token is therefore the average color value for the recorded color values of the constituent pixels of the token.

[0142] Thus: $\log(I) = \log(ML) = \log(M) + \log(L)$. This can be restated as $i = m + l$, where i represents $\log(I)$, m represents $\log(M)$ and l represents $\log(L)$. In the constraining relationship of the same illumination constraint, in an example where three Type C tokens, a , b and c , (see figure 13) are adjacent (and not within the same Type B token, (as can be shown by a comparison of row and column numbers for all constituent pixels)), $i_a = i_b = i_c$. Since: $i_a = i_a - m_a$, $i_b = i_b - m_b$, and $i_c = i_c - m_c$, these mathematical relationships can be expressed as $(1)m_a + (-1)m_b + (0)m_c = (i_a - i_b)$, $(1)m_a + (0)m_b + (-1)m_c = (i_a - i_c)$ and $(0)m_a + (1)m_b + (-1)m_c = (i_b - i_c)$.

[0143] Figure 13 shows a representation of an $[A][x] = [b]$ matrix equation for the mathematical relationships of the example of the three adjacent Type C tokens a , b and c described above, as constrained by the same illumination constraint: the adjacent Type C tokens a , b and c are at the same illumination. In the matrix equation of figure 15, the various values for the $\log(I)$, in the $[b]$ matrix, are known from the average recorded pixel color values for the constituent pixels of the adjacent Type C tokens a , b and c , generated by the Type C tokenization block 35 from the image selected for segregation. The $[A]$ matrix of 0's, 1's and -1's, is defined by the set of equations expressing the selected same illumination constraint, as described above. The number of rows in the $[A]$ matrix, from top to bottom, corresponds to the number of actual constraints imposed on the tokens, in this case three, the same illumination between three adjacent Type C tokens. The number of columns in the $[A]$ matrix, from left to right, corresponds to the number of unknowns to be solved for, again, in this case, three. Therefore, the values for the material components of each Type C token a , b and c , in the $[x]$ matrix, can be solved for in the matrix equation. It should be noted that each value is actually a vector of three values corresponding to the RGB color bands of our example.

[0144] Accordingly, the matrix equation of figure 13, as arranged by the constraint builder 26, is input by the constraint builder 26 to the solver 30 for an optimized solution for the values of the material components of the adjacent Type C tokens a , b and c of the selected image. As noted above, in the exemplary GUI embodiment of the present invention, a user selects one of several mathematical techniques for finding the optimal solution to the system of constraint equations, $[A][x] = [b]$. The CPU 12 configures the solver 30 according to the mathematical operation selected by the user.

[0145] For example, in a standard least squares solver, the matrix equation is restated as $\underset{x}{\min}(Ax - b)^2$. The solver 30 then executes the least squares operation to determine optimized values for each of m_a , m_b and m_c . The solver 30 can then proceed to generate and display a material image based upon the optimal m_a , m_b and m_c values. In the material image, the m_a , m_b and m_c values are substituted for the originally recorded RGB values, for each pixel of the respective tokens. The solver 30 can proceed to also generate an illumination image from the known recorded image values i_a , i_b , i_c , and the determined m_a , m_b and m_c values, utilizing the model expressed by $i = m + l$.

[0146] Each of the material and illumination images are displayed on the monitor 20, via, for example, the GUI (see figure 5) and can be stored by the service provider 24, and cross-referenced to the original image file 18. According to a feature of the present invention, intrinsic images generated from the representations image files 18a ... 18n by performance of the above described spatio-spectral operator/constraint/solver process, are also stored by the service provider 24.

[0147] As noted above, the spatio-spectral operator/constraint/solver process can be efficiently executed by the CPU 12 first in respect of the relatively coarse representation depicted in image file 18a. Thereafter, the solver (executing via the CPU 12) is constrained by the weighted average for color equation described above, when solving for an illumination image and/or material image for the fine scale original image file 18. Moreover, according to a further exemplary embodiment of the present invention, the spatio-spectral operator/constraint/solver process can be executed by the CPU 12 selectively relative to the various representations of the image. For example, only regions of the image file 18 having edges, are tokenized, while all regions of uniform color within the relatively coarse image file 18a are tokenized, for a division of processing, as described above.

[0148] According to a further feature of the present invention, the solver 30 can be configured to introduce factors including bounds that capture the limits of real world illumination and material phenomena, to keep material/illumination values determined by the optimization procedures as solutions, $[x]$, to within physically plausible ranges. This can be implemented, for example, in an iterative technique to introduce additional inequality constraints on out-of-bounds values in $[x]$, at each iteration, and executed to resolve towards values within the defined bounds. Thus, the above described

least squares technique can be augmented to include minimum and maximum bounds on individual material estimates (as expressed by the entries of $[x]$). Moreover, the entries of $[x]$ can be regularized such that the material estimates are consistent with *a priori* knowledge of material properties.

[0149] In an exemplary embodiment of the present invention, the matrices used in the least squares solver to specify the selected constraints, $[A]$ and $[b]$ are subject to the following bounds, expressed by the problem:

a linear least squares formulation : $\min_{x'} : \sum_i (A_i^T x' - t_i)^2$ subject to:

$$x' \geq \alpha_m \mathbf{1}$$

$$x' \leq \omega_m \mathbf{1}$$

$$x' \geq \text{img}_j$$

where $\mathbf{1}$ denotes the vector of all ones, α_m , the darkest possible material value (for example, a material cannot be darker than coal), and ω_m , the brightest possible material value. The img_j value is the log intensity value at a particular token j , to provide a constraint based upon the real world observation that a segregated material color cannot be darker than it appeared in the original image, since illumination can only brighten the apparent color of an observed material.

[0150] In the linear least squares formulation, the unique minimum solution for x' is the material map that minimizes, in a linear system expressed by $A^T A x' = A^T t$, the average squared difference between the target material differences t_i and the estimated differences $A_i^T x'$. For example, if the "ith" constraint A_i dictates that two tokens a & b are the same material, $A^T A x'$ takes the difference between the values of tokens a & b in x' and computes the distortion from the target value $t_i = 0$.

[0151] The inequalities expressed by the "subject to" bounds set forth above, form a feasible set of material solutions x' which satisfy the real world constraints of possible maximum and minimum material color values. This differs from the standard, known least squares solution in that x' , if not further constraint by the "subject to" bounds, could take on a value at a given location of an image (for example, at a particular pixel or token) that violates the real world observations of reflectance, yet achieves a more optimal solution for the $\min x'$ formulation.

[0152] In the optimization process executed by the solver 30, whenever any tokens have material color values that violate the "subject to" inequalities, at a particular iteration of the process, additional temporary constraints are added that pin the material values in violation, to values that satisfy the bounding conditions. Thus, the original matrices $[A]$ and $[b]$ are augmented with new matrices specifying the new bounding constraints A_{bounds} and b_{bounds} (as an expression of the "subject to" bounds) to define a new augmented system of matrix equations $[A; A_{\text{bounds}}] [x] = [b, b_{\text{bounds}}]$. The augmented system of equations can be solved analogously to the original system, for example, using the known least squares procedure.

[0153] In accordance with the above described bounded feature of the present invention, additional, temporary constraints are added whenever color values violate real world phenomena. A re-solving of the augmented equations can be repeated, as necessary, starting with the original system $A^T A x' = A^T t$, each time (i.e. the temporary bounding constraints need not be carried over between iterations), and iteratively solving augmented systems $A^T A x' = A^T t'$ until the "subject to" bounds are satisfied.

[0154] In accordance with yet another feature of the present invention, an L_1, L_∞ objective function provides a regularization of the optimized solution by encoding a preference for a small number of material changes. In effect, the L_1, L_∞ solver includes the *a priori* belief that material maps should contain a small number of materials in a figure-of-merit. In the solver of the system, there is a distinction between the objective function, a formula that assigns a figure-of-merit to every possible solution, and the algorithm used to find a solution, an optimal value according to a given objective function. As the problem in our exemplary embodiment is stated as a minimization, $\min_{x'} : \sum_i (A_i^T x' - t_i)^2$, the value an objective function assigns can be characterized as a "cost."

[0155] In our problem, let x' be a matrix of a number of rows of tokens and a number of columns of color bands, where x'^c denotes the c^{th} column associated with the c^{th} color band. The least squares objective function, in formula, is augmented, as follows:

$$\min_{x'} : \sum_c \sum_i (A_i^T x'^c - t_i)^2 + \gamma \sum_{k|tk} \max_c |A_k^T x'^c|$$

where $\gamma/\gamma > 0$ governs the trade-off between the cost associated with the least squares term and the L_1, L_∞ penalty. The expression $\sum_{k|tk} \max_c |A_k^T x'^c|$ accumulates the maximum per-channel absolute difference over all the same material

constraints in [A].

[0156] For example, given a same material constraint between tokens a & b , the L_1 , L_∞ function will only include a term for a color channel with the largest difference in between x_a^c and x_b^c over color channel c . In an exemplary embodiment of the present invention, the optimization procedure, for example as expressed by the objective function $\min_{x'} : \sum_c \sum_i (A_i^T x^c - t^c)^2 + \gamma \sum_{k|tk} \max_c |A_i^T x^c|$, is a shrinkage technique. That is, a sequence of least squares problems is solved in a manner wherein, at each round, constraint targets determined to violate the same material constraint are shrunk. At the end of the sequence, constraints with a value below a given threshold are culled from the constraint system, and a new least squares solution is computed. It should be noted that bounds such as the "subject to" bounds discussed above, can be added to the objective function to provide a bounded L_1 , L_∞ solver.

[0157] Figure 14 is a generalized functional block diagram for the service provider 24 and constraint builder 26. To summarize the above described constraint examples in a general scheme, a selection is made of an image 32, (one or more of image file 18, 18a, ... 18n) and, for each selected image 32, a number of constraint generators from a set of constraint generators 1, 2, ...N, (the constraint generator software modules) for example, by a user, via the GUI. The set of constraint generators 1-N includes the constraints described above, and any additional constraining relationships developed as a function of spatio-spectral information for an image. The above described set of constraints is provided as an example. The present invention contemplates any constraining relationship based upon spatio-spectral operators, that provides a logical deduction regarding material and illumination aspects of an image, and thus a basis for constructing matrices [A] and [b] to define a set of equations whose optimal solution captures intrinsic illumination and material components of a given image.

[0158] Likewise, a set of operators 1-M, generated by the Type C tokenization block 35 or the operators block 28, includes all operators defined in the constraint generator modules 1-N. As shown in figure 14, the service provider 24 provides all of the operators 1-M, as required by the selected constraint generators 1-N and further couples the selected constraint generators 1-N to a constraint assembly 39 via a logical switch 40 (both configured within the constraint builder 26). In the event any of the operators 1-M for a selected image 32 are not already stored by the service provider 24, the service provider 24 utilizes the operators block 28 to compute such operators on demand, in the manner described above. The constraint assembly 39 constructs a separate $[A] [x] = [b]$ matrix for each one of the selected constraint generators, as a function of the operators and the constraining relationships defined in the respective constraint generators 1-N. In each case, the $[A] [x] = [b]$ matrix is constructed in a similar manner as described above for the same illumination example.

[0159] Upon completion of the construction of the system of equations $[A]_i [x] = [b]_i$, for each of the selected constraint generators, $i = \{1, 2, \dots, N\}$, the constraint assembly 39 concatenates the constituent matrices $[A]_i$, $[b]_i$, from each constraint generator. Since each of the concatenated equations may contain a different subset of the unknowns, $[x]$, the assembly is performed such that corresponding columns of individual matrices $[A]_i$, that constrain particular unknowns in $[x]$, are aligned. The concatenated matrices, $[A] [x] = [b]$, are then input to the solver 30, for solution of the unknowns in the complete $[x]$ vector, pursuant to the selected optimization procedure, for output of intrinsic images 34. The individual constraints within the concatenated matrices, $[A] [x] = [b]$, can be weighted relative to one another as a function of factors such as perceived importance of the respective constraint, strength or empirically determined confidence level.

[0160] The above described example of a same illumination constraint utilizes Type C token and Type B token spatio-spectral operators. These token operators provide an excellent representation of images that include large surface areas of a single material, such as are often depicted in images including man-made objects. As such, processing of the image at a coarse scale of resolution, for example, at the level of image file 18a, is possible. However, in many natural scenes there are often large areas of highly textured regions, such as sand, grass, stones, foliage, and so on. Each texture has an intrinsic scale, which can be correlated to one of the scale representations of image files 18a, ...18n. As noted above, identification of Type B tokens using Type C tokens, can be difficult in an image texture, particularly at a fine scale of resolution. According to a further feature of the present invention, a texton histogram operator provides a mechanism for capturing statistically uniform spatial variations of textured regions in a manner that is useful in a constraint based optimization, for example, as expressed by the $[A] [x] = [b]$ matrix equation.

[0161] Thus, according to this feature of the present invention, rather than generating Type C tokens in textured regions of an image, from intensity histograms, for use in identifying Type B tokens, as described above, texture tokens are generated as a species of Type B tokens, for use in a constraint. Such a texture token analysis can be performed, for example, at a preselected scale suitable to the scale of the texture, for a more efficient and accurate analysis. In an exemplary embodiment of the texton histogram operator, the operators block 28 converts each pixel of the image (or pixels of those regions of an image identified as comprising a texture) from the recorded color band representation of the respective image file 18, such as, for example, RGB color band values, to a two band representation wherein the two bands comprise a texton label and a texton histogram label. The two band representations for the pixels are then used to identify texture tokens, as will appear.

[0162] A texton label for each pixel is generated through execution of a clustering process. A texture can be characterized by a texture primitive (for example, in a grass texture, a single blade of grass), and the spatial distribution of the

primitive. A texton analysis is an analytical method for characterizing a texture primitive, for example via a clustering algorithm. Clustering is a process for locating centers of natural groups or clusters in data. In an exemplary embodiment of the present invention, the data comprises pixel patches selected from among the pixels of an image being segregated into material and illumination components. For example, 3 X 3 pixel patches are clustered into K different groups, with each group being assigned a designating number (1,2,3, ...K). The texton label for each pixel of the 3 X 3 array is the group number of the group to which the respective patch was assigned during the clustering process.

[0163] To expedite execution of a clustering algorithm, random samples of 3 X 3 patches can be selected throughout the image, or region of the image identified as comprising a texture, for processing in a clustering algorithm. After execution of the clustering algorithm by the CPU 12 (operating as the operators block 28), each 3 X 3 patch of the image is assigned the texton label of the closest one of the K group centers identified in the clustering process, as executed in respect of the selected random samples.

[0164] To advantage, prior to execution of a clustering algorithm, the pixels of the image are subject to an image intensity normalization. In a clustering process utilizing an intensity-based distance matrix, dark areas of an image may be placed in a single group, resulting in an under representation of groups for shadowed areas of a textured region of an image. A normalization of the image provides a more accurate texton representation for texture regions under varying illumination. A normalized intensity for a pixel can be expressed by:

$$i_{\text{norm}}(n,m) = \log (i(n,m)/i_b(n,m)),$$

where $i_{\text{norm}}(n,m)$ is the normalized intensity for a pixel $p(n,m)$, $i(n,m)$ is the intensity for the pixel $p(n,m)$, as recorded in the image file 18, and $i_b(n,m)$ is a blurred or low passed filtered version of the pixel $p(n,m)$. For example, a 10 pixel blur radius can be used in any standard blurring function.

[0165] Clustering can be executed according to any known clustering algorithm, such as, for example, K means clustering where there are K clusters or groups S_i , $i = 1,2, \dots, K$, and μ_i is the mean point or center point of all the data points $x_j \in S_i$. In our example, each x_j comprises a selected 3 X 3 pixel patch arranged as a 9 X 1 vector of the nine pixels in the patch (27 elements total, including the RGB values of each of the nine pixels of the vector). As noted above, each mean point μ_i is assigned a texton label, 1, 2, 3 ...K, that becomes the texton label for any pixel of a 3 X 3 patch clustered into the group for which the respective mean point is the center.

[0166] According to an exemplary embodiment of the present invention, the CPU 12 executes the algorithm by initially partitioning the selected 9 X 1 vectors, representing 3 X 3 pixel patches of the image, into K initial groups S_i . The CPU 12 then calculates a center point μ_i , for each group S_i , utilizing an intensity-based distance matrix. After determining a center point μ_i for each group S_i , the CPU 12 associates each 9 X 1 vector to the closest center point μ_i , changing groups if necessary. Then the CPU 12 recalculates the center points μ_i . The CPU 12 executes iterations of the steps of associating each 9 X 1 vector to the closest center point μ_i , and recalculating the center points μ_i , until convergence. Convergence is when there is no need to change the group for any of the 9 X 1 vectors. At that point, the CPU 12 assigns the group number for the respective center point μ_i , as the texton label for the pixels of each vector in that group.

[0167] As noted above, pixels of 3 X 3 patches not selected as samples for clustering are assigned the texton label of the closest one of the K group centers μ_i , identified in the clustering process, as executed in respect of the selected random samples. A texton label map is stored by the service provider 24, and is coextensive with the pixel array of figure 2. In the texton label map, for each pixel location, there is an indication of the respective texton label.

[0168] Upon completion of the texton label assignment for pixels of the image, the CPU 12 operates to generate a texton histogram for each pixel to provide a representation of the spatial variation of texton representations within a textured region of the image. To that end, the CPU 12 accesses the texton label map. At each pixel location within the texton label map, a pixel patch of, for example, 21 X 21 pixels, is set up around the current location. The 21 X 21 patch size is far greater than the 3 X 3 patch sized used to generate the texton representations, so as to capture the spatial variations of the texture. A texton histogram is then generated for the pixel location at the center of the 21 X 21 patch, in a similar manner as the intensity histogram described above. However, rather than bins based upon color band values, in the texton histogram, there is a bin for each texton label value, 1,2, 3 ...K. The count for each bin corresponds to the number of pixels in the 21 X 21 patch having the texton label value for the respective bin.

[0169] When a texton histogram is generated for each pixel of the texton label map, the CPU 12 executes a second clustering step. In the second clustering step, the texton histograms are clustered using spectral clustering. Spectral clustering techniques use a spectrum of a similarity matrix of data of interest, (in our example, the texton histograms) to reduce the dimensionality for clustering in fewer dimensions. A similarity matrix for a given set of data points A can be defined as a matrix S where S_{ij} represents a measure of the similarity between points $i, j \in A$. In our example, eigenvectors of the Laplacian are clustered using a mean shift. The distance metric is a chi-squared distance of the histograms.

[0170] A texton histogram label (1,2 ...) is assigned to each cluster group defined by the clustering procedure. For

each pixel of the texton label map, the texton histogram label for the cluster group corresponding to a texton histogram that is nearest the texton histogram for the respective pixel, is assigned to that pixel. Distance is defined as the chi-squared histogram distance. Upon completion of the assignment of a texton histogram label to each pixel, each pixel is now represented by a two band, texton label, texton histogram label representation.

[0171] According to a feature of the present invention, the two band, texton label, texton histogram label representations for pixels of an image file 18 (or any of the selectively varied representations 18a, ... 18n) can be utilized in a constraint for construction of an $[A]_i [x] = [b]_i$ constituent within the concatenated matrices, $[A] [x] = [b]$. For example, it can be assumed that a region of an image wherein contiguous pixels within the region all have the same two band, texton label, texton histogram label representation, comprises a region of the same mean material of a texture depicted in the image. Such a region can be referred to as a texture token, a species of a Type B token. Thus, a constraint can be imposed that all Type C tokens within the same texture token are of the same mean material. In this constraint, the Type C tokens are the Type C tokens generated from the color band values of the constituent pixels by the Type C tokenization block 35.

[0172] While the above exemplary embodiment of the present invention has been described with a user selecting constraint generators and mathematical operations via a GUI, the image segregation processing can be done in other operating modes, such as automatically, with images, constraint generators and mathematical operations being automatically selected, for example, as a function of image parameters.

[0173] Referring once again to figure 3a, in block 1006, the CPU 12 outputs the intrinsic images generated in the image segregation block 1004. Such intrinsic images can be advantageously used in such applications as object recognition, or for improved results in an image modification. For example, a color correct gamma correction can be achieved by performing an intensity adjustment on the illumination image, and merging the intensity adjusted illumination image with the corresponding material image, for a color correct, intensity adjusted output image.

[0174] In summary, the selectively varied representation embodiment of the present invention provides modes of operation that permit spatio-spectral operators to be computed and applied at any one of several different representations of the image, for example, a representation at a scale of resolution that matches the intrinsic scale of an image feature. Also, different areas of the image can be processed at different scales, such as edge regions at a fine scale, and non-edge regions at a coarse scale. Moreover, a constraint can be applied over a limited extent of one representation of the image and projected to a more coarse scale to extend the range of the constraint. All of the various modes afforded by the selectively varied representation embodiment of the present invention improve the accuracy of an image segregation operation, while providing efficiencies that speed up the performance of the computer system 10.

[0175] 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 (32), comprising the steps of:

providing an image file (18) depicting an image (32), in a computer memory (16);
forming a set of selectively varied representations (18a, 18b, ..., 18n) of the image file (18) by forming multi-resolution representations of the image (32); and
performing an image segregation operation on at least one preselected representation (18a, 18b, ..., 18n) of the image (32) of the image file (18) as a function of an image segregation of another one of the set of selectively varied representations (18a, 18b, ..., 18n) of the image (32), to generate intrinsic images (34) corresponding to the image (32), by utilizing a spatio-spectral operator/constraint/solver model, which includes
generating spatio-spectral operators, spatio-spectral operators being a mathematical representation of spatio-spectral information, spatio-spectral information being information on spectral relationships among contiguous pixels, in terms of color bands, and the spatial extent of the pixel spectral relationships, introducing constraints between spatio-spectral operators, each constraint defining the spatio-spectral operators used by the constraint and a constraining relationship imposed upon these used spatio-spectral operators, expressing the constraints in a matrix equation $Ax=b$, wherein A is a matrix of values determined by the definition of the spectral operators and the constraining relationships for the spectral operators, b is a matrix of values observed in the image (32) or in a selectively varied representation (18a, 18b, ..., 18n) of the image (32), and x is a matrix of values to be solved for for generating the intrinsic images (34), and solving for the matrix x,
wherein one constraint constrains a weighted average of material color at a local area of a relatively finer scale representation to equal the material color of a corresponding local area of a relatively coarser scale representation.

2. The method of claim 1 wherein the multi-resolution representations of the image (32) comprise a scale-spaced pyramid of representations.
3. The method of claim 1 wherein the intrinsic images (34) comprise a material image and an illumination image.
4. The method of claim 1 wherein the step of utilizing a spatio-spectral operator/constraint/solver model is carried out by generating spatio-spectral operators at each of different ones of the selectively varied representations (18a, 18b, ..., 18n) of the image (32), and implementing a constraint as a function of the spatio-spectral operators generated at each of different ones of the selectively varied representations (18a, 18b, ..., 18n) of the image (32).
5. The method of claim 4 wherein the spatio-spectral operators include a weighted color average.
6. The method of claim 1 wherein the step of performing an image segregation operation on at least one preselected representation (18a, 18b, ..., 18n) of the image (32) of the image file (18), to generate intrinsic images (34) corresponding to the image (32), is carried out by performing an image segregation operation in respect of preselected regions of at least one preselected representation (18a, 18b, ..., 18n) of the image (32) of the image file (18).
7. The method of claim 6 wherein the step of performing an image segregation operation in respect of preselected regions of at least one preselected representation (18a, 18b, ..., 18n) of the image (32) of the image file (18) is carried out by performing an image segregation operation in respect of different preselected regions of at least two of the set of selectively varied representations (18a, 18b, ..., 18n) of the image (32).
8. The method of claim 1 wherein the step of performing an image segregation operation on at least one preselected representation (18a, 18b, ..., 18n) of the image (32) of the image file (18) is carried out over a limited extent of the at least one preselected representation (18a, 18b, ..., 18n) of the image (32).
9. The method of claim 8 wherein the function of an image segregation of another one of the set of selectively varied representations (18a, 18b, ..., 18n) of the image (32) includes a projection of the image segregation operation carried out over a limited extent of the at least one preselected representation (18a, 18b, ..., 18n) of the image (32) to another one of the set of selectively varied representations (18a, 18b, ..., 18n) of the image (32).
10. A computer system (10) which comprises:
 - a CPU (12); and
 - a memory (16) storing an image file (18) containing an image (32);

the CPU (12) is arranged and configured to execute a routine to perform the method of claim 1.

Patentansprüche

1. Automatisiertes, computergestütztes Verfahren zum Verarbeiten eines Bilds (32), umfassend die folgenden Schritte:
 - Bereitstellen einer Bilddatei (18), die ein Bild (32) zeigt, in einem Computerspeicher (16);
 - Bilden eines Satzes von selektiv variierten Darstellungen (18a, 18b, ..., 18n) der Bilddatei (18) durch Bilden von Mehrfachauflösungsdarstellungen des Bilds (32); und
 - Durchführen eines Bildaufspaltungsvorgangs an mindestens einer vorausgewählten Darstellung (18a, 18b, ..., 18n) des Bilds (32) der Bilddatei (18) als eine Funktion einer Bildaufspaltung einer anderen des Satzes selektiv variierten Darstellungen (18a, 18b, ..., 18n) des Bilds (32), um intrinsische Bilder (34) zu erzeugen, die dem Bild (32) entsprechen, unter Verwendung eines räumlich-spektralen Operator-/Einschränkungs-/Lösermodells, das das Erzeugen räumlich-spektraler Operatoren umfasst, wobei räumlich-spektrale Operatoren eine mathematische Darstellung von räumlich-spektralen Informationen sind, wobei räumlich-spektrale Informationen Informationen über spektrale Beziehungen zwischen benachbarten Pixeln in Bezug auf Farbbänder sind, und die räumliche Reichweite der spektralen Beziehungen der Pixel, wodurch Einschränkungen zwischen den räumlich-spektralen Operatoren entstehen, wobei jede Einschränkung, die die räumlich-spektralen Operatoren definiert, die von der Einschränkung verwendet werden, und eine Einschränkungsbeziehung, die diesen verwendeten räumlich-spektralen Operatoren auferlegt ist, die Einschränkungen in einer Matrixgleichung $Ax = b$ ausdrückt, wobei A eine Matrix von Werten ist, die durch die Definition der Spektraloperatoren und der einschränkenden

Beziehungen für die Spektraloperatoren bestimmt ist, b eine Matrix von Werten ist, die in dem Bild (32) oder in einer selektiv variierten Darstellung (18a, 18b, ..., 18n) des Bilds (32) beobachtet werden, und x eine Matrix von Werten ist, die zum Erzeugen der intrinsischen Bilder (34) gelöst werden sollen, wobei eine Einschränkung einen gewichteten Mittelwert von Materialfarbe an einem lokalen Bereich einer relativ feineren skalierten Darstellung einschränkt, um die Materialfarbe eines entsprechenden lokalen Bereichs einer relativ gröberen skalierten Darstellung auszugleichen.

2. Verfahren nach Anspruch 1, wobei die Mehrfachauflösungsdarstellungen des Bilds (32) eine skaliert-beabstandete Pyramide von Darstellungen umfassen.

3. Verfahren nach Anspruch 1, wobei die intrinsischen Bilder (34) ein Materialbild und ein Beleuchtungsbild umfassen.

4. Verfahren nach Anspruch 1, wobei der Schritt der Verwendung eines räumlich-spektralen Operator-/Einschränkungs-/Lösermodells durch Erzeugen von räumlich-spektralen Operatoren an jedem von verschiedenen der selektiv variierten Darstellungen (18a, 18b, ..., 18n) des Bilds (32) und durch Implementieren einer Einschränkung als eine Funktion der räumlich-spektralen Operatoren, die an jeder von verschiedenen der selektiv variierten Darstellungen (18a, 18b, ..., 18n) des Bilds (32) erzeugt werden, durchgeführt wird.

5. Verfahren nach Anspruch 4, wobei die räumlich-spektralen Operatoren einen gewichteten Farbmittelwert umfassen.

6. Verfahren nach Anspruch 1, wobei der Schritt des Ausführens eines Bildaufspaltungsvorgangs an mindestens einer vorausgewählten Darstellung (18a, 18b, ..., 18n) des Bilds (32) der Bilddatei (18) zum Erzeugen intrinsischer Bilder (34) entsprechend dem Bild (32) durch Ausführen eines Bildaufspaltungsvorgangs in Bezug auf vorausgewählte Bereiche mindestens einer vorausgewählten Darstellung (18a, 18b, ..., 18n) des Bilds (32) der Bilddatei (18) durchgeführt wird.

7. Verfahren nach Anspruch 6, wobei der Schritt des Ausführens eines Bildaufspaltungsvorgangs in Bezug auf vorausgewählte Bereiche mindestens einer vorgewählten Darstellung (18a, 18b, ..., 18n) des Bilds (32) der Bilddatei (18) durch Ausführen eines Bildaufspaltungsvorgangs in Bezug auf verschiedene vorausgewählte Bereiche von mindestens zwei aus dem Satz selektiv variierten Darstellungen (18a, 18b, ..., 18n) der Bilddatei (18) durchgeführt wird.

8. Verfahren nach Anspruch 1, wobei der Schritt des Ausführens eines Bildaufspaltungsvorgangs an mindestens einer vorgewählten Darstellung (18a, 18b, ..., 18n) des Bilds (32) der Bilddatei (18) über einen begrenzten Umfang der mindestens einen vorausgewählten Darstellung (18a, 18b, ..., 18n) des Bilds (32) durchgeführt wird.

9. Verfahren nach Anspruch 8, wobei die Funktion einer Bildaufspaltung einer anderen aus dem Satz selektiv variierten Darstellungen (18a, 18b, ..., 18n) des Bilds (32) eine Projektion des Bildaufspaltungsvorgangs umfasst, die über einen begrenzten Umfang der mindestens einen vorausgewählten Darstellung (18a, 18b, ..., 18n) des Bilds (32) auf eine andere aus dem Satz selektiv variierten Darstellungen (18a, 18b, ..., 18n) des Bilds (32) durchgeführt wird.

10. Computersystem (10), das Folgendes umfasst:

eine CPU (12); und

einen Speicher (16), der eine Bilddatei (18) speichert, die ein Bild (32) enthält;

wobei die CPU (12) angeordnet und konfiguriert ist, um eine Routine auszuführen, um das Verfahren nach Anspruch 1 auszuführen.

Revendications

1. Procédé automatisé informatisé de traitement d'une image (32), comprenant les étapes consistant à:

fournir un fichier d'image (18) décrivant une image (32) dans une mémoire d'ordinateur (16);

former un ensemble de représentations (18a, 18b, ..., 18n) sélectivement variées du fichier d'image (18) en formant des représentations à résolutions multiples de l'image (32); et

effectuer une opération de ségrégation d'image sur au moins une représentation présélectionnée (18a, 18b, ...,

18n) de l'image (32) du fichier d'image (18) en fonction d'une ségrégation d'image d'une autre de l'ensemble des représentations (18a, 18b, ..., 18n) sélectivement variées de l'image (32), pour générer des images intrinsèques (34) correspondant à l'image (32), en utilisant un modèle opérateur/contrainte/solveur spatio-spectral, qui inclut la génération d'opérateurs spatio-spectraux, les opérateurs spatio-spectraux étant une représentation mathématique d'informations spatio-spectrales, les informations spatio-spectrales étant des informations sur les relations spectrales entre des pixels contigus, en termes des bandes de couleur et de l'étendue spatiale des relations spectrales en pixels, introduisant des contraintes entre les opérateurs spatio-spectraux, chaque contrainte définissant les opérateurs spatio-spectraux utilisés par la contrainte et une relation de contrainte imposée aux opérateurs spatio-spectraux utilisés, exprimant les contraintes dans une équation matricielle $Ax = b$, dans laquelle A est une matrice de valeurs déterminée par la définition des opérateurs spectraux et des relations de contrainte pour les opérateurs spectraux, b est une matrice de valeurs observées dans l'image (32) ou dans une variable modifiée de manière sélective représentation (18a, 18b, ..., 18n) de l'image (32), et x est une matrice de valeurs à résoudre pour générer les images intrinsèques (34), dans laquelle une contrainte contraint une moyenne pondérée de la couleur du matériau au niveau d'une zone locale d'une représentation à échelle relativement fine pour correspondre à la couleur du matériau d'une zone locale correspondante d'une représentation à l'échelle relativement plus grossière.

2. Procédé selon la revendication 1, dans lequel les représentations à résolutions multiples de l'image (32) comprennent une pyramide de représentations espacées à l'échelle.

3. Procédé selon la revendication 1, dans lequel les images intrinsèques (34) comprennent une image matérielle et une image d'éclairage.

4. Procédé selon la revendication 1, dans lequel l'étape d'utilisation d'un modèle opérateur/contrainte/solveur spatio-spectral est réalisée en générant des opérateurs spatio-spectraux au niveau de chacune des différentes représentations parmi les représentations sélectivement variées (18a, 18b, ..., 18n) de l'image (32), et mettant en oeuvre une contrainte en fonction des opérateurs spatio-spectraux générés au niveau de différentes représentations parmi les représentations sélectivement variées (18a, 18b, ..., 18n) de l'image (32).

5. Procédé selon la revendication 4, dans lequel les opérateurs spatio-spectraux comprennent une moyenne de couleur pondérée.

6. Procédé selon la revendication 1, dans lequel l'étape consistant à effectuer une opération de ségrégation d'images sur au moins une représentation présélectionnée (18a, 18b, ..., 18n) de l'image (32) du fichier d'image (18), pour générer des images intrinsèques (34) correspondant à l'image (32), est réalisée en effectuant une opération de ségrégation d'image pour des régions présélectionnées d'au moins une représentation présélectionnée (18a, 18b, ..., 18n) de l'image (32) du fichier image (18).

7. Procédé selon la revendication 6, dans lequel l'étape consistant à effectuer une opération de ségrégation d'image en ce qui concerne des régions présélectionnées d'au moins une représentation présélectionnée (18a, 18b, ..., 18n) de l'image (32) du fichier d'image (18) est réalisée en effectuant une opération de ségrégation d'images en ce qui concerne différentes régions présélectionnées d'au moins deux des ensembles de représentations (18a, 18b, ..., 18n) sélectivement variées du fichier d'image (18).

8. Procédé selon la revendication 1, dans lequel l'étape consistant à effectuer une opération de ségrégation d'image sur au moins une représentation présélectionnée (18a, 18b, ..., 18n) de l'image (32) du fichier d'image (18) est effectuée sur une étendue limitée de la au moins une représentation présélectionnée (18a, 18b, ..., 18n) de l'image (32).

9. Procédé selon la revendication 8, dans lequel la fonction de ségrégation d'image d'un autre parmi l'ensemble des représentations sélectivement variées (18a, 18b, ..., 18n) de l'image (32) comprend une projection de l'opération de ségrégation d'image effectuée sur une étendue limitée de la au moins une représentation présélectionnée (18a, 18b, ..., 18n) de l'image (32) par rapport à une autre de l'ensemble des représentations sélectivement variées (18a, 18b, ..., 18n) de la image (32).

10. Un système informatique (10) qui comprend:

un processeur (12); et

EP 2 435 956 B1

une mémoire (16) stockant un fichier d'image (18) contenant une image (32);

le processeur (12) est agencé et configuré pour exécuter une routine pour exécuter le procédé de la revendication 1.

5

10

15

20

25

30

35

40

45

50

55

Computer System Configured to Operate on Images






FIG. 1

Pixel Array for Storing Image Data

$P(1, 1)$	$P(1, 2)$	\dots	\dots	$P(1, M)$
$P(2, 1)$	$P(2, 2)$			\vdots
$P(3, 1)$	$P(3, 2)$			\vdots
\vdots				
\vdots				\vdots
$P(N, 1)$	\dots		\dots	$P(N, M)$

Image File
18

FIG. 2A






Figure 3a




Figure 3b


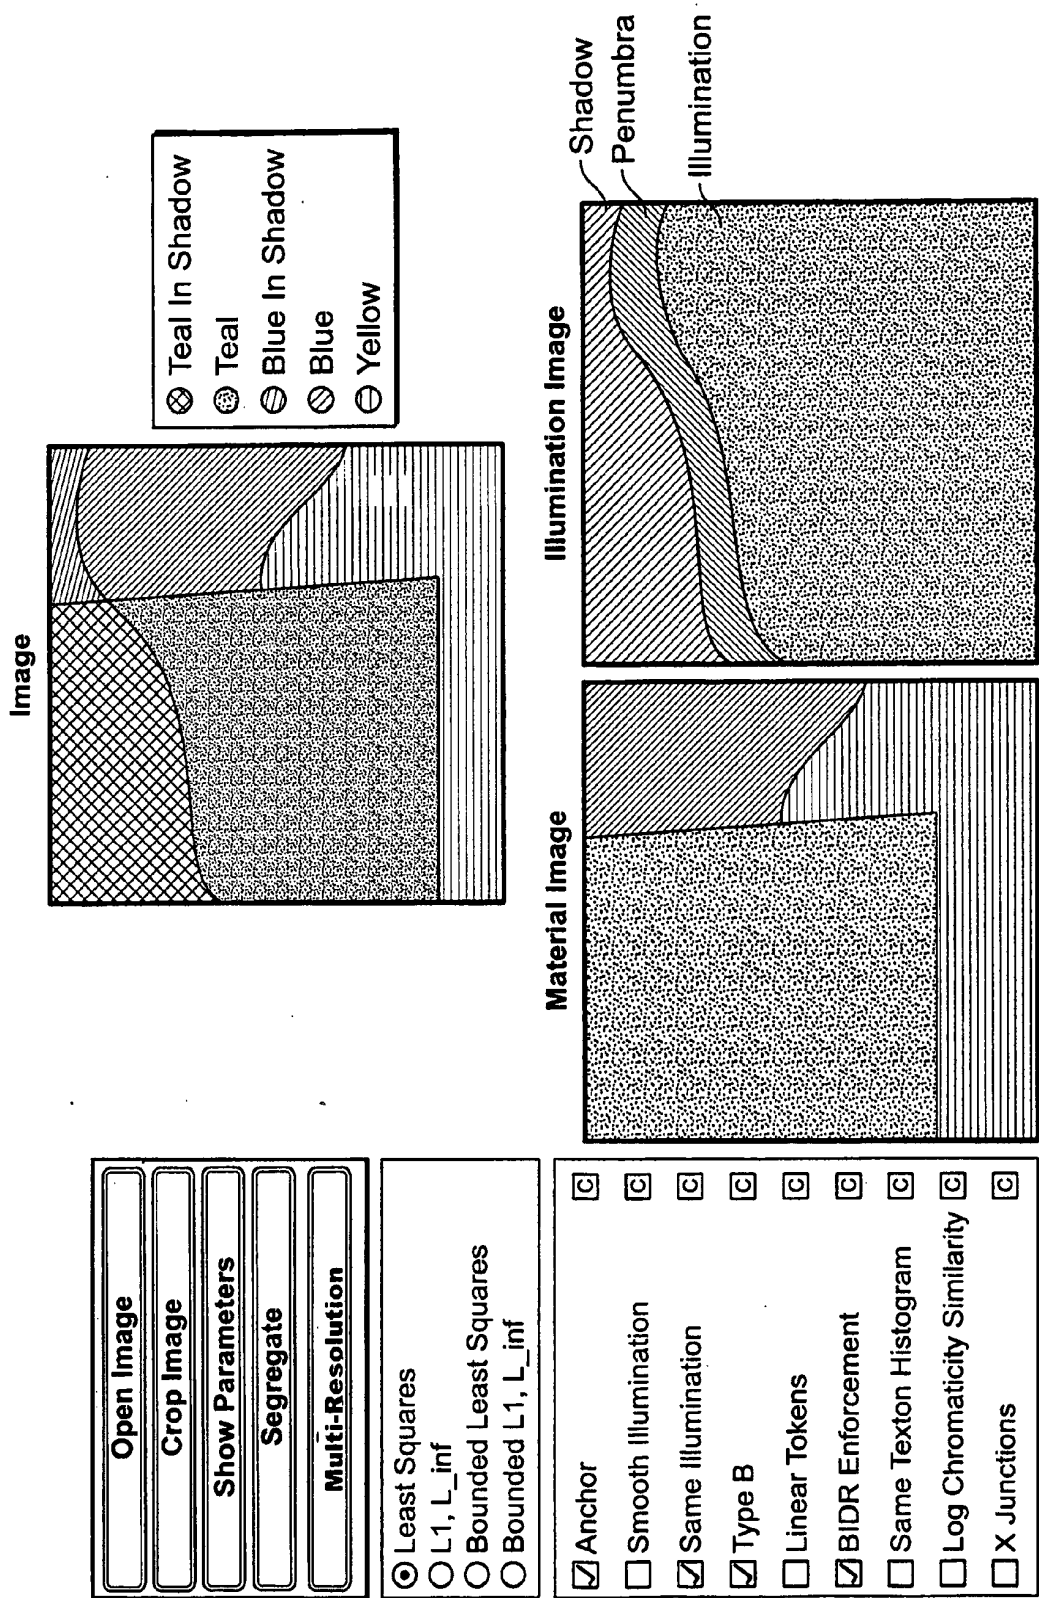


FIG. 4



Identifying Token Regions in an Image

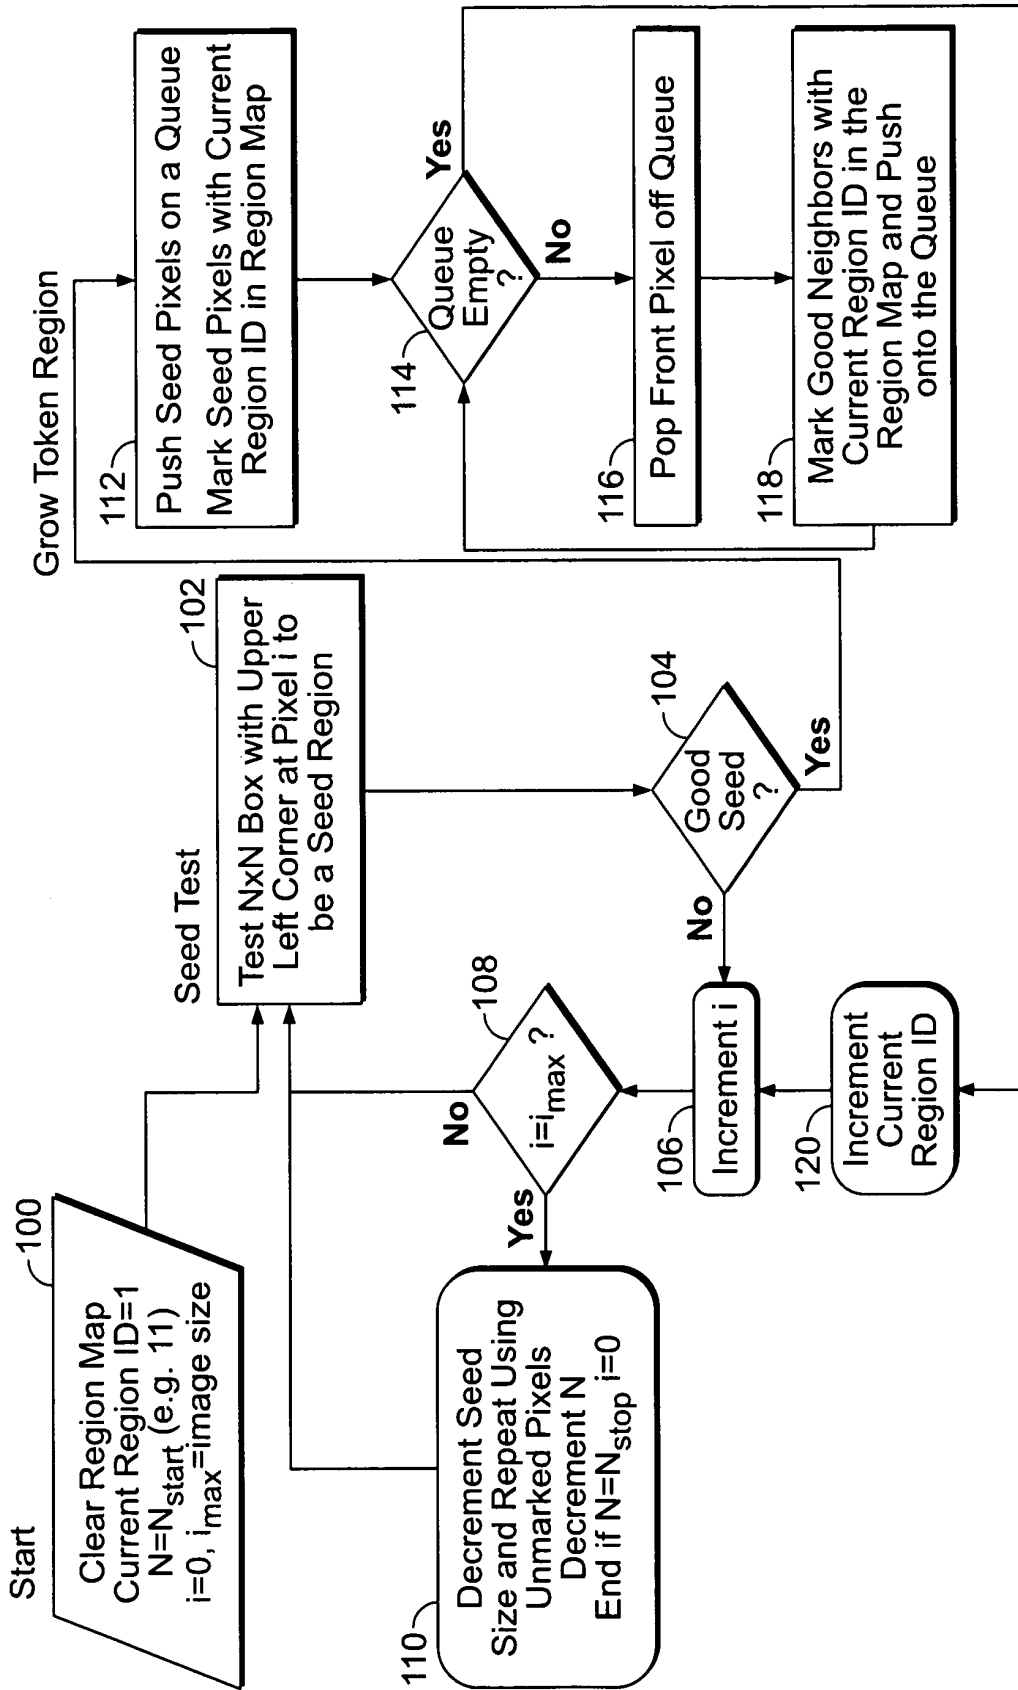



FIG. 6A

Examples of Identifying Token Regions in an Image

Original Image



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

FIG. 6B

Token Regions




FIG. 6C




FIG. 6D

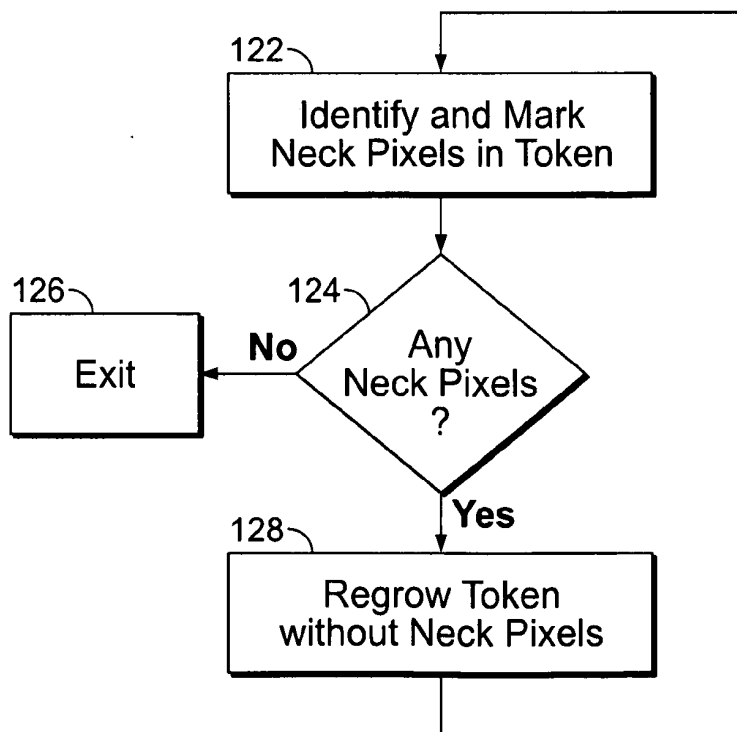


FIG. 7

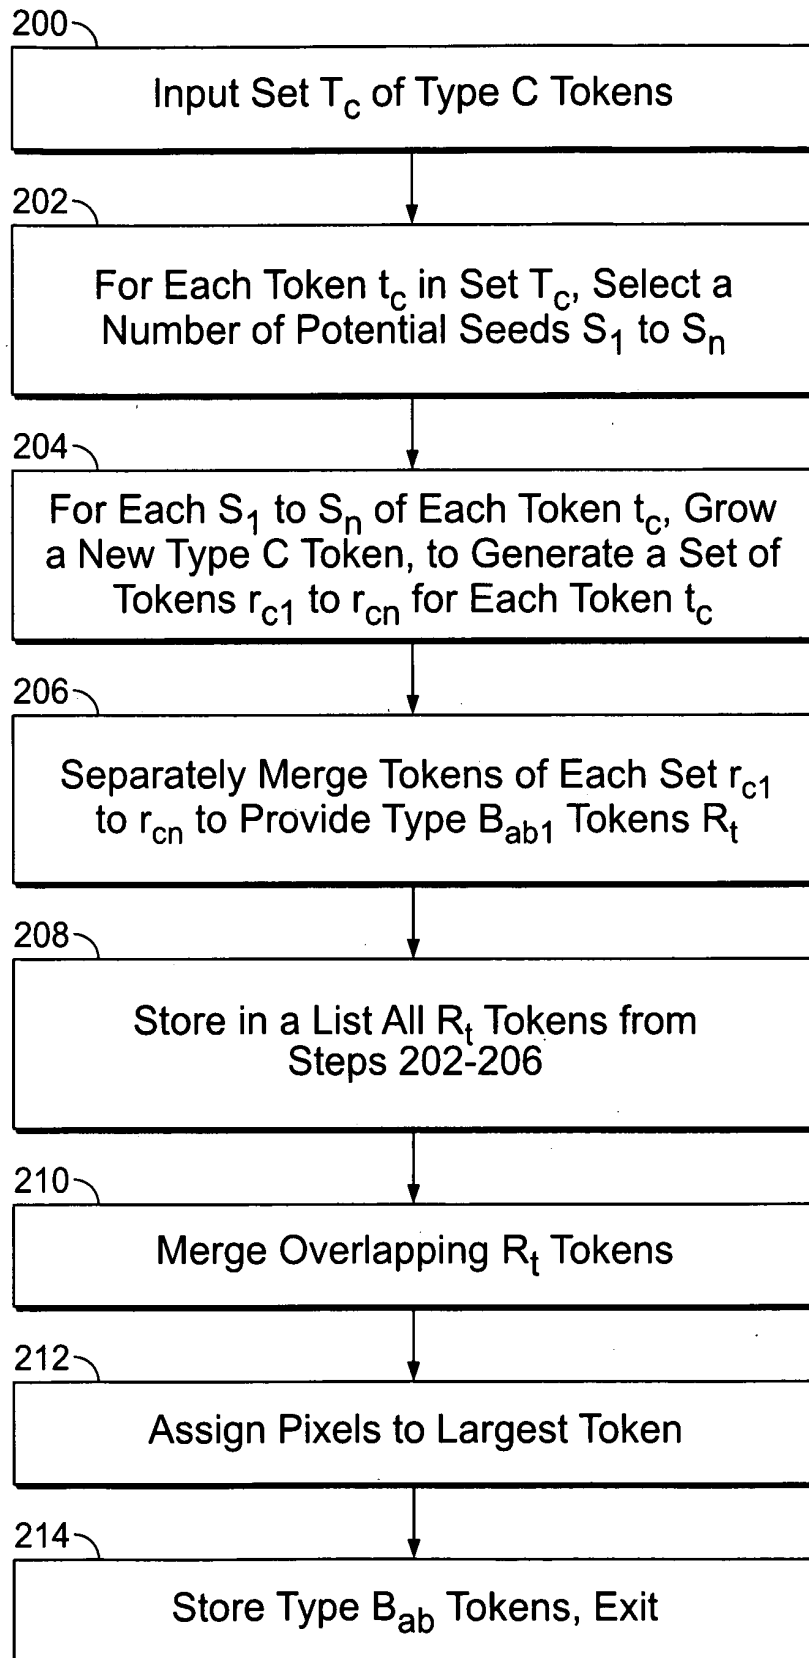


FIG. 8

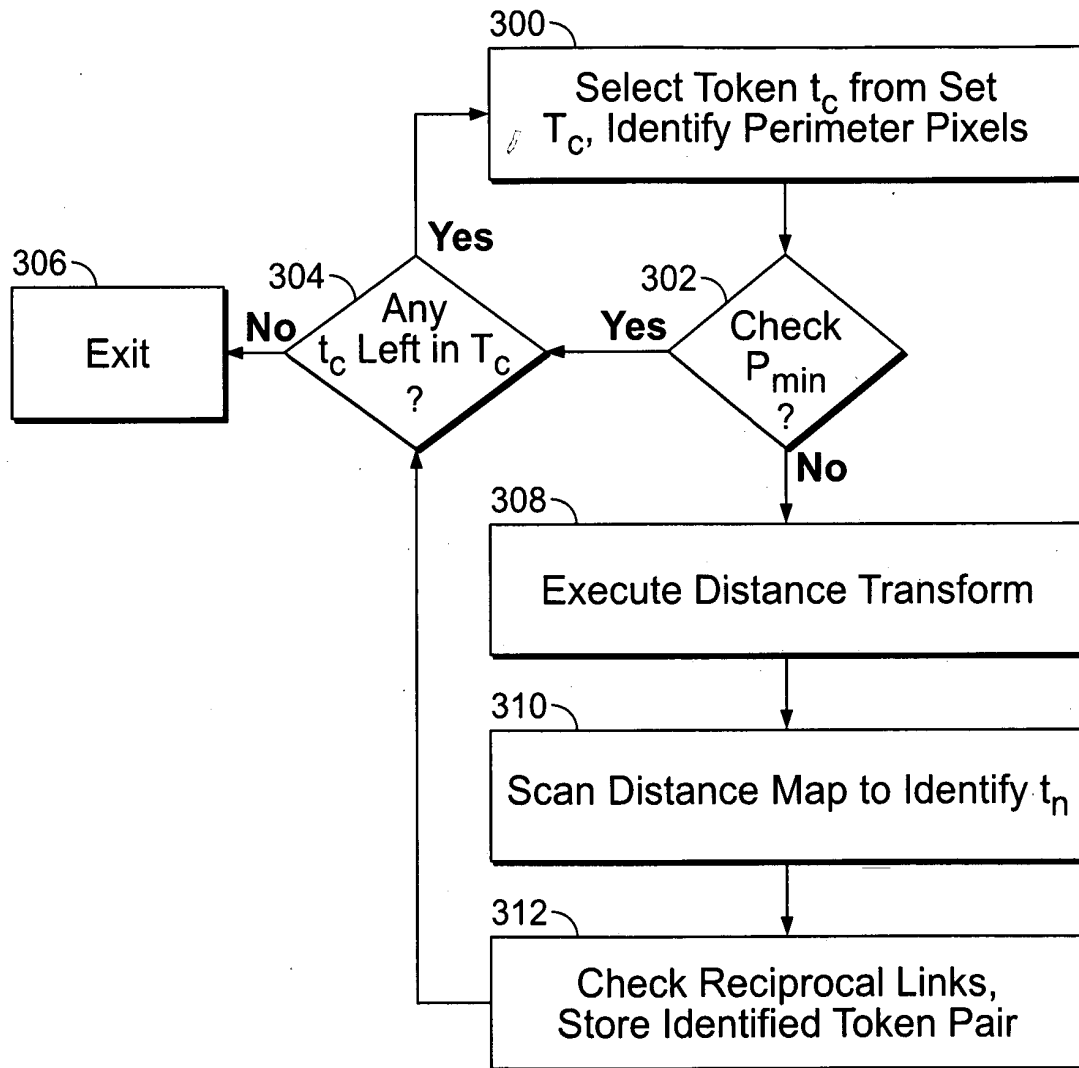


FIG. 9

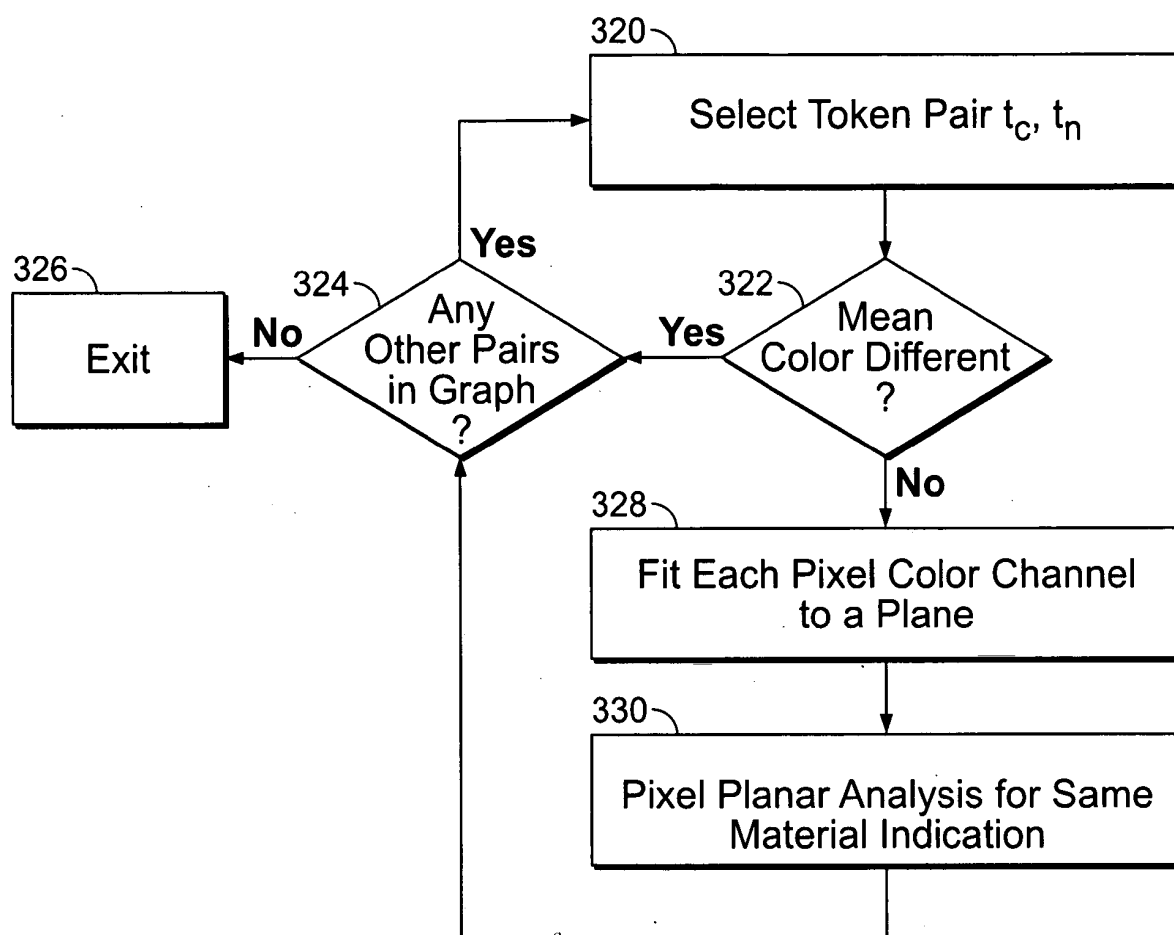
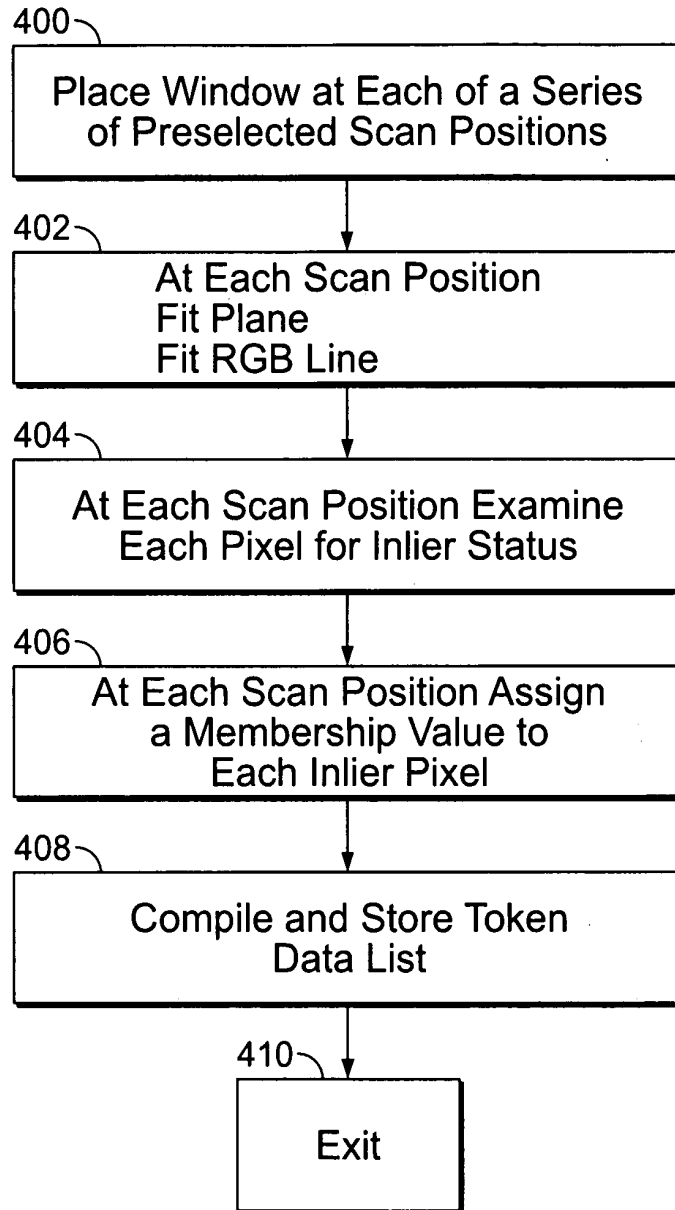


FIG. 10

**FIG. 11**

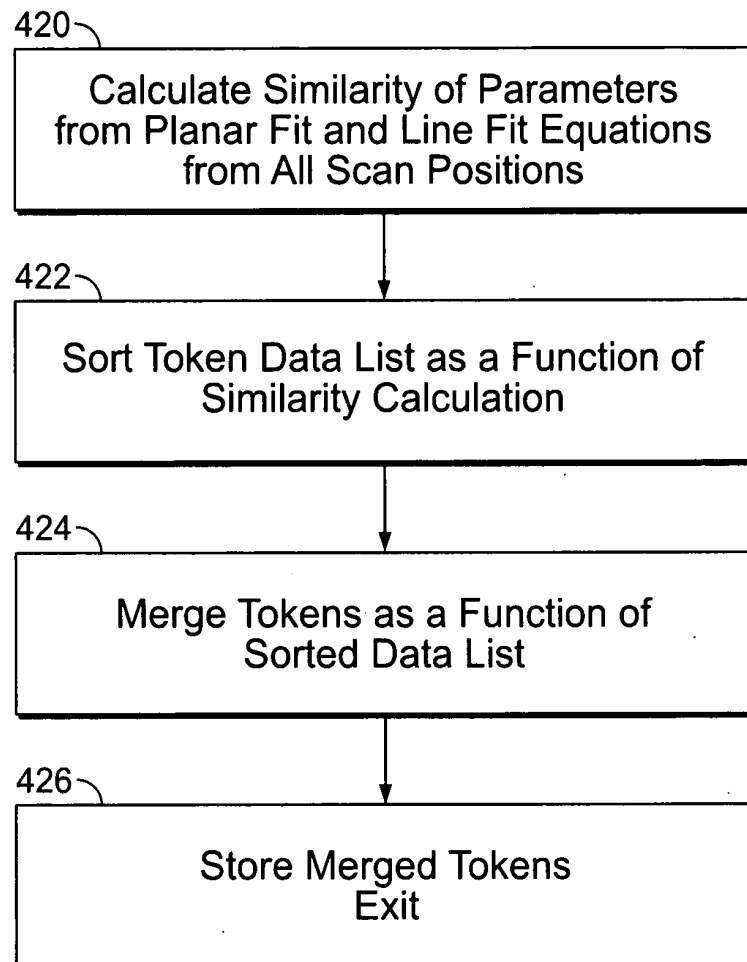



FIG. 12



$$\begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \\ 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} m_a \\ m_b \\ m_c \end{bmatrix} = \begin{bmatrix} i_a - i_b \\ i_a - i_c \\ i_b - i_c \end{bmatrix}$$

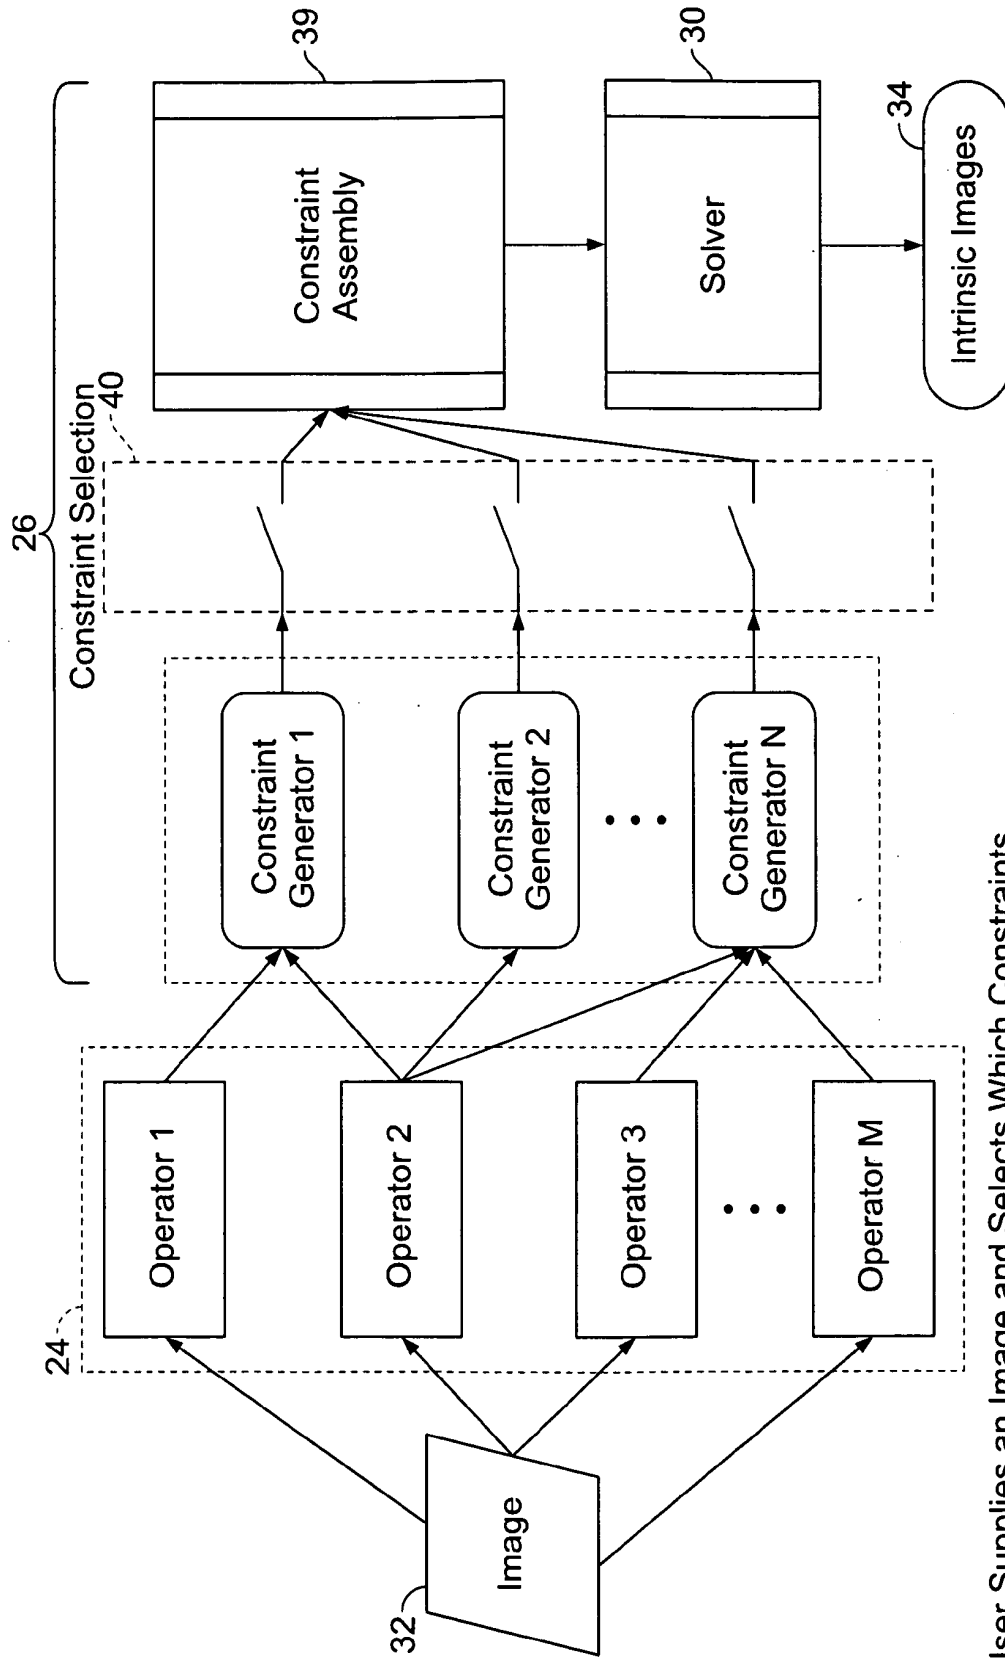
[A]

[x]

=

[b]

FIG. 13



User Supplies an Image and Selects Which Constraints.
 The Operators Process the Image Data.
 The Constraint Generators Use the Operators to Create Individual Constraints.
 The Constraints are Assembled into one Consistent Set and Passed to the Solver.
 The Solver Produces the Intrinsic Components of the Input Image.

FIG. 14

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 34175106 A [0033]
- US 20070176940 A [0033]
- US 40371906 A [0035]
- US 20070242878 A [0035]
- US 34174206 A [0052]
- US 20060177149 A [0052]
- US 07987808 A [0053]

Non-patent literature cited in the description

- Learning Face Appearance under Different Lighting Conditions. **BRENDAN MOORE et al.** Biometrics: Theory, Applications and Systems, 2008, BTAS. IEEE, 2008 [0002]
- Estimating Intrinsic Component Images using Non-Linear Regression. **MARSHALL F TAPPEN et al.** Conference on Computer Vision and Pattern Recognition. IEEE, 2006, vol. 2, 1992-1999 [0003]
- **RANDEN, T.**[TRYGVE ; **HUSØY, J.H.** [**JOHN HÅKON**. Filtering for Texture Classification: A Comparative Study. *PAMI*, 04 April 1999, vol. 21, 291-310 [0077]
- *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2003, vol. 2, 691-698 [0077]
- **P. FELZENSZWALB ; D. HUTTENLOCHER**. *Distance Transforms of Sampled Functions*, Cornell Computing and Information Science Technical Report TR2004-1963, September 2004 [0115]
- **ABDI, H.** Encyclopedia of Measurement and Statistics. Thousand Oaks, 2007 [0119]
- **PETER J. ROUSSEEUW**. *Least Median of Squares Regression*, *Journal of the American Statistical Association*, December 1984, vol. 79 (388), 871-880 [0121]
- **ZVI GALIL ; GIUSEPPE F. ITALIANO**. *Data structures and algorithms for disjoint set union problems*, *ACM Computing Surveys*, September 1991, vol. 23 (3), 319-344 [0123]
- **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 [0127]