



(11) **EP 3 373 245 A2**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
12.09.2018 Bulletin 2018/37

(51) Int Cl.:
G06T 7/00 (2017.01) **G06T 15/00 (2011.01)**
G06F 19/00 (2018.01) **G06T 15/08 (2011.01)**

(21) Application number: **18160956.1**

(22) Date of filing: **09.03.2018**

(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 RS SE SI SK SM TR
Designated Extension States:
BA ME
Designated Validation States:
KH MA MD TN

- **MIAO, Shun**
Princeton, NJ New Jersey 08540 (US)
- **PETKOV, Kaloian**
Lawrenceville, NJ New Jersey 08648 (US)
- **SUDARSKY, Sandra**
Bedminster, NJ New Jersey 07921 (US)
- **YU, Daphne**
Yardley, PA Pennsylvania 19067 (US)
- **MANSI, Tommaso**
Plainsboro, NJ New Jersey 08536 (US)

(30) Priority: **10.03.2017 US 201762469703 P**
03.01.2018 US 201815827263

(74) Representative: **Patentanwälte Bals & Vogel**
Sendlinger Strasse 42A
80331 München (DE)

(71) Applicant: **Siemens Healthcare GmbH**
91052 Erlangen (DE)

(72) Inventors:
• **LIU, Chen**
1024 Ecublens (CH)

(54) **CONSISTENT 3D RENDERING IN MEDICAL IMAGING**

(57) For three-dimensional rendering (32), a machine-learned model is trained (14) to generate representation vectors for rendered images formed with different rendering parameter settings. The distances between representation vectors of the images to a reference are used to select (36) the rendered image and corresponding rendering parameters that provides a consistency

with the reference. In an additional or different embodiment, optimized pseudorandom sequences are used for physically-based rendering (52). The random number generator seed is selected (46) to improve the convergence speed of the renderer and to provide higher quality images, such as providing images more rapidly for training compared to using non-optimized seed selection.

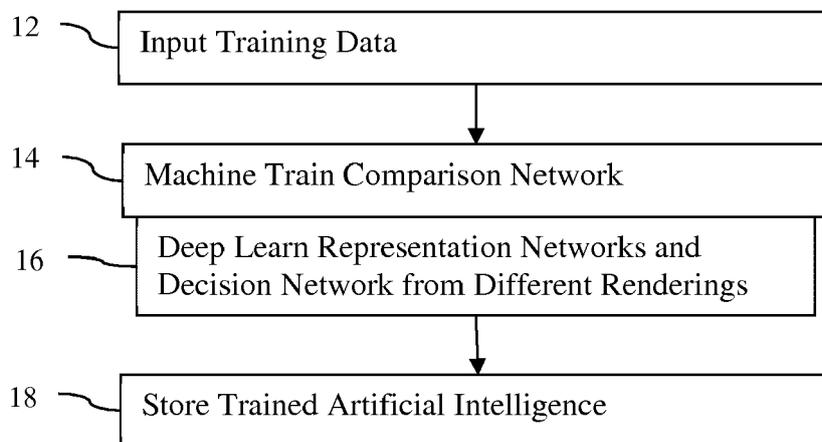


FIG. 1

EP 3 373 245 A2

Description

BACKGROUND

[0001] The present embodiments relate to medical imaging of three-dimensional (3D) scans. Data representing a volume is rendered for visualization. Due to the many different scan settings and patient variability, renderings for different patients or at different times appear different.

[0002] To reduce variability, most existing medical rendering provides a set of static visualization presets for specific workflows or diagnostic contexts. Even with static presets, the rendering may require manual adjustment and may not provide consistent rendered results. Patient and scanner variability may also continue to contribute to inconsistency. Semi-automatic machine learning approaches have been used to provide transfer function design in traditional volume rendering (e.g., ray casting or alpha blending), but may still not sufficiently reduce variability.

[0003] Variability is particularly troublesome for physically-based volume rendering, which relies on the physical simulation of light propagation (e.g., unbiased path tracing). Physically-based visualization techniques produce global illumination effects in computer-generated graphics that mimic the real-world interaction of light with various 3D objects or tissues. This results in physically plausible images that are often easier for the human brain to interpret when compared to the more analytical images from traditional rendering. These physically-based visualization techniques are more sensitive to changes in the classification or the interpretation of the underlying medical data. As a result, small changes to the rendering parameters may have a more pronounced impact on the perception and interpretation of 3D structures in the final image as compared to the traditional volume rendering techniques. Obtaining very high quality reproducible images with diagnostic significance is then partially at the user's discretion. While existing techniques for providing visualization presets may help, the resulting rendered images may not have consistent quantitative properties (color, hue, reflectance, etc.) across datasets. Physically-based rendering also takes longer to generate a rendered image, so alteration to reduce variability is time consuming.

SUMMARY

[0004] By way of introduction, the preferred embodiments described below include methods, systems, instructions, and computer readable media for three-dimensional rendering. A machine-learned model is trained to generate representation vectors for rendered images formed with different rendering parameter settings. The distances between representation vectors of the images to a reference are used to select the rendered image and corresponding rendering parameters that provides a con-

sistency with the reference. In an additional or different embodiment, optimized pseudo-random sequences are used for physically-based rendering. The random number generator seed is selected to improve the convergence speed of the renderer and to provide higher quality images, such as providing images more rapidly for training compared to using non-optimized seed selection.

[0005] In a first aspect, a method is provided for three-dimensional rendering in a rendering system. A medical imaging system acquires a medical dataset representing a three-dimensional region of a patient. A renderer renders, using different combinations of rendering settings, a plurality of images from the medical dataset. An image processor applies different pairs of the images of the plurality and a reference image to a machine-learned model. The machine-learned model is trained to generate representation vectors of the images and the reference image. The image processor selects one of the plurality of images based on the representation vectors. The selected one of the plurality of images is displayed.

Preferred is a method, wherein acquiring comprises acquiring from a magnetic resonance, computed tomography, ultrasound, single photon emission computed tomography, or positron emission tomography system as the medical imaging system.

Preferred is a further method, wherein rendering comprises rendering with path tracing emulating photon interaction with the medical dataset. Rendering can comprise rendering with Monte Carlo path tracing using a random number generator seed identified based on an optimization. Preferred is a method, wherein rendering comprises rendering with the rendering settings comprising transfer function, window function, shading, and clip plane position, and wherein a first subset of the rendering settings is held constant in the rendering of the plurality of images and a second subset of the rendering settings are varied.

Further, preferred is a method, wherein applying comprises applying with the machine-learned model comprising a machine-learned neural network or wherein applying comprises applying with the machine-learned neural network comprising a convolutional neural network taught to output a feature vector for each of the images of the pair and the reference image or wherein applying comprises applying with the machine-learned model trained with ground truth from a perceptual hash metric, or wherein applying comprises applying with the machine-learned model trained with random variation of the rendering settings and camera positions.

Preferred is a method, wherein selecting comprises selecting based on Euclidian distances of the representation vectors and/or wherein selecting comprises selecting the one as most similar to the reference image, and wherein displaying comprises displaying the selected one as an image of the patient on a display and/or wherein selecting comprises searching with optimization applied to the different combinations of the rendering settings

through the rendering and applying.

[0006] In a second aspect, a system is provided for three-dimensional rendering in medical imaging. A medical scanner is configured to scan a volume of a patient. A graphics processing unit is configured to three-dimensionally render images from the scan of the volume. The images are rendered with different values for rendering parameters. An image processor is configured to search for one of the images matching a reference image based on a machine-learned neural network. A display is configured to display the one image of the volume of the patient. Preferred is a system, wherein the graphics processing unit is configured to three-dimensionally render the images with path tracing or wherein the graphics processing unit is configured to render with a random number generator seed, the random number generator seed having been selected in an optimization.

Further, a system is preferred, wherein the image processor is configured to search with the machine-learned neural network including a convolutional neural network for the reference image and the images and configured to determine distances between representation vectors output by the convolutional neural network.

[0007] In a third aspect, a method is provided for three-dimensional rendering in a rendering system. A reference image is three-dimensionally rendered with path tracing using a first number of samples per pixel and a first random number generator seed. Other images are three-dimensionally rendered with the path tracing using a second number or less of samples per pixel and other random number generator seeds. The second number is less than the first by a factor of ten or more. An error between the other images and the reference image is measured. One of the other random number generator seeds with a lessor error than the others of the random number generator seeds is identified. The rendering system is configured for path tracing with the identified one of the random number generator seeds. Preferred is such a method, further comprising scanning a patient with a medical imaging system, and three-dimensionally rendering, by the rendering system as configured, an image from output of the scanning with path tracing using the second number or less of the samples per pixel and the identified one of the random number generator seeds. The method can further comprise performing the measuring and identifying as part of interactive rendering from the output of the scanning.

Preferred is a method, wherein measuring the error comprises measuring the number of samples per pixel where the error reaches a threshold.

Preferred is a method, wherein three-dimensionally rendering other images comprises emulating coherent scattering.

Preferred is further a method, the method further comprising training a neural network with deep learning to identify similar images with an architecture using a representation network connected to a decision network, the deep learning using training renderings generated with

the identified one of the random number generator seeds.

[0008] The present invention is defined by the following claims, and nothing in this section should be taken as a limitation on those claims. Further aspects and advantages of the invention are discussed below in conjunction with the preferred embodiments and may be later claimed independently or in combination.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The components and the figures are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention. Moreover, in the figures, like reference numerals designate corresponding parts throughout the different views.

Figure 1 is a flow chart diagram of one embodiment of a method for machine training for a comparison network;

Figure 2 is a flow chart diagram of an embodiment of a method for rendering a reference consistent image;

Figure 3 illustrates a neural network architecture for a comparison network;

Figure 4 is a flow chart diagram of an embodiment of a method for improving physically-based rendering speed by seed selection;

Figures 5A-D show different example physically-based renderings;

Figures 6A and 6B show different path samplings due to differences in seed;

Figure 7 shows example perceptual errors based on seed; and

Figure 8 is a block diagram of one embodiment of a system for three-dimensional rendering.

DETAILED DESCRIPTION OF THE DRAWINGS AND PRESENTLY PREFERRED EMBODIMENTS

[0010] Two general approaches are described herein. In one approach, a machine-learned representation network provides for consistent three-dimensionally rendered medical images. This machine-learned representation approach is described below with respect to Figures 1-3. In another approach, pseudo-random sequences for physically-based rendering are optimized. This pseudo-random sequence approach is described below with respect to Figures 4-7. These different approaches may be used separately. In other embodiments, these different approaches are combined, such as using the increased rendering speed provided by optimization of the pseudo-random sequence to generate many samples for training for the representation network.

[0011] For machine-learned comparison, an unsupervised neural network maps an image x to a vector representation, denoted $f(x)$. The vector representation has a property. The Euclidean distance between representations of two images is correlated to the difference of

their visual effects (i.e., representations of images with similar visual effect will have small Euclidean distance). This representation is used to automatically select rendering parameters of the photorealistic (i.e., physically-based) or other renderer to produce desired visual effects.

[0012] Consistent visual interpretation for 3D images may potentially simplify diagnostic workflow and may promote the clinical use of 3D data for decision support. Consistent photorealistic rendered images may be provided in new platforms and workflows, such as fully automated reporting and cloud-based rendering services. Consistency may be provided based on an expert or practice selected reference. A general, adaptable approach using a machine-learned comparison to a reference is scalable, allowing variance in any rendering parameter. The same approach may be used in a wide variety of clinical problems or circumstances.

[0013] Figure 1 shows one embodiment of a method for machine training for comparison. Optimal rendering parameters may be obtained by representation learning to compare different rendered images with a reference image. The input of the artificial intelligence system includes multiple images to compare to a reference, so the training learns to indicate which of the multiple images is most like the reference. Using many examples, the machine learns to provide a scalar indicating the more similar of the multiple images to the reference. This training creates a representation network that provide vectors that may be compared to identify perceptual similarity. The trained network maps the input image to a vector with the property that the Euclidean distance of two vectors correlates with the visual effect difference between the two images. The most similar rendering to the reference is obtained by searching the rendering that leads to the minimum Euclidean distance. The resulting image of the multiple images more likely includes diagnostically useful information or other characteristic reflected in the reference, so is used for diagnosis.

[0014] The method of Figure 1 is implemented by a machine training system, such as by an image processor and a graphics processing unit in a server, computer, or other machine. In one embodiment, the system of Figure 8 implements the method of Figure 1, with or without the medical scanner and/or display. For example, data is input by loading from a memory, the graphics processing unit (GPU) renders images to be compared, the image processor performs the training based on the data, and the memory stores the learnt artificial intelligence. As another example, the medical scanner and/or memory input the data to a server, which performs the training and outputs to a memory.

[0015] The acts are performed in the order shown (top to bottom) or another order. For some forms of learning, results may be fed back as input in a repetition of the acts.

[0016] Additional, different, or fewer acts may be provided. For example, act 18 is not provided, such as where the trained artificial intelligence is used rather than stored

for later use.

[0017] In act 12, training data is input. The training data is gathered from one or more memories, sensors, and/or scanners. The gathered data may be stored and/or transmitted to a buffer, memory, cache, processor, or other device for training. This gathered training data is input to the machine learning algorithm.

[0018] Many samples of the same type of data are input. To learn statistical or other relationships that may be sufficiently complex that a human cannot perceive the relationships in a same level, tens, hundreds, or thousands of samples are provided.

[0019] Any data may be used. For representation learning, a pool of three-dimensional volumes is collected. The volumes are scan volumes of one or more patients. The volumes are for a given application, such as vascular computed tomography (CT). Alternatively, the volumes are generic to the application, such as being whole body volumes or volumes representing different parts.

[0020] Since the artificial intelligence is to be trained to provide relative similarity or comparison with a reference, three-dimensionally rendered images are used as the training data. For each sample, an image rendered from the scan data representing a patient is used. Any rendering may be used, such as ray casting, surface rendering, projection, maximum intensity projection, or physically-based rendering (i.e., path tracing). Physically-based rendering simulates light propagation through the volume. This is different than illumination or shading used in ray casting or surface rendering, instead using pseudo-random simulation of different paths that may be followed by light. Any path tracing, such as unbiased path tracing, Monte Carlo rendering, global illumination, or other simulations of light propagation may be used.

[0021] Since the representation learning is to compare two different images to a reference image, the samples used for training are grouped in sets of three. Where representations for additional images are included in the network architecture, samples in other size groupings (e.g., four images) may be provided.

[0022] In one embodiment, six images are rendered for each volume, V . Three images are rendered using the same camera position Φ_0 , and with random rendering parameters θ_i , $i = 1, 2, 3$. The images are denoted as $R(V, \Phi_0, \theta_i)$, $i = 1, 2, 3$. Then, three camera positions Φ_i , $i = 1, 2, 3$ are randomly generated, and three images $R(V, \Phi_i, \theta_i)$, $i = 1, 2, 3$ are rendered for the three camera positions. The rendering parameters used for the images from the three different camera positions are the same rendering parameters used for the images from the same camera position. The six images generated in this way are referred to as a "dataset." The process is repeated to generate N datasets for each 3D volume, V . Other numbers of images, sets of random rendering parameters, numbers of camera positions, and/or numbers N (e.g., tens, hundreds, or thousands) may be used. The same volume may be used to create more than one dataset or set of

six images.

[0023] Any rendering parameter may be varied. Camera position is treated separately, but may be a rendering parameter. Other rendering parameters include the transfer function, window, bounding box, scaling, level compression, data normalization, type of virtual lights, position of the virtual light sources, orientation of the virtual light sources, image-based lighting sources, and/or clip plane position. For physically-based rendering of photorealistic images, additional rendering parameters include material properties, such as reflectance, a global albedo, transparency, and/or surface roughness. Other renderer-specific parameters may be used, such as the seed for random number generation in Monte Carlo path tracing.

[0024] The values for one or more rendering parameters are randomly selected. Non-random or patterned selection of the values may be used. Some, one, or all but one rendering parameters may be assigned a constant value, such as changing the transfer function while maintaining the clip plane in a user set position for each of the renderings. A range of values or group of values may be used for random or patterned selection to generate the images with different settings of the rendering parameters. As discussed below for selection, the value or values may be determined based on an optimization.

[0025] To generate the samples, many renderings are performed. Since the renderings are for training, any amount of time may be used to render. Physically-based rendering is more processor and time consuming than ray casting. To allow a greater number of renderings and corresponding samples for training and/or reduce the time needed for physically-based rendering, a random number seed optimized to allow rendering with fewer light-path samples is used. For example, a seed optimized as discussed below for Figure 4 is used.

[0026] For training, the goal or ground truth information is input. The artificial intelligence system is trained with labeled data (e.g., input triplets of images and corresponding desired output). For training for comparison, triplets of images are input. The ground truth is which of two of the images is more similar to the third. This binary ground truth is a 0 or negative for one image and a 1 or positive for the other image. In other embodiments, a non-binary scalar is used, such as a relative amount of similarity to the reference.

[0027] The images from the different camera angles for each triplet sample are assigned the same ground truth. The triplet with different angles is assigned the same ground truth as the triplet with the same angle. To train the network, triplets of images from different camera positions and their ground truth relationship are acquired.

[0028] The ground truth is calculated from the triplets. Denoting the three images as R (reference), A, B, the ground truth relationship tells if A or B is more similar to R. The ground truth relation is obtained in the following steps. Given the triplet of images from the same camera position, denoted as R', A' and B', the perceptual hash

values between the R' and A', and R' and B' are computed. A smaller perceptual hash value indicates a similar visual effect and indicates if A' or B' is more similar to R'. The perceptual hash may be calculated on a pair of images that are aligned (e.g., images from the same camera position) and may not be directly applied on the triplet with difference camera positions. Since R, A, B have the same rendering parameters as R', A', B', and the same rendering parameters indicate same visual effect, the ground truth relationship calculated for R', A', B' is directly transferred to R, A, B.

[0029] The perceptual hash metric is denoted as $PH(\cdot, \cdot)$. For each dataset, a level of perceptual similarity between each of two images (2 and 3) to a reference image (1) of the triplet from a same camera angle. These two perceptual hashes are calculated as: $D_{12} = PH(R(V, \phi_0, \theta_1), R(V, \phi_0, \theta_2))$ and $D_{13} = PH(R(V, \phi_0, \theta_1), R(V, \phi_0, \theta_3))$. While the perceptual hash value may be used as the ground truth, the perceptual hash values are used to generate the binary ground truth in other embodiments. The binary ground truth or scalar value, y , is calculated as:

$$y = \begin{cases} 1, & \text{if } D_{12} < D_{13} \\ 0, & \text{otherwise} \end{cases}$$

If y is 1, then image 3 is more similar to the reference image 1 than image 2. If y is 0, then image 2 is more similar to the reference image 1 than image 3. Other metrics than perceptual hash may be used, such as other perceptual metrics, minimum sum of absolute differences or other correlation.

[0030] Based on the provided datasets and calculated ground truth, the artificial intelligence system is trained to produce the desired output from any input. In act 14, the artificial intelligence or intelligence is machine trained. The collected samples (e.g., the triplets from different camera angles), including ground truth (e.g., binary relative perceptual similarity calculated with perceptual hash), are used to train a comparison network. A machine, such as an image processor, computer, server, or other device, learns from the samples to indicate which of two input images are more similar to a reference image. Using machine-learning, complex statistical relationships between large numbers (e.g., tens, hundreds, thousands, or more) of input variables to any number of output variables are extracted from the large number of samples.

[0031] The training is for imaging in any context. One model may be learned for any number of imaging situations. Alternatively, different models are trained for different situations. The different situations may include different scan modalities (e.g., different model for computed tomography, magnetic resonance, ultrasound, positron emission tomography, and single photon emission computed tomography). The different situations may include different types of tissue of interest (e.g., liver versus kidney), different diagnostic purpose or workflow (e.g., cancerous lesion versus bone calcification), and/or different

users (e.g., different operators may have different preferences for visualization).

[0032] Any now known or later developed machine learning may be used. Regression, classification, and/or reinforcement learning are used. Regression training learns a range or continuous output by minimization of a metric. Classification learns disparate outputs. Reinforcement learning learns a sequence of actions with feedback. Neural network, Bayes network, or support vector machine training may be used. Hierarchical or other approaches may be used. Supervised or semi-supervised machine learning may be used.

[0033] To train, features are extracted from the input data. Haar wavelet, steerable, or other features may be extracted from the scan dataset or images. Alternatively, the input data itself (e.g., pixel or color values of the rendered image) is used and the learning determines features, such as with deep learning.

[0034] Act 16 represents an embodiment of act 14. The comparison network is trained with deep learning. The deep learning provides representative learning. The training learns convolution kernels from the input images to generate a representation that may be used for a comparison decision. Deep learning (e.g., deep structured learning, hierarchical learning, or deep machine learning) models high-level abstractions in data by using multiple processing layers with structures composed of multiple non-linear transformations, where the input data features are not engineered explicitly. A deep neural network processes the input via multiple layers of feature extraction to produce features used to derive comparison results. The deep learning provides the features used by other machine training to learn to output the comparison. Other deep learnt, sparse auto-encoding models may be trained and applied. The machine training is unsupervised in learning the features to use and how to classify given an input sample (i.e., feature vector).

[0035] Using deep learning, the input features required for the machine learning are not engineered by the user explicitly, as in shallow learning algorithms. This is especially useful with render settings, which may not have an explicit natural relationship with the input data. Deep learning figures out the data features that end up providing a good classifier or model for the comparison network. Deep learning is effective in extracting high level abstractions (features) from low level input data (e.g., raw images) that is suitable for the trained task, so may be more convenient than and superior to hand-crafted features for identifying features for render settings.

[0036] Any neural network architecture for the deep learning may be used. The neural network architecture defines the neural network used to train and the resulting trained or machine-learned comparison model. For example, a set of layers (e.g., representation network) is used to generate features at any level of abstraction from the input images, and another layer or group of layers is used to make a comparison based on the values of the features.

[0037] Figure 3 shows one example architecture used for the comparison network. A convolutional neural network (a.k.a., representation network) 20 is created. The same convolutional neural network is used for each input image of the triplet. The representation network 20 is shown in Figure 3 three times, once for each of a triplet of images (i.e., reference image I_r , comparison image 1, I_1 , or comparison image 2, I_2). The representation network 20 is trained to output a vector of a fixed length. In the example of Figure 3, the vector includes 100 values for images that are 256x256 pixels. The convolutional neural network is denoted as $f(v)$, and its output vector, v , is referred to as the "representation" of the image (i.e., an image vector). In the example of Figure 3, the representation network 20 has five convolutional and pooling layers and two fully connected layers, but other layer arrangements may be used.

[0038] The architecture includes a network for comparing based on the vectors from the representation network. A neural network (a.k.a., decision network) 22 receives as input the three "representations", denoted as $f(v_1)$, $f(v_2)$, $f(v_3)$, and learns to output a scalar between 0 and 1. The scalar represents the possibility that the visual effects of v_1 and v_2 are closer than that of v_1 and v_3 . In the example of Figure 3, the neural network 22 for the decision network is formed of three fully connected layers, but other layer arrangements may be used. The decision network 22 is used for training the representation networks.

[0039] Softmax represents the function where the scalar values for the two comparison images sums to 1. For example, the decision network learns to output two values that sum to 1, such as outputting 0.25 and 0.75 for one triplet and 0.4 and 0.6 for another triplet. In another embodiment, one scalar value is output for a comparison with the difference from 1 being assigned as the scalar value for the other comparison. Other output functions may be used, such as a binary indication for each comparison image or other values that may not sum to 1.

[0040] The entire neural network is formed by the arrangement of the three separate representation networks without outputs connected to the decision network. This neural network architecture takes three images as input, and outputs a scalar value. The neural network is trained using $R(V, \Phi_i, \theta_i)$, $i = 1,2,3$ as input, and y as output. Other architectures may be used, such as replacing the decision network with triplet loss options designed to train the network with triplets.

[0041] Any deep learning may be used. The representation and decision networks are trained together, but may be separately trained. In one embodiment shown in Figure 3, a cross-entropy loss function is used to train. Other functions measuring the difference of a current output from the ground truth and optimizing the neural network to minimize the difference may be used.

[0042] The machine trains to output a comparison that may be used to find a rendered image most similar or sufficiently similar (e.g., within a threshold amount based

on the scalar value) to a reference image. The machine-learned comparison model may be used to provide consistent imaging. The path tracing, other physically-based rendering, or other rendering with settings resulting in a similar or consistent image is identified using the trained comparison model.

[0043] In act 18, the trained artificial intelligence is stored. The result of the training is series of tensors. Each tensor represents the weights of one layer of the neural network. The tensors represent the learned knowledge through machine training. In one embodiment, the representation network is stored. The decision network is used to train, but not for application so is not stored.

[0044] The trained representation model is stored in a memory. Any memory may be used. The memory used for the training data may be used. For application, the memory may be in another device. For example, the trained model is stored in a memory of a medical imager or workstation. The medical imager or workstation uses the trained model to select an image and/or corresponding rendering parameter settings. A server implementing one copy of the trained comparison model may be used to identify for different patients. Multiple copies of the trained model may be provided to different physicians, medical scanners, and/or workstations for use by different physicians for different patients.

[0045] Figure 2 shows one embodiment of a method for three-dimensional rendering in a rendering system. The method applies a machine-trained representation model or neural network to identify an image or rendering settings for producing an image consistent with a reference.

[0046] The rendering system is the system of Figure 8 or a different system. For example, the rendering system includes a renderer (e.g., GPU) for rendering images and an image processor for selecting an image using the machine-learned comparison model.

[0047] The acts are performed in the order shown (top to bottom or numerical) or other orders. Additional, different, or fewer acts may be provided. For example, acts for optimizing the seed selection for the random number generator in path tracing rendering are provided. As another example, the selected image is stored or transferred over a network instead of or in addition to display in act 38. In another example, acts 32 and 34 are repeated any number of times before or as part of performing act 36.

[0048] In act 30, a medical scanner acquires medical data. The medical scanner provides medical data representing a 3D region of the patient. The medical scanner may directly provide the medical data, such as providing in an image processing pipeline of the medical scanner. The medical scanner may indirectly provide the medical data, such as routing the medical data through a memory or computer network. The medical data may be accessed from a picture archiving and communications server.

[0049] The medical data is for a given patient. While the training may use samples from many patients, the

learned model is applied to the medical data for a patient to identify an image or corresponding rendering settings for consistent imaging of that patient. The same or different machine-learned models are applied to data for any number of patients.

[0050] Any type of medical data is obtained. Scan data representing a volume is loaded. The scan data may be from multiple two-dimensional scans or may be formatted from a three-dimensional scan. The medical data represents a three-dimensional region of the patient, such as being a set of voxel values for voxels distributed along three spatial dimensions.

[0051] Any imaging modality and corresponding medical imaging system may be used. Example medical scanners include magnetic resonance, computed tomography, ultrasound, positron emission tomography, and/or single photon emission computed tomography.

[0052] In act 32, a renderer renders a plurality of images from the medical dataset. The volume of the patient is rendered in different ways. Different combinations of rendering settings are used. One or more rendering settings may be constant for rendering the images, but at least one rendering parameter is varied. The training may vary the same or greater number of rendering parameters. Alternatively, all the rendering parameters may be varied.

[0053] The settings may be randomly assigned or assigned using a pattern, such as rendering images with a representative sampling of the possible settings. The user may control or select the settings for parameters and/or parameters to be varied. Any number of images may be rendered, such as five or more, tens, or hundreds. For each combination of rendering settings, an image is rendered.

[0054] Any type of rendering may be used, such as the same or different type of rendering used for training. For example, ray casting is used. As another example, a physically-based renderer renders photorealistic images of the volume region of the patient using different sets of values for the rendering parameters. The physically-based renderer is a graphics processing unit, array of graphics processing units, workstation, server, computer, or other machine capable of the required computations. Any physically-based rendering algorithm capable of computing the light transport may be used, such as path tracing, unbiased path tracing, Monte-Carlo path tracing, or Metropolis light transport. The physically-based rendering simulates the physics of light propagation (e.g., emulates photon interaction) to create an image instead of accumulating volume samples with local illumination along viewing rays from the camera, as is done for traditional volume rendering. In path tracing implemented for physically-based rendering, the paths of light or photons, including due to scattering and/or absorption, are modeled and used for rendering. Figures 6A and 6B show example light paths, including many paths having at least one instance of scattering as represented by a change in direction. The physically-based

rendering result may be built up over time as the rendering may rely on probabilistic scattering and tracing many light paths. For each pixel of the image being rendered, any number of light path samples are used. For example, 16-2000 samples are used for each pixel.

[0055] In one embodiment, the renderer renders with Monte Carlo path tracing. The random number generator of the Monte Carlo path tracing generates a sequence of pseudo random numbers used to determine scattering or the paths of light. One seed may result in fewer samples needed before a rendered image provides sufficient information. Using the method of Figure 4, various seeds may have been tested in an optimization. The seed used for rendering for a given patient is selected as part of a previous optimization so that the images may be rendered more rapidly given a same rendering capability than if a different seed were used. Non-optimized seeds may be used for rendering. The seed optimization may be performed as part of rendering for a given patient, such as using different seeds for different renderings and testing as described for Figure 4.

[0056] The rendering results in a photorealistic image. For a given rendering or settings of rendering parameters, a sequence of images may be provided as the image is built or rendered. Alternatively, for a given set of values of rendering parameters, a single image is output. Multiple, such as ten, twenty, or more images are rendered using different settings. A sequence or group of images due to alteration of values for one or more rendering parameters may be output.

[0057] In act 34, an image processor applies the rendered images for the volume and a reference image (i.e., applies a triplet) to the representation network part of machine-learned model. Any reference image may be used. For example, an expert selects or creates a reference image. As another example, a hospital, medical association, physician, or other person determines an ideal or desired image for a given application (e.g., CT vascular). A previously created image for a given physician may be used. An image from a clinical study, literature, or medical atlas may be used. In one embodiment, a panel or plurality of possible reference images are displayed, and a user selects the reference image from the panel. In another embodiment, a previous image of the patient, such as from a previous scan or examination, is used.

[0058] The rendered images are applied in sequence. As each image of the new or current volume is rendered, the new image is compared with the reference image.

[0059] For applying the images in the representation part of the machine-learned model, the images are processed by the image processor, such as a computer, medical scanner, workstation, or server. The image processor uses the internal representation of the representation network of the machine-learned model. Given a reference image provided by the user, denoted as R , the representation network is applied on the image to get a vector, denoted as $f(R)$. For any given rendering parameter, θ ,

the rendered image is denoted as $X(\theta)$. The rendered image is applied to the representation network to get a vector $f(X(\theta))$. The image processor inputs the images, resulting in the machine-learned model outputting the vectors from the representation network for the two images.

[0060] Where the reference image is previously determined, the representation V_r may be previously determined. Alternatively, the reference image is applied at a same time as the rendered images for the patient.

[0061] The machine-learned model was previously learned as discussed above. The machine-learned model was trained to indicate ones of the pairs of images that is more similar to the reference image. The decision network part of the machine-learned model is not used for application. Instead, the representation network is used as the machine-learned model to provide vectors representing the reference image and images rendered differently. Random or other variation in rendering settings and/or camera positions are used to train the machine-learned model to compare with common or possible variation in rendering. The trained model is used to find representations that reflect similarity in perception. In one embodiment, the machine-learned model is a neural network with an architecture shown as the representation network 20 in Figure 3. The convolutional neural network taught to output a feature vector for each of the images is used. The machine-learned model was trained with any metric of similarity, such as a perceptual hash metric, based on the representation networks 20 and the decision network 22.

[0062] For consistency, the application of the machine-learned model is more likely to provide a vector for rendered images (e.g., photorealistic image) of the patient volume that may be used to compare to a vector for a standard or reference image despite differences in the volume, medical scanner, medical scanner configuration, or other differences. The machine-learned model is trained to output representation vectors of fixed length, which may be used for comparison. The machine-learned model (i.e., representation network) is used to provide consistency over time for a patient, in a practice, for an application, and/or across patients.

[0063] In act 36, the image processor selects one of the rendered images based on the representation vectors from application of the machine-learned model. The most similar one of the rendered images of the patient volume to the reference image is found through repetition of acts 32 and 34. A sufficiency threshold may be used, such as searching for an image that has a similarity above a threshold. The repetition continues until a rendered image that is perceptually similar to the reference image is found. By comparing representation vectors from multiple images to the representation vector of the reference image, an image more similar than any of the other images may be identified.

[0064] To select, the distance between representation vectors is calculated. For example, the Euclidean distance between the two vectors, $D(\theta) = \|f(X(\theta)) - f(R)\|_2$ is

determined. Other distances between representation vectors may be used. The distance is a function of θ . The rendered image with the minimum distance is selected.

[0065] In one embodiment, the rendering of act 32 and the selection of act 36 are performed with a search strategy. The rendering of act 32 and application of act 34 are performed with different combinations of the rendering settings selected as part of an optimization. The graphics processing unit renders new images to mimic the effect of the picked reference by searching for the rendering parameters that minimize the Euclidean distance between the representation of the rendered image and the representation of the picked reference. To find the setting of the rendering parameter that provides the most similar visual effect to R, an optimization problem is solved. θ that minimizes $D(\theta)$ is found. Such optimization problem may be solved using standard optimization techniques, such as hill climbing or gradient descent. Which rendering parameters to change and/or how much to change the setting is based on optimizing. Based on the distances between vectors of different images from the reference image, different magnitudes of change and/or selections of rendering parameters to change are tested by rendering and performing further comparison. The results are used to inform the magnitude and/or selection of rendering parameter for the next rendering.

[0066] In act 38, the selected rendered image is displayed. Other images of the plurality of rendered images, the reference image, or other images may be displayed. The rendering parameter settings for the selected image may be used for rendering subsequent images, such as from different camera positions or other rendering interactions from the user.

[0067] The rendered image of the patient is displayed. This displayed image is more consistent in appearance with the reference image due to the application of the machine-learned comparison model. By displaying the selected image, a rendered image more consistent with the reference is provided. As a result, the image may more likely include diagnostically recognizable information, may be more easily analyzed, and/or may be used for comparison with expected or past results. The time spent reviewing may be reduced due to consistency. The ability or likelihood to recognize diagnostically useful information may be increased.

[0068] The renderer, host computer of the renderer, or other device transmits the photorealistic image or images to a display device. The transmission is over a bus, through a network interface, over a network, from a memory, or from a display plane. For example, the image is loaded into a buffer and output from the buffer to a display for viewing by a physician to aid diagnosis or pre-operative planning. A computer-generated representation of the volume of the patient suitable for viewing by human users (e.g., 2D images, videos, stereoscopic image pairs, volumes, holograms, or lightfields) is output. Alternatively, the output may be for storage and/or for further machine analysis.

[0069] In another embodiment, the display of the image is an initial image for interactive viewing. The machine-learned model is used to initialize interactive viewing applications in the clinical workflow. By providing values for rendering parameters so that the images are consistent with an established visualization design, the amount of user interaction may be reduced.

[0070] Figure 4 is a flow chart diagram of one embodiment of a method for three-dimensional rendering in a rendering system. Due to the time needed to render (e.g., seconds or minutes) using physically-based rendering (e.g., global illumination volume rendering using Monte Carlo sampling), interactive performance of the rendering is limited. While this computationally expensive rendering may produce high quality images with improved perception of depth cues, the rendering uses a high number of samples per pixel to compute the rendering equation. The method of Figure 4 uses a self-optimizing random number sequence for the path tracing, allowing better approximation of the final image with a smaller number of samples. Path tracing randomly samples light paths between the virtual camera and the light sources, through the volumetric data or surfaces. The computation uses a sequence of random numbers to compute light scattering for a single light path. An optimized sequence may allow the initial images produced by the renderer, using only a small number of light path computations, to better approximate the final result of the light transport simulation using a greater number of light path computations. To compute the light paths in an optimized order, the sequence of light scattering direction is controlled indirectly by manipulating the seed used for the pseudo-random number generator.

[0071] The method is performed by the system of Figure 8 or a different system. For example, a renderer performs acts 40, 42, and 52, an image processor or controller performs acts 44-48, and a medical scanner performs act 50. Other devices may perform or be used in performance of any of the acts.

[0072] The acts are performed in the order shown or other orders. For example, acts 40 and 42 are performed in any order. Additional, different, or fewer acts may be performed. For example, acts 50 and/or 52 are not performed.

[0073] In act 40, a renderer three-dimensionally renders a reference image. The reference image is rendered with path tracing using a first number of samples per pixel and a first random number generator seed. Any seed may be used, such as a default seed, a seed based on internal operation of a processor (e.g., time), or other seed. Any number of samples per pixel (i.e., number of light-paths traced in the volume per pixel) may be used, such as hundreds or thousands. The rendering provides a reference or desired image. For example, Figure 5D shows a physically-based rendering using 2000 samples per pixel for a given seed.

[0074] In one embodiment, a user selects and loads a representative volume (e.g., scan data) and rendering

configuration that produces images with a variety of light paths. The volume used for the seed optimization may be an example of real data for which the best rendering is to be provided (e.g., vascular CT volume) or a special volume used to find a seed that produces improved results over a variety of data. The types, sizes, and/or numbers of surfaces, textures, boundaries and/or other characteristics represented in the volume may be considered in selecting the volume to create the reference image. For example, the selected volume contains structure for a variety of light interactions (e.g., interactions with convex surfaces, concave surfaces, semi-transparent regions, clip planes, areas in light and shadow, etc.).

[0075] The reference image is rendered to be application specific (e.g., for a particular organ, disease, or test) or may be generic to two or more applications. A full resolution reference image is generated with a large number of samples per pixel.

[0076] In act 42, the renderer three-dimensionally renders other images with the path tracing using a lower number of samples per pixel and other random number generator seeds. Any different candidate seeds may be used, such as random or pseudo-random selection of seeds. The seeds may be from a list, such as sequencing through the list. Any selection of the different seeds may be used, such as using time as each rendering is performed. Since the relationship between the values of the seed and the resulting image rendered using the seed is not direct, any selection may be used.

[0077] Any number of different seeds and corresponding different rendered images may be used, such as five or more, tens, hundreds, or thousands. The identification of an acceptable seed may be performed as a pre-process or before use of the seed for a given patient, so there may be any amount of time and/or processing power available. The settings or values of other rendering parameters are kept constant for the different renderings, but may also vary.

[0078] The number of samples per pixel is limited, such as being less than the number of samples per pixel used in the reference by a factor of two, four, ten, or more. For example, the reference uses 2000 samples per pixel and the renderings with other seeds use 200 or less (e.g., 20 or less). The rendered images of Figures 5A-C use 16 samples per pixel compared to the 2000 samples per pixel of Figure 5D. A small number of samples is used since a small number results in more rapid rendering time, allowing for more responsive interactive rendering. An image with sufficient information to allow a user to decide whether to change the rendering (e.g., camera position, transfer function, window, or material property) is rendered more quickly using fewer samples per pixel.

[0079] The rendering of acts 40 and 42 uses or emulates coherent scattering in the path tracing. Coherent scattering improves the memory caching efficiency, leading to significant performance improvement at runtime (e.g., 30% to 100+% improvement in rendering speed, depending on data classification, GPU architecture and

other factors). Coherent scattering also lowers the perceived spatial variance in the image at the cost of slower convergence speed and temporal artifacts during image refinement. Coherent scattering may be used during interactive rendering with a small number of samples per pixel, where sub-optimal light paths may lead to larger image differences compared to the non-coherent scattering.

[0080] In act 44 of Figure 4, an image processor measures an error between each of the rendered images and the reference image. Any metric may be used, such as a measure of similarity or correlation. In one embodiment, a perceptual metric is used, such as a perceptual hash. The error between the reference image and an image rendering with a given, fewer number of samples per pixel than the reference image is calculated.

[0081] In an alternative to fixing the number of samples, the error measurement of act 44 may be used with the rendering of act 42 to measure the number of samples per pixel when the error reaches a threshold. The rendering continues by adding more samples. For each iteration (e.g., every 10 or 100 additional samples per pixel), the measurement of act 44 is performed again. This continues until the error between the rendered image and the reference image is below a threshold. The RNG seed that produces faster convergence toward the reference image (i.e., fewer number of samples per pixel to reach the threshold error), rather than optimizing the interactive image quality, is found. For each candidate seed, any number of light paths are used in rendering based on the number of samples required to lower the error metric to a given threshold.

[0082] In act 46, the image processor identifies one of the RNG seeds. The RNG seed that results in a rendered image with a lesser error than the other seeds is selected. Alternatively, the seed that results in a sufficient error with the lowest number of samples per pixel is selected.

[0083] By repeating the rendering of act 42 and the measurement of error of act 44, the seed resulting in a rendering with a fewer number of samples is found. The repetition continues until a seed providing an error below a threshold is found. Alternatively, the repetition continues until a given number, N, of candidates are evaluated, and the seeds with the lowest metric or number of samples per pixel meeting a metric is selected.

[0084] In act 48, a rendering system is configured for path tracing with the identified seed. The image processor may configure the renderer. Alternatively, the seed is passed to a different system. A renderer of the different system is configured by loading or using the selected seed. The identification is performed as pre-processing. The identified seed is then passed to any number of other rendering systems for use in rendering volumes for patients.

[0085] Different seeds may be identified for different applications. The different seeds are used to configure the rendering system for rendering in respective different applications. Alternatively, a seed is identified generically

and used in any application.

[0086] The seed identified in act 46 allows the renderer to produce an image close to the reference image with fewer samples per pixel, so more rapidly. The optimized random number generator (RNG) sequence resulting from the selected seed produces improved results for interactive rendering with a low number of samples as compared to the rendering of the reference image.

[0087] Figures 5A-D illustrate the effect of the optimized RNG seed on the image quality during interactive rendering. Figures 5A-C are rendered with only 16 samples per pixel. In contrast, the result in Figure 5D is rendered with 2000 samples per pixel. Figure 5A illustrates a rendering with randomized scattering, resulting in significant spatial noise. Figure 5B illustrates a rendering with coherent scattering with a non-optimal seed (i.e., using a seed chosen without the Figure 4 method). In Figure 5B, the color differs significantly from the reference image since the most important light paths have not been sampled yet. In contrast, Figure 5C approximates the reference image more closely with reduced spatial noise at 16 samples per pixel by leveraging both coherent scattering and optimized light path evaluation. Coherent scattering is an independent feature, which improves the speed of rendering. The optimized seed allows further improvement using a small number of light-paths. Unlike in non-coherent sampling, light paths for nearby pixels are traced in such a way that they follow similar paths through the volume in coherent scattering. However, this means that entire images worth of light paths may all miss the important directions in the initial samples. By using the optimized seed, the initial samples used in coherent scattering are less likely to miss the important directions. With sufficient sampling, all three configurations reach the reference image. The optimization allows an image with many fewer samples per pixel to be closer to the reference image.

[0088] Figures 6A and 6B illustrate the initial 16 light paths computed per pixel. Figure 6A uses the non-optimized seed of Figure 5B. Figure 6B uses the optimized seed of Figure 5C. As shown in Figure 6B, a more uniform sampling of the light directions without explicit prior knowledge of the data or the lighting environment is provided due to seed selection.

[0089] Figure 7 illustrates the error progression over the number of samples per pixel. The error (e.g., perceptual hash difference from a reference image) decreases as more samples are integrated into the final image. The dashed line is for the unoptimized seed used for Figure 5B. The solid line is for the optimized seed used for Figure 5C. While in both cases, the image is visually indistinguishable from the reference image at above 2000 samples per pixel, the initial refinement which the user sees during interactive rendering is heavily dependent on the RNG stream. The seed selected in act 46 (i.e., optimized seed) provides a physically-based rendered image more similar to the reference image after a much fewer number of samples per pixel. Even after a few samples per pixel

(e.g., 5), the image from the optimized seed has a much smaller (e.g., perceptual hash error of 5 against the reference image) than the unoptimized seed (e.g., perceptual hash error of 40 against the reference image).

[0090] The more rapid rendering allows the user to interact and change rendering more quickly. Since fewer samples per pixel are needed to have an image closer to the final image, the user may more responsively adjust. Rather than waiting for the completion of the rendering at many samples per pixel, the changes may be made earlier in the progressive rendering used in path tracing. The effect allows coherent scattering to be used during interaction as well, leading to performance improvement and lowered spatial noise.

[0091] The optimized seed may be used in any of various situations. In an embodiment shown in Figure 4, the seed is used for rendering for a given patient. In act 50, the patient is scanned. The previously selected seed is used to configure the rendering system in act 48. The rendering system then three-dimensionally renders from the scan data of the patient in act 52. The image of the patient is rendered with path tracing using a fewer number or lesser number of samples per pixel than the reference image used for optimizing. An initial or progressively rendered image more closely appearing as the reference is provided. The progressive rendering may continue by adding additional samples per pixel or may be ceased to save time and processing power. The earlier, good quality rendered image allows for more responsive interactive rendering.

[0092] In another embodiment, acts 42-48 are performed as part of performing acts 50 and 52. The optimization of the seed occurs as part of rendering a sequence of images from a scan of the patient. As the user interacts, different renderings are performed. For example, as the user rotates the camera position relative to the volume, multiple (e.g., four or ten) renderings may be performed. Each rendering uses a different seed. The error is measured from the reference image to identify any seed providing for similar appearance after fewer number of samples per pixel. If or once such a seed is identified, that seed may then be used for rendering other images in the sequence for that patient.

[0093] In yet another embodiment, the seed is used to render for training a neural network or other machine learning. Rather than performing renderings using many samples per pixel, time and processing may be saved by using the optimized seed and limiting the number of samples per pixel. Many rendered images may be generated more rapidly due to the seed optimization. The machine training (e.g., deep learning) uses images generated with the identified random number generator seed. The seed may be used in other rendering situations. For example, the renderings of act 32 of Figure 2 use the optimized seed.

[0094] Figure 8 shows a block diagram of one embodiment of a system for three-dimensional rendering in medical imaging. The system implements the method of

Figures 1, 2, and/or 4. The system is for training with machine learning (e.g., Figure 1) and/or application of a machine-learned model (e.g., Figure 2). The machine learning is used for comparison of images with a reference image to identify an image and/or rendering settings for consistent rendering. Alternatively or additionally, the system is for identifying and/or applying a seed used for pseudo-random number generation in path tracing (e.g., Figure 4). Other methods or acts may be implemented, such as providing a user input (e.g., mouse, trackball, touch pad, and/or keyboard) and user interface for interactive rendering.

[0095] The system includes a GPU 80, a memory 82, a display 84, a medical scanner 86, and an image processor 88. Additional, different, or fewer components may be provided. For example, the medical scanner 86 and/or memory 82 are not provided. In another example, a network or network connection is provided, such as for networking with a medical imaging network or data archival system. A user interface may be provided for interacting with the GPU 80, image processor 88, or other components.

[0096] The GPU 80, memory 82, image processor 88, and/or display 84 are part of the medical scanner 86. Alternatively, the GPU 80, memory 82, image processor 88, and/or display 84 are part of a server, workstation, or computer separate from the medical scanner 86. In other embodiments, the image processor 88, memory 82, and/or GPU 80 are part of a remote server for interacting with the medical scanner 86, which includes the remaining components. The GPU 80, memory 82, image processor 88, and/or display 84 may be a personal computer, such as desktop or laptop, a workstation, a server, or combinations thereof. In yet other embodiments, the GPU 80 and memory 82 are part of a separate computer from the image processor 88.

[0097] The medical scanner 86 is a medical diagnostic imaging system. Ultrasound, CT, x-ray, fluoroscopy, positron emission tomography (PET), single photon emission computed tomography (SPECT), and/or MR systems may be used. The medical scanner 86 may include a transmitter and includes a detector for scanning or receiving data representative of the interior of the patient. Scan data is acquired and used for diagnosis or surgical planning, such as identifying a lesion or treatment location.

[0098] The medical scanner 86 acquires scan data representing the patient. The medical scanner 86 is configured by settings, such as a preset and/or user controlled settings, to scan the patient. The scan is a volume scan, such as scanning a three-dimensionally distributed region. The scan data may represent a volume of the patient. The medical scanner provides scan data representing locations distributed in three dimensions.

[0099] In alternative embodiments, the medical scanner 86 is not provided, but previously acquired scan data for a patient is stored in the memory 82. In yet other alternatives, many medical images are provided in the

memory 82 as the training data, which is gathered and stored in the memory 82 from the medical scanner 86 or other sources.

[0100] The GPU 80 is a renderer, such as a graphics card, graphic chip, multi-core processor, or other processor for running a software package that implements the photorealistic image generation functionality. The GPU 80 is configured by an application programming interface to render one or more images from the 3D scan data representing a patient. Using physically-based rendering, a photorealistic image is rendered. Path tracing based on a pseudo-random number sequence generated from a seed is used to model light scattering. Using the seed and other rendering parameter settings, path tracing is applied to render the image. Other types of rendering may be used. For machine learning, the GPU 80 may render training images.

[0101] In one embodiment, the GPU 80 is configured to three-dimensionally render images from the scan of the volume of a patient or multiple patients. The images are rendered with different values for rendering parameters with a same seed (e.g., optimization selected seed). In other embodiments, the same values for the rendering parameters are used, but different RNG seeds are used in the rendering using a limited number of samples. In yet other embodiments, both the seeds and one or more other rendering parameters are changed for different images.

[0102] The image processor 88 is a computer, workstation, server, processor, or other device configured to apply machine learning, to apply a machine-learned model, to select an image using optimization, and/or to optimize a RNG seed. The image processor 88 is configured by software, hardware, and/or firmware. For learning, the image processor 88 is configured by one or more machine learning algorithms. For applying a learnt model, the image processor 88 is configured, in part, by a learnt matrix or matrices associating input data to output values.

[0103] The image processor 88 applies a machine-learned model, such as one learnt with deep learning. The image processor 88 is configured to search for one of the images matching a reference image based on a machine-learned neural network, such as a representation network. The machine-learned model, as implemented by the image processor 88, generates representation vectors from a reference image. The image processor 88 is configured to select an image and/or rendering parameters sufficiently similar (e.g., threshold) to the reference or best matching from a distance between representation vectors for the images and the reference. Optimization by the image processor 88 may be used to control the rendering settings used.

[0104] In another embodiment, the image processor 88 measures an error from renderings with different seeds to a reference image. An optimized seed is determined from the error, such as based on the number of samples providing sufficiently similar to a reference (e.g., threshold applied to the error) or the seed providing a

best match (e.g., least error).

[0105] The memory 82 is a graphics processing memory, a video random access memory, a random-access memory, system memory, cache memory, hard drive, optical media, magnetic media, flash drive, buffer, database, combinations thereof, or other now known or later developed memory device for storing training data, rendering parameters values, one or more seeds, medical datasets, scan data, rendered images, error values, scalar values, and/or other information. The memory 82 is part of the medical scanner 86, part of a computer associated with the GPU 80, part of a computer associated with the image processor 88, a database, part of another system, a picture archival memory, or a standalone device.

[0106] The memory 82 stores scan data representing one or more patients. For example, data from the medical scanner 86 is stored. The data is in a scan format or reconstructed to a volume or three-dimensional grid format. The scan data is stored for training and/or stored for rendering an image or images of a patient.

[0107] The memory 82 or other memory is alternatively or additionally a non-transitory computer readable storage medium storing data representing instructions executable by the programmed GPU 80 and/or image processor 88 for learning or applying the machine-learned model, for selecting a similar image, and/or for identifying a seed. The instructions for implementing the processes, methods, and/or techniques discussed herein are provided on non-transitory computer-readable storage media or memories, such as a cache, buffer, RAM, removable media, hard drive, or other computer readable storage media. Non-transitory computer readable storage media include various types of volatile and nonvolatile storage media. The functions, acts or tasks illustrated in the figures or described herein are executed in response to one or more sets of instructions stored in or on computer readable storage media. The functions, acts or tasks are independent of the particular type of instructions set, storage media, processor or processing strategy and may be performed by software, hardware, integrated circuits, firmware, micro code and the like, operating alone, or in combination. Likewise, processing strategies may include multiprocessing, multitasking, parallel processing, and the like.

[0108] In one embodiment, the instructions are stored on a removable media device for reading by local or remote systems. In other embodiments, the instructions are stored in a remote location for transfer through a computer network or over telephone lines. In yet other embodiments, the instructions are stored within a given computer, CPU, GPU, or system.

[0109] The display 84 is a monitor, LCD, projector, plasma display, CRT, printer, or other now known or later developed device for displaying the rendered image or images (e.g., photorealistic image or images). The displayed image or images may be of a volume of a patient. The display 84 receives images from the GPU 80, mem-

ory 82, image processor 88, or medical scanner 86. The images of the tissue captured by the medical scanner 86 are displayed. Other information may be displayed as well, such as generated graphics, text, or quantities as a virtual overlay.

[0110] Additional images may be displayed. For example, an image is rendered with physically-based rendering. An initial image is displayed. As further samples per pixel are performed in path tracing, the image is updated (e.g., one or more pixel colors are replaced with new colors or the entire image is replaced).

[0111] While the invention has been described above by reference to various embodiments, it should be understood that many changes and modifications can be made without departing from the scope of the invention. It is therefore intended that the foregoing detailed description be regarded as illustrative rather than limiting, and that it be understood that it is the following claims, including all equivalents, that are intended to define the spirit and scope of this invention.

Claims

1. A method for three-dimensional rendering in a rendering system, the method comprising:
 - acquiring (30), by a medical imaging system, a medical dataset representing a three-dimensional region of a patient;
 - rendering (32), by a renderer using different combinations of rendering (32) settings, a plurality of images from the medical dataset;
 - applying (34), by an image processor (88), different pairs of the images of the plurality and a reference image to a machine-learned model, the machine-learned model trained to generate representation vectors of the images and the reference image;
 - selecting (36), by the image processor (88), one of the plurality of images based on the representation vectors; and
 - displaying (38) the selected one of the plurality of images.
2. The method according to claim 1, wherein acquiring (30) comprises acquiring (30) from a magnetic resonance, computed tomography, ultrasound, single photon emission computed tomography, or positron emission tomography system as the medical imaging system.
3. The method according to claims 1 or 2, wherein rendering (32) comprises rendering (32) with path tracing emulating photon interaction with the medical dataset.
4. The method according to claim 3, wherein rendering

- (32) comprises rendering (32) with Monte Carlo path tracing using a random number generator seed identified (46) based on an optimization.
5. The method according to claims 1 or 2, wherein rendering (32) comprises rendering (32) with the rendering (32) settings comprising transfer function, window function, shading, and clip plane position, and wherein a first subset of the rendering (32) settings is held constant in the rendering (32) of the plurality of images and a second subset of the rendering (32) settings are varied. 5
 6. The method according to any of the preceding claims, wherein applying (34) comprises applying (34) with the machine-learned model comprising a machine-learned neural network 10
or
wherein applying (34) comprises applying (34) with the machine-learned neural network comprising a convolutional neural network taught to output a feature vector for each of the images of the pair and the reference image 20
or
wherein applying (34) comprises applying (34) with the machine-learned model trained with ground truth from a perceptual hash metric 25
or
wherein applying (34) comprises applying (34) with the machine-learned model trained with random variation of the rendering (32) settings and camera positions. 30
 7. The method according to any of the preceding claims, wherein selecting (36) comprises selecting (36) based on Euclidian distances of the representation vectors 35
and/or
wherein selecting (36) comprises selecting (36) the one as most similar to the reference image, and wherein displaying (38) comprises displaying (38) the selected one as an image of the patient on a display 40
and/or
wherein selecting (36) comprises searching with optimization applied to the different combinations of the rendering (32) settings through the rendering (32) and applying (34). 45
 8. A system for three-dimensional rendering in medical imaging, the system comprising: 50

a medical scanner (86) configured to scan a volume of a patient;
a graphics processing unit (80) configured to three-dimensionally render images from the scan of the volume, the images rendered with different values for rendering parameters; 55
 - an image processor (88) configured to search for one of the images matching a reference image based on a machine-learned neural network; and
a display (84) configured to display the one image of the volume of the patient.
 9. The system according to claim 8, wherein the graphics processing unit (80) is configured to three-dimensionally render the images with path tracing or wherein the graphics processing unit (80) is configured to render with a random number generator seed, the random number generator seed having been selected in an optimization.
 10. The system according to claims 8 or 9, wherein the image processor (88) is configured to search with the machine-learned neural network including a convolutional neural network for the reference image and the images and configured to determine distances between representation vectors output by the convolutional neural network.
 11. A method for three-dimensional rendering in a rendering system, the method comprising:

three-dimensionally rendering (40) a reference image with path tracing using a first number of samples per pixel and a first random number generator seed;
three-dimensionally rendering (42) other images with the path tracing using a second number or less of samples per pixel and other random number generator seeds, the second number less than the first by a factor of ten or more;
measuring (44) an error between the other images and the reference image;
identifying (46) one of the other random number generator seeds with a lessor error than the others of the random number generator seeds; and
configuring (48) the rendering system for path tracing with the identified one of the random number generator seeds.
 12. The method according to claim 11, further comprising scanning (50) a patient with a medical imaging system, and three-dimensionally rendering (52), by the rendering system as configured, an image from output of the scanning with path tracing using the second number or less of the samples per pixel and the identified one of the random number generator seeds.
 13. The method according to claim 12, further comprising performing the measuring (44) and identifying (46) as part of interactive rendering (42) from the output of the scanning.

14. The method according to any of the claims 11 to 13, wherein measuring (44) the error comprises measuring (44) the number of samples per pixel where the error reaches a threshold. 5
15. The method according to any of the claims 11 to 14, wherein three-dimensionally rendering (40, 42) other images comprises emulating coherent scattering. 10
16. The method according to any of the claim 11 to 15, further comprising training (14) a neural network with deep learning to identify similar images with an architecture using a representation network connected to a decision network, the deep learning using training renderings generated with the identified one of the random number generator seeds. 15

20

25

30

35

40

45

50

55

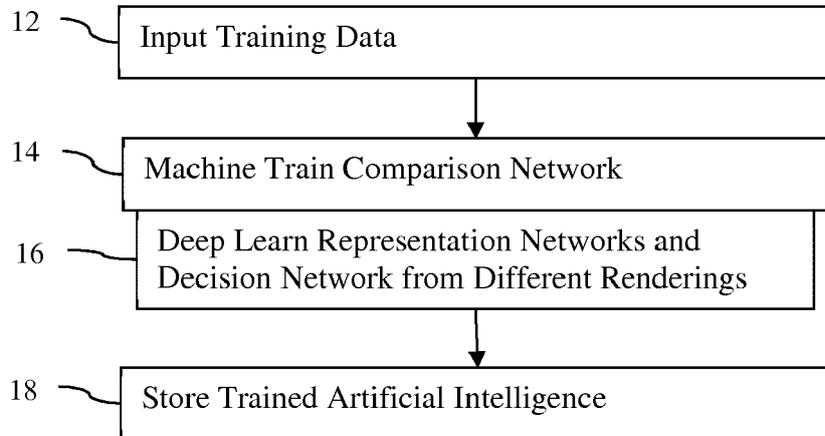


FIG. 1

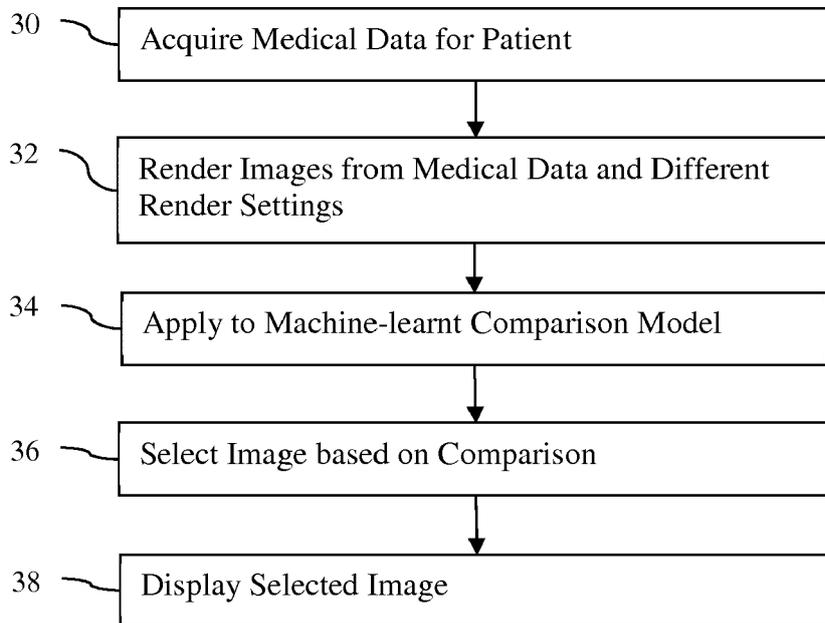


FIG. 2

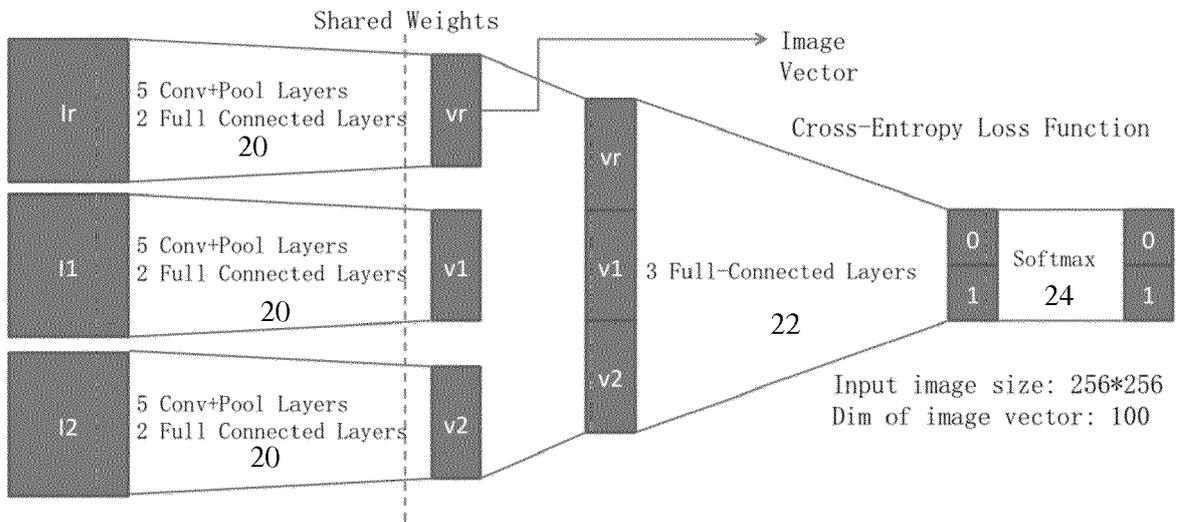


FIG. 3

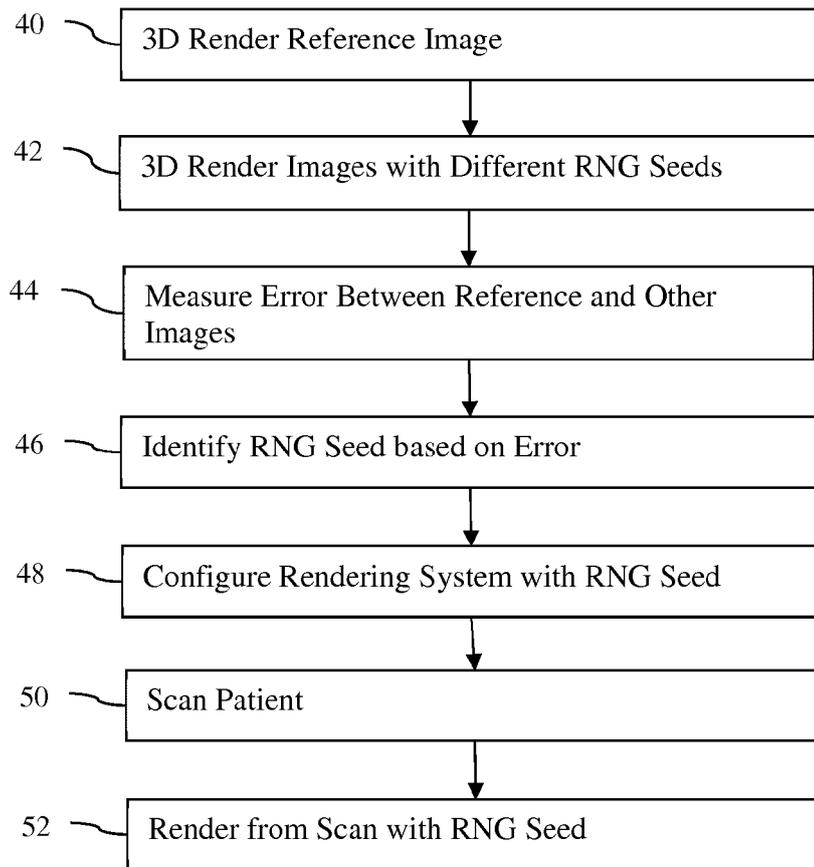


FIG. 4

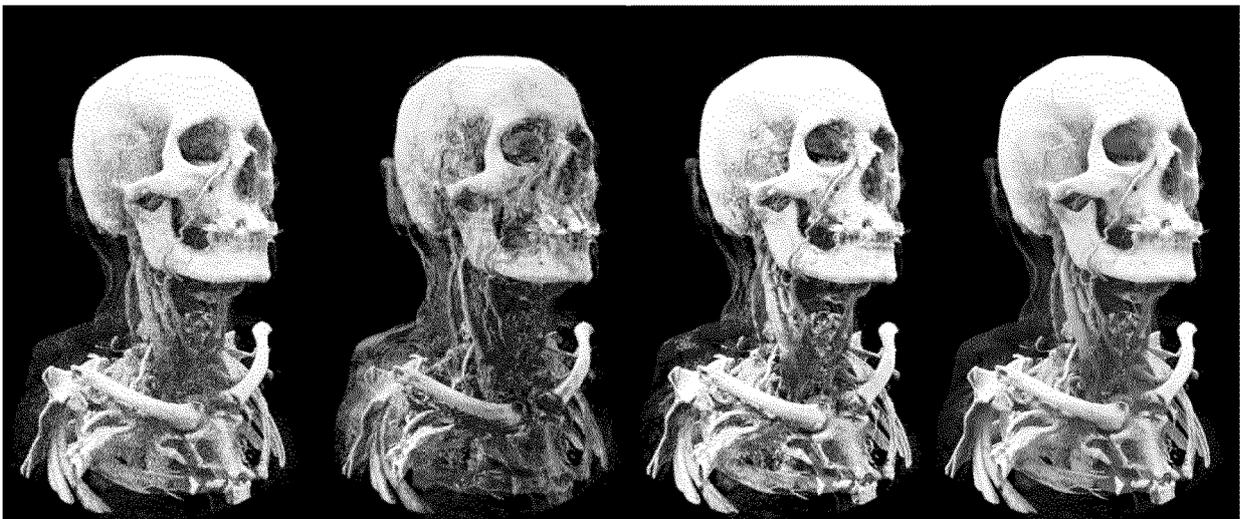


FIG. 5A

FIG. 5B

FIG. 5C

FIG. 5D

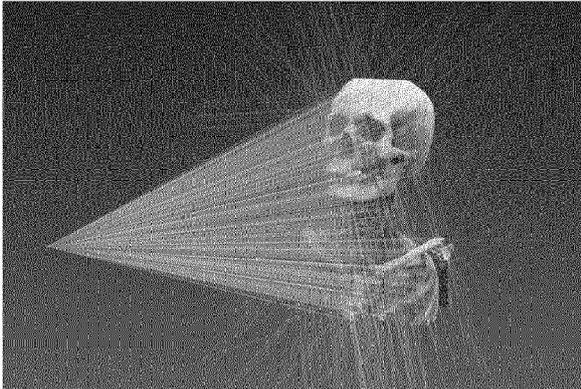


FIG. 6A

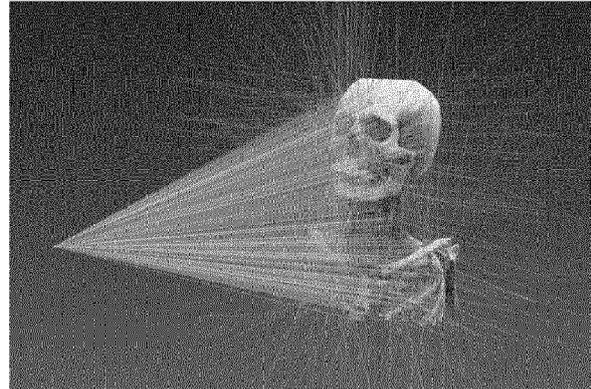


FIG. 6B

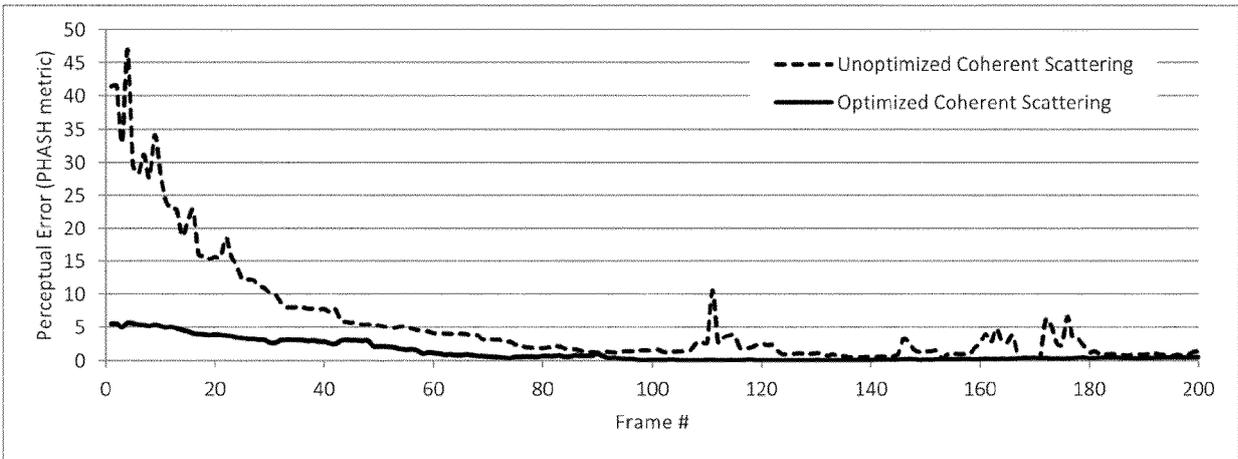


FIG. 7

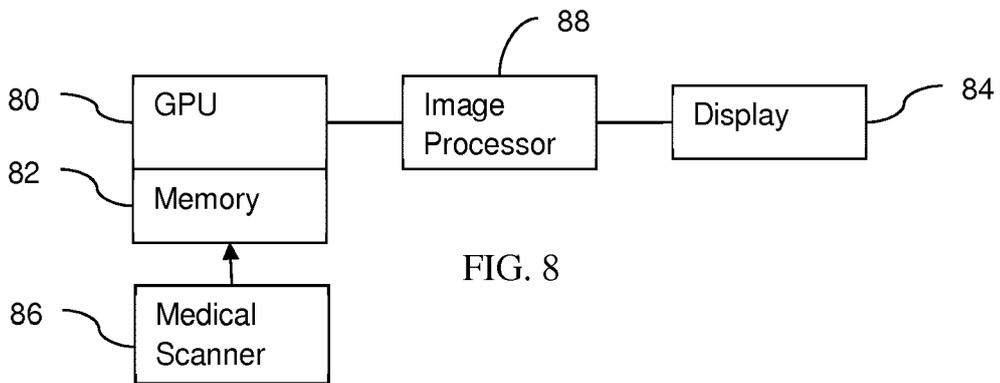


FIG. 8