As illustrated, a generator network 202 may take, as an input, a random input 204. Generator network 202 may then generate, as an output, a synthetic (i.e., fake) sample 206. Referring now to discriminator network 212, a real sample 210 may be selected from a set of real samples (e.g., real images 208). Discriminator network 212 may be configured to take, as inputs, real sample 210 and synthetic sample 206 and generate an output indicating which of the two inputs is the real sample. An error associated with the output of discriminator network 212 may then be used to determine discriminator loss 214 and generator loss 216. For example generator loss 216 may be relatively higher in instances in which discriminator network 212 correctly identifies real sample 210 as the real sample, effectively penalizing generator network 202 for not generating a realistic enough synthetic sample. In contrast, discriminator loss 214 may be relatively higher in instances in which discriminator network 212 incorrectly identifies the real sample, effectively penalizing discriminator network 212 for not correctly discriminating between real and synthetic inputs. By concurrently updating generator network 202 and discriminator network 212, generator network 202 may be trained to generate highly realistic synthetic samples. Note that, once training of the GAN has been completed, the trained GAN may include a trained generator network that is now (e.g., after training) configured to generate realistic synthetic samples.

[0042] In some implementations, a GAN (e.g., a GAN used to generate synthetic segmented images and/or a GAN used to generate virtual images from a synthetic segmented image) may have a pyramidal structure comprising a set of layers. Using such a pyramidal structure, the GAN may be trained using as few as a single real image. Each layer may include a generator network and a discriminator network. Each layer may operate on a patch of the real image used to train the GAN such that lower level layers operate on patches having lower (e.g., coarser) resolution, and higher level layers operate on patches having higher (e.g., finer) resolution. Accordingly, by training all of the layers of the GAN, the GAN may, in aggregate, learn to synthetically generate both coarse and fine features in synthetic images. In some embodiments, for a given layer of the GAN, the generator network may receive both a random seed image (e.g., noise) having the size corresponding to the patch that layer of the GAN operates on, as well as the synthetically generated image generated by the generator network of the preceding (e.g., lower level) GAN. For a given layer, the discriminator network may receive the synthetically generated image of the generator network from the same layer. In some implementations, generator networks and discriminator networks may be trained sequentially for the different layers for a given image, such that lower level layer networks (corresponding to lower, or coarser resolution) are trained before higher level layer networks

[0071] The computer-storage media may include both volatile and nonvolatile media and/or removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules, and/or other data types. For example, memory 704 may store computer-readable instructions (e.g., that represent a program(s) and/or a program element(s), such as an operating system. Computer-storage media may include, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computing device 700. As used herein, computer storage media does not comprise signals per se.

[0072] The communication media may embody computer-readable instructions, data structures, program modules, and/or other data types in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" may refer to a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, the communication media may include wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer-readable media.

[0073] CPU(s) 706 may be configured to execute the computer-readable instructions to control one or more components of the computing device 700 to perform one or more of the methods and/or processes described herein. CPU(s) 706 may each include one or more cores (e.g., one, two, four, eight, twenty-eight, seventy-two, etc.) that are capable of handling a multitude of software threads simultaneously. CPU(s) 706 may include any type of processor and may include different types of processors depending on the type of computing device 700 implemented (e.g., processors with fewer cores for mobile devices and processors with more cores for servers). For example, depending on the type of computing device 700, the processor may be an ARM processor implemented using Reduced Instruction Set Computing (RISC) or an x86 processor implemented using Complex Instruction Set Computing (CISC). Computing device 700 may include one or more CPUs 706 in addition to one or more microprocessors or supplementary co-processors, such as math co-processors.

[0074] GPU(s) 708 may be used by computing device 700 to render graphics (e.g., 3D graphics). GPU(s) 708 may include many (e.g., tens, hundreds, or thousands) of cores that are