



- (51) **International Patent Classification:**
G06F 3/00 (2006.01) *G06F 15/00* (2006.01)
- (21) **International Application Number:**
PCT/US2010/021883
- (22) **International Filing Date:**
22 January 2010 (22.01.2010)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
61/148,884 30 January 2009 (30.01.2009) US
12/391,150 23 February 2009 (23.02.2009) US
- (71) **Applicant (for all designated States except US): MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (72) **Inventors:** **LATTA, Stephen G.**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **TSUNODA, Kudo**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **GEISNER, Kevin**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **MARKOVIC, Relja**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **BENNETT, Darren**; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.


(84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

[Continued on next page]

(54) Title: STANDARD GESTURES

**FIG. 8**

(57) **Abstract:** Systems, methods and computer readable media are disclosed for grouping complementary sets of standard gestures into gesture libraries. The gestures may be complementary in that they are frequently used together in a context or in that their parameters are interrelated. Where a parameter of a gesture is set with a first value, all other parameters of the gesture and of other gestures in the gesture package that depend on the first value may be set with their own value which is determined using the first value.



Published:

- *without international search report and to be republished
upon receipt of that report (Rule 48.2(g))*

STANDARD GESTURES

BACKGROUND OF THE INVENTION

[0001] Many computing applications such as computer games, multimedia applications, office applications or the like use controls to allow users to manipulate game characters or other aspects of an application. Typically such controls are input using, for example, controllers, remotes, keyboards, mice, or the like. Unfortunately, such controls can be difficult to learn, thus creating a barrier between a user and such games and applications. Furthermore, such controls may be different than actual game actions or other application actions for which the controls are used. For example, a game control that causes a game character to swing a baseball bat may not correspond to an actual motion of swinging the baseball bat.

SUMMARY OF THE INVENTION

[0002] Disclosed herein are systems and methods for receiving data reflecting skeletal movement of a user, and determining from that data whether the user has performed one or more gestures. Packages of standard gestures are disclosed from which application developers can incorporate gesture recognition into their applications.

[0003] In an exemplary embodiment, a gesture library comprises a plurality of gestures. Where these gestures are complementary with each other, they may be grouped into gesture packages. These gesture packages are then provided to applications for use by a gesture recognizer engine. An application may utilize one or more gesture packages.

[0004] The application may assign a value to a first parameter of a gesture, such as an arm velocity minimum threshold that must be reached for a football throw gesture to be recognized. The recognizer engine sets the first parameter with the value, and also sets the value of any other parameters of that gesture or any other gestures in the gesture package that are dependent upon the value of the first gesture. For instance, where the gesture package is a sports gesture package that includes the football throw gesture, the package may also include a curveball baseball throw gesture and a fastball baseball throw gesture that are interrelated with the football throw gesture. Where it has been determined that the curveball baseball throw gesture should have an arm velocity minimum threshold of 80% and the fastball baseball throw gesture should have an arm velocity minimum threshold of 90% of the football throw gesture, then those parameters may be set to 80% and 90% of the value, respectively.

[0005] It can be appreciated by one of skill in the art that one or more various aspects of the disclosure may include but are not limited to circuitry and/or programming for effecting the herein-referenced aspects of the present disclosure; the circuitry and/or programming can be virtually any combination of hardware, software, and/or firmware configured to effect the herein-referenced aspects depending upon the design choices of the system designer.

[0006] The foregoing is a summary and thus contains, by necessity, simplifications, generalizations and omissions of detail. Those skilled in the art will appreciate that the summary is illustrative only and is not intended to be in any way limiting.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The systems, methods, and computer readable media for gesture packages of complementary gestures in accordance with this specification are further described with reference to the accompanying drawings in which:

[0008] FIGs. 1A and 1B illustrate an example embodiment of a target recognition, analysis, and tracking system with a user playing a game.

[0009] FIG. 2 illustrates an example embodiment of a capture device that may be used in a target recognition, analysis, and tracking system.

[0010] FIG. 3A illustrates an example embodiment of a computing environment that may be used to interpret one or more gestures in a target recognition, analysis, and tracking system.

[0011] FIG. 3B illustrates another example embodiment of a computing environment that may be used to interpret one or more gestures in a target recognition, analysis, and tracking system.

[0012] FIG. 4A illustrates a skeletal mapping of a user that has been generated from the target recognition, analysis, and tracking system of FIG. 2.

[0013] FIG. 4B illustrates further details of the gesture recognizer architecture shown in FIG. 2.

[0014] FIGs. 5A and 5B illustrate how gesture filters may be stacked to create more complex gesture filters.

[0015] FIGs. 6A, 6B, 6C, 6D and 6E illustrate an example gesture that a user may make to signal for a “fair catch” in football video game.

[0016] FIGs. 7A, 7B, 7C, 7D and 7E illustrate the example “fair catch” gesture of FIGs. 6A-E as each frame of image data has been parsed to produce a skeletal map of the user.

[0017] FIG. 8 illustrates how generic gestures from a gesture library are grouped into genre packages of complementary gestures for a particular task.

[0018] FIG. 9 illustrates exemplary operational procedures for tuning complementary gestures in a gesture package when an application provides a value for one parameter of one gesture.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

[0019] As will be described herein, a user may control an application executing on a computing environment such as a game console, a computer, or the like by performing one or more gestures. According to one embodiment, the gestures may be received by, for example, a capture device. For example, the capture device may capture a depth image of a scene. In one embodiment, the capture device may determine whether one or more targets or objects in the scene corresponds to a human target such as the user. To determine whether a target or object in the scene corresponds a human target, each of the targets may be flood filled and compared to a pattern of a human body model. Each target or object that matches the human body model may then be scanned to generate a skeletal model associated therewith. The skeletal model may then be provided to the computing environment such that the computing environment may track the skeletal model, render an avatar associated with the skeletal model, and may determine which controls to perform in an application executing on the computer environment based on, for example, gestures of the user that have been recognized from the skeletal model. A gesture recognizer engine, the architecture of which is described more fully below, is used to determine when a particular gesture has been made by the user.

[0020] FIGs. 1A and 1B illustrate an example embodiment of a configuration of a target recognition, analysis, and tracking system 10 with a user 18 playing a boxing game. In an example embodiment, the target recognition, analysis, and tracking system 10 may be used to recognize, analyze, and/or track a human target such as the user 18.

[0021] As shown in FIG. 1A, the target recognition, analysis, and tracking system 10 may include a computing environment 12. The computing environment 12 may be a computer, a gaming system or console, or the like. According to an example embodiment, the computing environment 12 may include hardware components and/or software components such that the computing environment 12 may be used to execute applications such as gaming applications, non-gaming applications, or the like.

[0022] As shown in FIG. 1A, the target recognition, analysis, and tracking system 10 may further include a capture device 20. The capture device 20 may be, for example, a

camera that may be used to visually monitor one or more users, such as the user 18, such that gestures performed by the one or more users may be captured, analyzed, and tracked to perform one or more controls or actions within an application, as will be described in more detail below.

5 **[0023]** According to one embodiment, the target recognition, analysis, and tracking system 10 may be connected to an audiovisual device 16 such as a television, a monitor, a high-definition television (HDTV), or the like that may provide game or application visuals and/or audio to a user such as the user 18. For example, the computing environment 12 may include a video adapter such as a graphics card and/or an audio
10 adapter such as a sound card that may provide audiovisual signals associated with the game application, non-game application, or the like. The audiovisual device 16 may receive the audiovisual signals from the computing environment 12 and may then output the game or application visuals and/or audio associated with the audiovisual signals to the user 18. According to one embodiment, the audiovisual device 16 may be connected to
15 the computing environment 12 via, for example, an S-Video cable, a coaxial cable, an HDMI cable, a DVI cable, a VGA cable, or the like.

[0024] As shown in FIGs. 1A and 1B, the target recognition, analysis, and tracking system 10 may be used to recognize, analyze, and/or track a human target such as the user 18. For example, the user 18 may be tracked using the capture device 20 such that the
20 movements of user 18 may be interpreted as controls that may be used to affect the application being executed by computer environment 12. Thus, according to one embodiment, the user 18 may move his or her body to control the application.

[0025] As shown in FIGs. 1A and 1B, in an example embodiment, the application executing on the computing environment 12 may be a boxing game that the user 18 may
25 be playing. For example, the computing environment 12 may use the audiovisual device 16 to provide a visual representation of a boxing opponent 22 to the user 18. The computing environment 12 may also use the audiovisual device 16 to provide a visual representation of a player avatar 24 that the user 18 may control with his or her movements. For example, as shown in FIG. 1B, the user 18 may throw a punch in
30 physical space to cause the player avatar 24 to throw a punch in game space. Thus, according to an example embodiment, the computer environment 12 and the capture device 20 of the target recognition, analysis, and tracking system 10 may be used to recognize and analyze the punch of the user 18 in physical space such that the punch may be interpreted as a game control of the player avatar 24 in game space.

[0026] Other movements by the user 18 may also be interpreted as other controls or actions, such as controls to bob, weave, shuffle, block, jab, or throw a variety of different power punches. Furthermore, some movements may be interpreted as controls that may correspond to actions other than controlling the player avatar 24. For example, the player
5 may use movements to end, pause, or save a game, select a level, view high scores, communicate with a friend, etc.

[0027] In example embodiments, the human target such as the user 18 may have an object. In such embodiments, the user of an electronic game may be holding the object such that the motions of the player and the object may be used to adjust and/or control
10 parameters of the game. For example, the motion of a player holding a racket may be tracked and utilized for controlling an on-screen racket in an electronic sports game. In another example embodiment, the motion of a player holding an object may be tracked and utilized for controlling an on-screen weapon in an electronic combat game.

[0028] According to other example embodiments, the target recognition, analysis, and
15 tracking system 10 may further be used to interpret target movements as operating system and/or application controls that are outside the realm of games. For example, virtually any controllable aspect of an operating system and/or application may be controlled by movements of the target such as the user 18.

[0029] FIG. 2 illustrates an example embodiment of the capture device 20 that may be
20 used in the target recognition, analysis, and tracking system 10. According to an example embodiment, the capture device 20 may be configured to capture video with depth information including a depth image that may include depth values via any suitable technique including, for example, time-of-flight, structured light, stereo image, or the like. According to one embodiment, the capture device 20 may organize the calculated depth
25 information into “Z layers,” or layers that may be perpendicular to a Z axis extending from the depth camera along its line of sight.

[0030] As shown in FIG. 2, the capture device 20 may include an image camera component 22. According to an example embodiment, the image camera component 22 may be a depth camera that may capture the depth image of a scene. The depth image
30 may include a two-dimensional (2-D) pixel area of the captured scene where each pixel in the 2-D pixel area may represent a length in, for example, centimeters, millimeters, or the like of an object in the captured scene from the camera.

[0031] As shown in FIG. 2, according to an example embodiment, the image camera component 22 may include an IR light component 24, a three-dimensional (3-D) camera

26, and an RGB camera 28 that may be used to capture the depth image of a scene. For example, in time-of-flight analysis, the IR light component 24 of the capture device 20 may emit an infrared light onto the scene and may then use sensors (not shown) to detect the backscattered light from the surface of one or more targets and objects in the scene
5 using, for example, the 3-D camera 26 and/or the RGB camera 28. In some embodiments, pulsed infrared light may be used such that the time between an outgoing light pulse and a corresponding incoming light pulse may be measured and used to determine a physical distance from the capture device 20 to a particular location on the targets or objects in the scene. Additionally, in other example embodiments, the phase of the outgoing light wave
10 may be compared to the phase of the incoming light wave to determine a phase shift. The phase shift may then be used to determine a physical distance from the capture device to a particular location on the targets or objects.

[0032] According to another example embodiment, time-of-flight analysis may be used to indirectly determine a physical distance from the capture device 20 to a particular
15 location on the targets or objects by analyzing the intensity of the reflected beam of light over time via various techniques including, for example, shuttered light pulse imaging.

[0033] In another example embodiment, the capture device 20 may use a structured light to capture depth information. In such an analysis, patterned light (i.e., light displayed as a known pattern such as grid pattern or a stripe pattern) may be projected onto the scene via,
20 for example, the IR light component 24. Upon striking the surface of one or more targets or objects in the scene, the pattern may become deformed in response. Such a deformation of the pattern may be captured by, for example, the 3-D camera 26 and/or the RGB camera 28 and may then be analyzed to determine a physical distance from the capture device to a particular location on the targets or objects.

25 **[0034]** According to another embodiment, the capture device 20 may include two or more physically separated cameras that may view a scene from different angles, to obtain visual stereo data that may be resolved to generate depth information

[0035] The capture device 20 may further include a microphone 30. The microphone 30 may include a transducer or sensor that may receive and convert sound into an electrical
30 signal. According to one embodiment, the microphone 30 may be used to reduce feedback between the capture device 20 and the computing environment 12 in the target recognition, analysis, and tracking system 10. Additionally, the microphone 30 may be used to receive audio signals that may also be provided by the user to control applications

such as game applications, non-game applications, or the like that may be executed by the computing environment 12.

[0036] In an example embodiment, the capture device 20 may further include a processor 32 that may be in operative communication with the image camera component 22. The processor 32 may include a standardized processor, a specialized processor, a microprocessor, or the like that may execute instructions that may include instructions for receiving the depth image, determining whether a suitable target may be included in the depth image, converting the suitable target into a skeletal representation or model of the target, or any other suitable instruction.

[0037] The capture device 20 may further include a memory component 34 that may store the instructions that may be executed by the processor 32, images or frames of images captured by the 3-D camera or RGB camera, or any other suitable information, images, or the like. According to an example embodiment, the memory component 34 may include random access memory (RAM), read only memory (ROM), cache, Flash memory, a hard disk, or any other suitable storage component. As shown in FIG. 2, in one embodiment, the memory component 34 may be a separate component in communication with the image capture component 22 and the processor 32. According to another embodiment, the memory component 34 may be integrated into the processor 32 and/or the image capture component 22.

[0038] As shown in FIG. 2, the capture device 20 may be in communication with the computing environment 12 via a communication link 36. The communication link 36 may be a wired connection including, for example, a USB connection, a Firewire connection, an Ethernet cable connection, or the like and/or a wireless connection such as a wireless 802.11b, g, a, or n connection. According to one embodiment, the computing environment 12 may provide a clock to the capture device 20 that may be used to determine when to capture, for example, a scene via the communication link 36.

[0039] Additionally, the capture device 20 may provide the depth information and images captured by, for example, the 3-D camera 26 and/or the RGB camera 28, and a skeletal model that may be generated by the capture device 20 to the computing environment 12 via the communication link 36. The computing environment 12 may then use the skeletal model, depth information, and captured images to, for example, recognize user gestures and in response control an application such as a game or word processor. For example, as shown, in FIG. 2, the computing environment 12 may include a gestures recognizer engine 190. The gestures recognizer engine 190 may include a collection of

gesture filters, each comprising information concerning a gesture that may be performed by the skeletal model (as the user moves). The data captured by the cameras 26, 28 and device 20 in the form of the skeletal model and movements associated with it may be compared to the gesture filters in the gesture recognizer engine 190 to identify when a user (as represented by the skeletal model) has performed one or more gestures. Those gestures may be associated with various controls of an application. Thus, the computing environment 12 may use the gesture recognizer engine 190 to interpret movements of the skeletal model and to control an application based on the movements.

[0040] FIG. 3A illustrates an example embodiment of a computing environment that may be used to interpret one or more gestures in a target recognition, analysis, and tracking system. The computing environment such as the computing environment 12 described above with respect to FIGs. 1A-2 may be a multimedia console 100, such as a gaming console. As shown in FIG. 3A, the multimedia console 100 has a central processing unit (CPU) 101 having a level 1 cache 102, a level 2 cache 104, and a flash ROM (Read Only Memory) 106. The level 1 cache 102 and a level 2 cache 104 temporarily store data and hence reduce the number of memory access cycles, thereby improving processing speed and throughput. The CPU 101 may be provided having more than one core, and thus, additional level 1 and level 2 caches 102 and 104. The flash ROM 106 may store executable code that is loaded during an initial phase of a boot process when the multimedia console 100 is powered ON.

[0041] A graphics processing unit (GPU) 108 and a video encoder/video codec (coder/decoder) 114 form a video processing pipeline for high speed and high resolution graphics processing. Data is carried from the graphics processing unit 108 to the video encoder/video codec 114 via a bus. The video processing pipeline outputs data to an A/V (audio/video) port 140 for transmission to a television or other display. A memory controller 110 is connected to the GPU 108 to facilitate processor access to various types of memory 112, such as, but not limited to, a RAM (Random Access Memory).

[0042] The multimedia console 100 includes an I/O controller 120, a system management controller 122, an audio processing unit 123, a network interface controller 124, a first USB host controller 126, a second USB controller 128 and a front panel I/O subassembly 130 that are preferably implemented on a module 118. The USB controllers 126 and 128 serve as hosts for peripheral controllers 142(1)-142(2), a wireless adapter 148, and an external memory device 146 (e.g., flash memory, external CD/DVD ROM drive, removable media, etc.). The network interface 124 and/or wireless adapter 148

provide access to a network (e.g., the Internet, home network, etc.) and may be any of a wide variety of various wired or wireless adapter components including an Ethernet card, a modem, a Bluetooth module, a cable modem, and the like.

[0043] System memory 143 is provided to store application data that is loaded during the boot process. A media drive 144 is provided and may comprise a DVD/CD drive, hard drive, or other removable media drive, etc. The media drive 144 may be internal or external to the multimedia console 100. Application data may be accessed via the media drive 144 for execution, playback, etc. by the multimedia console 100. The media drive 144 is connected to the I/O controller 120 via a bus, such as a Serial ATA bus or other high speed connection (e.g., IEEE 1394).

[0044] The system management controller 122 provides a variety of service functions related to assuring availability of the multimedia console 100. The audio processing unit 123 and an audio codec 132 form a corresponding audio processing pipeline with high fidelity and stereo processing. Audio data is carried between the audio processing unit 123 and the audio codec 132 via a communication link. The audio processing pipeline outputs data to the A/V port 140 for reproduction by an external audio player or device having audio capabilities.

[0045] The front panel I/O subassembly 130 supports the functionality of the power button 150 and the eject button 152, as well as any LEDs (light emitting diodes) or other indicators exposed on the outer surface of the multimedia console 100. A system power supply module 136 provides power to the components of the multimedia console 100. A fan 138 cools the circuitry within the multimedia console 100.

[0046] The CPU 101, GPU 108, memory controller 110, and various other components within the multimedia console 100 are interconnected via one or more buses, including serial and parallel buses, a memory bus, a peripheral bus, and a processor or local bus using any of a variety of bus architectures. By way of example, such architectures can include a Peripheral Component Interconnects (PCI) bus, PCI-Express bus, etc.

[0047] When the multimedia console 100 is powered ON, application data may be loaded from the system memory 143 into memory 112 and/or caches 102, 104 and executed on the CPU 101. The application may present a graphical user interface that provides a consistent user experience when navigating to different media types available on the multimedia console 100. In operation, applications and/or other media contained within the media drive 144 may be launched or played from the media drive 144 to provide additional functionalities to the multimedia console 100.

[0048] The multimedia console 100 may be operated as a standalone system by simply connecting the system to a television or other display. In this standalone mode, the multimedia console 100 allows one or more users to interact with the system, watch movies, or listen to music. However, with the integration of broadband connectivity made available through the network interface 124 or the wireless adapter 148, the multimedia console 100 may further be operated as a participant in a larger network community.

[0049] When the multimedia console 100 is powered ON, a set amount of hardware resources are reserved for system use by the multimedia console operating system. These resources may include a reservation of memory (e.g., 16MB), CPU and GPU cycles (e.g., 5%), networking bandwidth (e.g., 8 kbs), etc. Because these resources are reserved at system boot time, the reserved resources do not exist from the application's view.

[0050] In particular, the memory reservation preferably is large enough to contain the launch kernel, concurrent system applications and drivers. The CPU reservation is preferably constant such that if the reserved CPU usage is not used by the system applications, an idle thread will consume any unused cycles.

[0051] With regard to the GPU reservation, lightweight messages generated by the system applications (e.g., popups) are displayed by using a GPU interrupt to schedule code to render popup into an overlay. The amount of memory required for an overlay depends on the overlay area size and the overlay preferably scales with screen resolution. Where a full user interface is used by the concurrent system application, it is preferable to use a resolution independent of application resolution. A scaler may be used to set this resolution such that the need to change frequency and cause a TV resynch is eliminated.

[0052] After the multimedia console 100 boots and system resources are reserved, concurrent system applications execute to provide system functionalities. The system functionalities are encapsulated in a set of system applications that execute within the reserved system resources described above. The operating system kernel identifies threads that are system application threads versus gaming application threads. The system applications are preferably scheduled to run on the CPU 101 at predetermined times and intervals in order to provide a consistent system resource view to the application. The scheduling is to minimize cache disruption for the gaming application running on the console.

[0053] When a concurrent system application requires audio, audio processing is scheduled asynchronously to the gaming application due to time sensitivity. A multimedia

console application manager (described below) controls the gaming application audio level (e.g., mute, attenuate) when system applications are active.

[0054] Input devices (e.g., controllers 142(1) and 142(2)) are shared by gaming applications and system applications. The input devices are not reserved resources, but are to be switched between system applications and the gaming application such that each will have a focus of the device. The application manager preferably controls the switching of input stream, without knowledge the gaming application's knowledge and a driver maintains state information regarding focus switches. The cameras 26, 28 and capture device 20 may define additional input devices for the console 100.

[0055] FIG. 3B illustrates another example embodiment of a computing environment 220 that may be the computing environment 12 shown in FIGs. 1A-2 used to interpret one or more gestures in a target recognition, analysis, and tracking system. The computing system environment 220 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the presently disclosed subject matter. Neither should the computing environment 220 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary operating environment 220. In some embodiments the various depicted computing elements may include circuitry configured to instantiate specific aspects of the present disclosure. For example, the term circuitry used in the disclosure can include specialized hardware components configured to perform function(s) by firmware or switches. In other examples embodiments the term circuitry can include a general purpose processing unit, memory, etc., configured by software instructions that embody logic operable to perform function(s). In example embodiments where circuitry includes a combination of hardware and software, an implementer may write source code embodying logic and the source code can be compiled into machine readable code that can be processed by the general purpose processing unit. Since one skilled in the art can appreciate that the state of the art has evolved to a point where there is little difference between hardware, software, or a combination of hardware/software, the selection of hardware versus software to effectuate specific functions is a design choice left to an implementer. More specifically, one of skill in the art can appreciate that a software process can be transformed into an equivalent hardware structure, and a hardware structure can itself be transformed into an equivalent software process. Thus, the selection of a hardware implementation versus a software implementation is one of design choice and left to the implementer.

[0056] In FIG. 3B, the computing environment 220 comprises a computer 241, which typically includes a variety of computer readable media. Computer readable media can be any available media that can be accessed by computer 241 and includes both volatile and nonvolatile media, removable and non-removable media. The system memory 222

5 includes computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) 223 and random access memory (RAM) 260. A basic input/output system 224 (BIOS), containing the basic routines that help to transfer information between elements within computer 241, such as during start-up, is typically stored in ROM 223. RAM 260 typically contains data and/or program modules that are
10 immediately accessible to and/or presently being operated on by processing unit 259. By way of example, and not limitation, FIG. 3B illustrates operating system 225, application programs 226, other program modules 227, and program data 228.

[0057] The computer 241 may also include other removable/non-removable, volatile/nonvolatile computer storage media. By way of example only, FIG. 3B illustrates
15 a hard disk drive 238 that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive 239 that reads from or writes to a removable, nonvolatile magnetic disk 254, and an optical disk drive 240 that reads from or writes to a removable, nonvolatile optical disk 253 such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in
20 the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 238 is typically connected to the system bus 221 through a non-removable memory interface such as interface 234, and magnetic disk drive 239 and optical disk drive 240 are typically connected to the system
25 bus 221 by a removable memory interface, such as interface 235.

[0058] The drives and their associated computer storage media discussed above and illustrated in FIG. 3B, provide storage of computer readable instructions, data structures, program modules and other data for the computer 241. In FIG. 3B, for example, hard disk drive 238 is illustrated as storing operating system 258, application programs 257, other
30 program modules 256, and program data 255. Note that these components can either be the same as or different from operating system 225, application programs 226, other program modules 227, and program data 228. Operating system 258, application programs 257, other program modules 256, and program data 255 are given different numbers here to illustrate that, at a minimum, they are different copies. A user may enter

commands and information into the computer 241 through input devices such as a keyboard 251 and pointing device 252, commonly referred to as a mouse, trackball or touch pad. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected
5 to the processing unit 259 through a user input interface 236 that is coupled to the system bus, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). The cameras 26, 28 and capture device 20 may define additional input devices for the console 100. A monitor 242 or other type of display device is also connected to the system bus 221 via an interface, such as a video
10 interface 232. In addition to the monitor, computers may also include other peripheral output devices such as speakers 244 and printer 243, which may be connected through a output peripheral interface 233.

[0059] The computer 241 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 246. The
15 remote computer 246 may be a personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 241, although only a memory storage device 247 has been illustrated in FIG. 3B. The logical connections depicted in FIG. 3B include a local area network (LAN) 245 and a wide area network (WAN) 249, but may also include
20 other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

[0060] When used in a LAN networking environment, the computer 241 is connected to the LAN 245 through a network interface or adapter 237. When used in a WAN
networking environment, the computer 241 typically includes a modem 250 or other
25 means for establishing communications over the WAN 249, such as the Internet. The modem 250, which may be internal or external, may be connected to the system bus 221 via the user input interface 236, or other appropriate mechanism. In a networked environment, program modules depicted relative to the computer 241, or portions thereof, may be stored in the remote memory storage device. By way of example, and not
30 limitation, FIG. 3B illustrates remote application programs 248 as residing on memory device 247. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

[0061] FIG. 4A depicts an example skeletal mapping of a user that may be generated from the capture device 20. In this embodiment, a variety of joints and bones are

identified: each hand 302, each forearm 304, each elbow 306, each bicep 308, each shoulder 310, each hip 312, each thigh 314, each knee 316, each foreleg 318, each foot 320, the head 322, the torso 324, the top 326 and bottom 328 of the spine, and the waist 330. Where more points are tracked, additional features may be identified, such as the
5 bones and joints of the fingers or toes, or individual features of the face, such as the nose and eyes.

[0062] Through moving his body, a user may create gestures. A gesture comprises a motion or pose by a user that may be captured as image data and parsed for meaning. A gesture may be dynamic, comprising a motion, such as mimicking throwing a ball. A
10 gesture may be a static pose, such as holding one's crossed forearms 304 in front of his torso 324. A gesture may also incorporate props, such as by swinging a mock sword. A gesture may comprise more than one body part, such as clapping the hands 302 together, or a subtler motion, such as pursing one's lips.

[0063] Gestures may be used for input in a general computing context. For instance,
15 various motions of the hands 302 or other body parts may correspond to common system wide tasks such as navigate up or down in a hierarchical list, open a file, close a file, and save a file. Gestures may also be used in a video-game-specific context, depending on the game. For instance, with a driving game, various motions of the hands 302 and feet 320 may correspond to steering a vehicle in a direction, shifting gears, accelerating, and
20 breaking.

[0064] A user may generate a gesture that corresponds to walking or running, by walking or running in place himself. The user may alternately lift and drop each leg 312-320 to mimic walking without moving. The system may parse this gesture by analyzing each hip 312 and each thigh 314. A step may be recognized when one hip-thigh angle (as
25 measured relative to a vertical line, wherein a standing leg has a hip-thigh angle of 0°, and a forward horizontally extended leg has a hip-thigh angle of 90°) exceeds a certain threshold relative to the other thigh. A walk or run may be recognized after some number of consecutive steps by alternating legs. The time between the two most recent steps may be thought of as a period. After some number of periods where that threshold angle is not
30 met, the system may determine that the walk or running gesture has ceased.

[0065] Given a "walk or run" gesture, an application may set values for application-determined parameters associated with this gesture. These parameters may include the above threshold angle, the number of steps required to initiate a walk or run gesture, a number of periods where no step occurs to end the gesture, and a threshold period that

determines whether the gesture is a walk or a run. A fast period may correspond to a run, as the user will be moving his legs quickly, and a slower period may correspond to a walk.

[0066] A gesture may be associated with a set of default parameters at first that the application may override with its own parameters. In this scenario, an application is not forced to provide parameters, but may instead use a set of default parameters that allow the gesture to be recognized in the absence of application-defined parameters.

[0067] There are a variety of outputs that may be associated with the gesture. There may be a baseline “yes or no” as to whether a gesture is occurring. There also may be a confidence level, which corresponds to the likelihood that the user’s tracked movement corresponds to the gesture. This could be a linear scale that ranges over floating point numbers between 0 and 1, inclusive. Wherein an application receiving this gesture information cannot accept false-positives as input, it may use only those recognized gestures that have a high confidence level, such as at least .95. Where an application must recognize every instance of the gesture, even at the cost of false-positives, it may use gestures that have at least a much lower confidence level, such as those merely greater than .2. The gesture may have an output for the time between the two most recent steps, and where only a first step has been registered, this may be set to a reserved value, such as -1 (since the time between any two steps must be positive). The gesture may also have an output for the highest thigh angle reached during the most recent step.

[0068] Another exemplary gesture is a “heel lift jump.” In this, a user may create the gesture by raising his heels off the ground, but keeping his toes planted. Alternatively, the user may jump into the air where his feet leave the ground entirely. The system may parse the skeleton for this gesture by analyzing the angle relation of the shoulders, hips and knees to see if they are in a position of alignment equal to standing up straight. Then these points and upper and lower spine points may be monitored for any upward acceleration. A sufficient combination of acceleration may trigger a jump gesture.

[0069] Given this “heel lift jump” gesture, an application may set values for application-determined parameters associated with this gesture. The parameters may include the above acceleration threshold, which determines how fast some combination of the user’s shoulders, hips and knees must move upward to trigger the gesture, as well as a maximum angle of alignment between the shoulders, hips and knees at which a jump may still be triggered.

[0070] The outputs may comprise a confidence level, as well as the user's body angle at the time of the jump.

[0071] Setting parameters for a gesture based on the particulars of the application that will receive the gesture is important in accurately identifying gestures. Properly

5 identifying gestures and the intent of a user greatly helps in creating a positive user experience. Where a gesture recognizer system is too sensitive, and even a slight forward motion of the hand 302 is interpreted as a throw, the user may become frustrated because gestures are being recognized where he has no intent to make a gesture, and thus, he lacks control over the system. Where a gesture recognizer system is not sensitive enough, the
10 system may not recognize conscious attempts by the user to make a throwing gesture, frustrating him in a similar manner. At either end of the sensitivity spectrum, the user becomes frustrated because he cannot properly provide input to the system.

[0072] Another parameter to a gesture may be a distance moved. Where a user's gestures control the actions of an avatar in a virtual environment, that avatar may be arm's
15 length from a ball. If the user wishes to interact with the ball and grab it, this may require the user to extend his arm 302-310 to full length while making the grab gesture. In this situation, a similar grab gesture where the user only partially extends his arm 302-310 may not achieve the result of interacting with the ball.

[0073] A gesture or a portion thereof may have as a parameter a volume of space in
20 which it must occur. This volume of space may typically be expressed in relation to the body where a gesture comprises body movement. For instance, a football throwing gesture for a right-handed user may be recognized only in the volume of space no lower than the right shoulder 310a, and on the same side of the head 322 as the throwing arm 302a-310a. It may not be necessary to define all bounds of a volume, such as with this
25 throwing gesture, where an outer bound away from the body is left undefined, and the volume extends out indefinitely, or to the edge of scene that is being monitored.

[0074] FIG. 4B provides further details of one exemplary embodiment of the gesture recognizer engine 190 of FIG. 2. As shown, the gesture recognizer engine 190 may comprise at least one filter 418 to determine a gesture or gestures. A filter 418 comprises
30 information defining a gesture 426 (hereinafter referred to as a "gesture"), and may also comprise parameters 428, or metadata, for that gesture. For instance, a throw, which comprises motion of one of the hands from behind the rear of the body to past the front of the body, may be implemented as a gesture 426 comprising information representing the movement of one of the hands of the user from behind the rear of the body to past the front

of the body, as that movement would be captured by the depth camera. Parameters 428 may then be set for that gesture 426. Where the gesture 426 is a throw, a parameter 428 may be a threshold velocity that the hand has to reach, a distance the hand must travel (either absolute, or relative to the size of the user as a whole), and a confidence rating by the recognizer engine that the gesture occurred. These parameters 428 for the gesture 426 may vary between applications, between contexts of a single application, or within one context of one application over time.

[0075] Filters may be modular or interchangeable. In an embodiment, a filter has a number of inputs, each of those inputs having a type, and a number of outputs, each of those outputs having a type. In this situation, a first filter may be replaced with a second filter that has the same number and types of inputs and outputs as the first filter without altering any other aspect of the recognizer engine architecture. For instance, there may be a first filter for driving that takes as input skeletal data and outputs a confidence that the gesture associated with the filter is occurring and an angle of steering. Where one wishes to substitute this first driving filter with a second driving filter – perhaps because the second driving filter is more efficient and requires fewer processing resources – one may do so by simply replacing the first filter with the second filter so long as the second filter has those same inputs and outputs – one input of skeletal data type, and two outputs of confidence type and angle type.

[0076] A filter need not have a parameter. For instance, a “user height” filter that returns the user’s height may not allow for any parameters that may be tuned. An alternate “user height” filter may have tunable parameters – such as to whether to account for a user’s footwear, hairstyle, headwear and posture in determining the user’s height.

[0077] Inputs to a filter may comprise things such as joint data about a user’s joint position, like angles formed by the bones that meet at the joint, RGB color data from the scene, and the rate of change of an aspect of the user. Outputs from a filter may comprise things such as the confidence that a given gesture is being made, the speed at which a gesture motion is made, and a time at which a gesture motion is made.

[0078] A context may be a cultural context, and it may be an environmental context. A cultural context refers to the culture of a user using a system. Different cultures may use similar gestures to impart markedly different meanings. For instance, an American user who wishes to tell another user to “look” or “use his eyes” may put his index finger on his head close to the distal side of his eye. However, to an Italian user, this gesture may be interpreted as a reference to the mafia.

[0079] Similarly, there may be different contexts among different environments of a single application. Take a first-person shooter game that involves operating a motor vehicle. While the user is on foot, making a fist with the fingers towards the ground and extending the fist in front and away from the body may represent a punching gesture.

5 While the user is in the driving context, that same motion may represent a “gear shifting” gesture. There may also be one or more menu environments, where the user can save his game, select among his character’s equipment or perform similar actions that do not comprise direct game-play. In that environment, this same gesture may have a third meaning, such as to select something or to advance to another screen.

10 [0080] The gesture recognizer engine 190 may have a base recognizer engine 416 that provides functionality to a gesture filter 418. In an embodiment, the functionality that the recognizer engine 416 implements includes an input-over-time archive that tracks recognized gestures and other input, a Hidden Markov Model implementation (where the modeled system is assumed to be a Markov process – one where a present state
15 encapsulates any past state information necessary to determine a future state, so no other past state information must be maintained for this purpose – with unknown parameters, and hidden parameters are determined from the observable data), as well as other functionality required to solve particular instances of gesture recognition.

[0081] Filters 418 are loaded and implemented on top of the base recognizer engine 416
20 and can utilize services provided by the engine 416 to all filters 418. In an embodiment, the base recognizer engine 416 processes received data to determine whether it meets the requirements of any filter 418. Since these provided services, such as parsing the input, are provided once by the base recognizer engine 416 rather than by each filter 418, such a service need only be processed once in a period of time as opposed to once per filter 418
25 for that period, so the processing required to determine gestures is reduced.

[0082] An application may use the filters 418 provided by the recognizer engine 190, or it may provide its own filter 418, which plugs in to the base recognizer engine 416. In an embodiment, all filters 418 have a common interface to enable this plug-in characteristic. Further, all filters 418 may utilize parameters 428, so a single gesture tool as described
30 below may be used to debug and tune the entire filter system 418.

[0083] These parameters 428 may be tuned for an application or a context of an application by a gesture tool 420. In an embodiment, the gesture tool 420 comprises a plurality of sliders 422, each slider 422 corresponding to a parameter 428, as well as a pectoral representation of a body 424. As a parameter 428 is adjusted with a

corresponding slider 422, the body 424 may demonstrate both actions that would be recognized as the gesture with those parameters 428 and actions that would not be recognized as the gesture with those parameters 428, identified as such. This visualization of the parameters 428 of gestures provides an effective means to both debug and fine tune a gesture.

[0084] FIG. 5 depicts more complex gestures or filters 418 created from stacked gestures or filters 418. Gestures can stack on each other, and a stacked filter may then be thought of as a filter comprising a plurality of other filters. That is, more than one gesture may be expressed by a user at a single time. For instance, rather than disallowing any input but a throw when a throwing gesture is made, or requiring that a user remain motionless save for the components of the gesture (e.g. stand still while making a throwing gesture that involves only one arm). Where gestures stack, a user may make a jumping gesture and a throwing gesture simultaneously, and both of these gestures will be recognized by the gesture engine.

[0085] FIG. 5A depicts a simple gesture filter 418 according to the stacking paradigm. The IFilter filter 502 is a basic filter 418 that may be used in every gesture filter. IFilter 502 takes user position data 504 and outputs a confidence level 506 that a gesture has occurred. It also feeds that position data 504 into a SteeringWheel filter 508 that takes it as an input and outputs an angle to which the user is steering (e.g. 40 degrees to the right of the user's current bearing) 510.

[0086] FIG. 5B depicts a more complex gesture that stacks filters 418 onto the gesture filter of FIG. 5A. In addition to IFilter 502 and SteeringWheel 508, there is an ITracking filter 512 that receives position data 504 from IFilter 502 and outputs the amount of progress the user has made through a gesture 514. ITracking 512 also feeds position data 504 to GreaseLightning 514 and EBrake 516, which are filters 418 regarding other gestures that may be made in operating a vehicle, such as using the emergency brake.

[0087] There exist other embodiments for stacking gestures. In an embodiment, position data 504 is passed directly to all filters of the stacked gesture, rather than through IFilter 502, and a component of the recognizer engine determines how each filter interoperates. For example, with a jump filter and a throw filter, each may be recognized independently when no other user action is occurring, but this component would recognize that a jump and some user input that may be a throw are occurring simultaneously based on the respective output from each filter. This component may then interpret the throw filter less rigorously (for instance, by expanding the acceptable range of values to satisfy a

parameter), and based on that the “some input that may be a throw” may be recognized as a throw.

[0088] FIG. 6 depicts an example gesture that a user 602 may make to signal for a “fair catch” in a football video game. These figures depict the user at points in time, with FIG.

5 6A being the first point in time, and FIG. 6E being the last point in time. Each of these figures may correspond to a snapshot or frame of image data as captured by a depth camera 402, though not necessarily consecutive frames of image data, as the depth camera 402 may be able to capture frames more rapidly than the user may cover the distance. For instance, this gesture may occur over a period of 3 seconds, and where a depth camera
10 captures data at 40 frames per second, it would capture 60 frames of image data while the user 602 made this fair catch gesture.

[0089] In FIG. 6A, the user 602 begins with his arms 604 down at his sides. He then raises them up and above his shoulders as depicted in FIG. 6B and then further up, to the approximate level of his head, as depicted in FIG. 6C. From there, he lowers his arms 604
15 to shoulder level, as depicted in FIG. 6D, and then again raises them up, to the approximate level of his head, as depicted in FIG. 6E. Where a system captures these positions by the user 602 without any intervening position that may signal that the gesture is cancelled, or another gesture is being made, it may recognize the fair catch gesture as having been made by the user 602.

[0090] FIG. 7 depicts the example “fair catch” gesture of FIG. 5 as each frame of image data has been parsed to produce a skeletal map of the user. The system, having produced a skeletal map from the depth image of the user, may now determine how that user’s body moves over time, and from that, parse the gesture.

[0091] In FIG. 7A, the user’s shoulders 310, are above his elbows 306, which in turn are
25 above his hands 302. The shoulders 310, elbows 306 and hands 302 are then at a uniform level in FIG. 7B. The system then detects in FIG. 7C that the hands 302 are above the elbows, which are above the shoulders 310. In FIG. 7D, the user has returned to the position of FIG. 7B, where the shoulders 310, elbows 306 and hands 302 are at a uniform level. In the final position of the gesture, shown in FIG. 7E, the user returns to the
30 position of FIG. 7C, where the hands 302 are above the elbows, which are above the shoulders 310.

[0092] While the depth camera 402 captures a series of still images, such that in any one image the user appears to be stationary, the user is moving in the course of performing this gesture (as opposed to a stationary gesture, as discussed *supra*). The system is able to take

this series of poses in each still image, and from that determine the moving gesture that the user is making.

[0093] In performing the gesture, a user is unlikely to be able to create an angle as formed by his right shoulder 310a, right elbow 306a and right hand 302a of, for example, between 140° and 145°. So, the application using the filter 418 for the fair catch gesture 428 may tune the associated parameters 426 to best serve the specifics of the application. For instance, the positions in FIGs. 7C and 7E may be recognized any time the user has his hands 302 above his shoulders 310, without regard to elbow 306 position. A set of parameters that are stricter may require that the hands 302 be above the head 310 and that the elbows 306 be both above the shoulders 310 and between the head 322 and the hands 302. Additionally, the parameters 426 for a fair catch gesture 428 may require that the user move from the position of FIG. 7A through the position of FIG. 7E within a specified period of time, such as 1.5 seconds, and if the user takes more than 1.5 seconds to move through these positions, it will not be recognized as the fair catch 418.

[0094] FIG. 8 depicts how generic gesture filters 806 from a gesture filter library 802 are grouped into genre packages 804 of complementary gesture filters for a particular task. The gesture filter library 802 aggregates all gesture filters 806 provided by the system. In an embodiment, an application may provide additional gesture filters for that application's use. Generic gesture filters comprise things such as "arm throw" 806a and "crouch down" 806b. These gesture filters are then grouped together in genre packages 804.

[0095] A genre package 804 may include those gestures that are commonly used within a genre. For instance, a first-person shooter (FPS) genre package 804a may have gesture filters for shooting a weapon 812c, throwing a projectile 812d, punching 812e, opening a door 812f, crouching 812g, jumping 812h, running 812i, and turning 812j. This FPS genre package 804a may be thought of as providing a generic FPS genre package 808a – one with gesture filter parameters tuned or set so that they will likely work acceptably with a large number of FPS applications.

[0096] A genre package is not limited to groups of complementary gesture filters that work for known genres or applications. A genre package may comprise gesture filters that comprise a subset of those filters used by an application or genre, or filters that are complementary, though an appropriate genre for them has yet to be identified.

[0097] An application may then tune those generic genre packages to meet the particulars of that application. The application may tune a generic genre package by setting values for parameters of filters in the genre package. For instance, the creators of

Game A 810a may decide that their game functions best when a demonstrative movement is required to register the opening a door gesture filter 812f, because otherwise it is too similar to the punching gesture filter 812e. However, the creators of Game B may decide that this is not a concern, and require only a more modest movement to register the opening a door gesture filter 812fB.

[0098] In the embodiment where a genre package comprises machine-readable instructions, a genre package may be provided as those instructions in source code form, or in a form reflecting some amount of compilation of those instructions.

[0099] FIG. 9 illustrates exemplary operational procedures for tuning complementary gesture filters in a filter package when an application provides a value for one parameter of one filter.

[0100] Operation 902 depicts providing a package comprising a plurality of filters, each filter comprising information about a gesture and at least one parameter, each filter being complementary with at least one other filter in the package. The package may represent gesture filters for a particular genre. For example, genre packages for video games may include genres such as first-person shooter, action, driving, and sports.

[0101] As used herein, and in at least one embodiment, “providing a package” may refer to allowing access to a programming language library file that corresponds to the filters in the package or allowing access to an application programming interface (API) to an application. The developer of the application may load the library file and then make method calls as appropriate. For instance, with a sports package there may be a corresponding sports package library file.

[0102] When included in the application, the application may then make calls that use the sports package according to the given API. Such API calls may include returning the value of a parameter for a filter, setting the value of a parameter for a filter, and correlating identification of a filter with triggering some part of the application, such as causing a user controlled tennis player to swing a tennis racket when the user makes the appropriate tennis racket swing gesture.

[0103] As described above, a gesture may comprise a wide variety of things. It may, for instance, be any of a crouch, a jump, a lean, an arm throw, a toss, a swing, a dodge, a kick, and a block. Likewise, a gesture may correspond to navigation of a user interface. For instance, a user may hold his hand with the fingers pointing up and the palm facing the 3D camera. He may then close his fingers towards the palm to make a fist, and this could be a

gesture that indicates that the focused window in a window-based user-interface computing environment should be closed.

[0104] As gestures may be used to indicate anything from that an avatar in an application should throw a punch to that a window in an application should be closed, a wide variety of applications, from video games to text editors may utilize gestures.

[0105] Complementary gesture filters – either complementary as in those that are commonly used together, or complementary as in a change in a parameter of one will change a parameter of another – may be grouped together into genre packages that are likely to be used by an application in that genre. These packages may be available or identified to an application, which may select at least one. The application may tune, or modify, the parameter(s) of a filter in a selected package to best fit the unique aspects of the application. When that parameter is tuned, a second, complementary parameter (in the inter-dependent sense) of either the filter or a second filter may also be tuned such that the parameters remain complementary.

[0106] The application-determined parameter may vary based on the context the application is in. To that end, an application may assign a plurality of values to an application-determined parameter for a filter, each value corresponding to a different context. As discussed *supra*, this context may be a cultural context or an environmental context.

[0107] Operation 904 depicts receiving an indication of assigning a value to a parameter of a first filter. An application-determined parameter may comprise any of a wide variety of characteristics of a filter, such as a body part, a volume of space, a velocity, a direction of movement, an angle, and a place where a movement occurs.

[0108] In an embodiment, the value of the application-determined parameter is determined by an end user of the application through making a gesture. For instance, an application may allow the user to train it, so that the user is able to specify what motions he believes a gesture should comprise. This may be beneficial to allow a user without good control over his motor skills to be able to link what motions he can make with a corresponding gesture. If this were not available, the user may become frustrated because he is unable to make his body move in the manner required by the application to produce the gesture.

[0109] In an embodiment where there exist complementary filters – a plurality of filters that have inter-related parameters – receiving from the application a value for an application-determined parameter of the first filter may include both setting the

application-determined parameter of the first filter with the value, and setting a complementary application-determined parameter of a second, complementary filter based on the value of the parameter of the first filter. For example, one may decide that a user who throws a football in a certain manner is likely to also throw a baseball in a certain manner. So, where it is determined that a certain application-determined parameter of one filter, such as a velocity parameter on a filter for a football throw gesture, should be set in a particular manner, other complementary application-determined parameters, such as the velocity parameter on a baseball throw gesture, may be set based on how that first application-determined parameter is set.

[0110] This need not be the same value for a given application-determined parameter, or even the same type of application-determined parameter across filters. For instance, it could be that when a football throw must be made with a forward arm velocity of X m/s, then a football catch must be made with the hands at least distance Y m away from the torso.

[0111] The value may be a threshold, such as arm velocity is greater than X . It may be an absolute, such as arm velocity equals X . There may be a fault tolerance, such as arm velocity equals within Y of X . It may also comprise a range, such as arm velocity is greater than or equal to X , but less than Z .

[0112] Operation 906 depicts assigning the value to the parameter of the first filter.

Where an association between parameters and their values is stored in a database, this may comprise storing the value in the database along with an association with the parameter.

[0113] Operation 908 depicts assigning a second value to a second parameter of a second filter, the second value determined using the value assigned to the parameter of the first filter. As discussed above, the second value may relate to the first value in a variety of ways. Where the two parameters involve something substantially similar such as a threshold jump height, the second value may be equal to the first value. The second value and the first value may have a variety of other relationships, such as a proportional relationship, an inversely proportional relationship, a linear relationship, an exponential relationship, and a function that takes the value as an input.

[0114] In an embodiment where filters may inherit characteristics from each other, such as in an object-oriented implementation, the second filter may comprise a child of the first filter, with the first filter likewise being a parent to the second filter. Take for example, a “hand slap” filter. This filter may serve as a parent to variations on hand slaps, such as the “high five,” the “high ten” and the “low five.” Where the “hand slap” has a “hand

movement distance threshold” parameter, when the value to that parameter is set, the “hand movement distance threshold” parameter for all child filters may be set with that same value.

[0115] Likewise, the complementary nature of two parameters may be due to one filter being stacked to be incorporated into another filter. One filter may be a steering filter, and that is stacked with other filters such as gear shift, accelerate and decelerate to create a driving filter. As the “minimum steering angle threshold” parameter of the steering filter is modified, the corresponding “minimum steering angle threshold” parameter of the driving filter may also be modified.

[0116] Operation 910 depicts the optional operation of receiving data comprising an image of a person, and when parsing filters to determine if the data matches a particular filter (and thus indicates a particular gesture), parsing the data for each filter in a selected package before parsing the data for a filter not in the package. Where an application selects a filter package for use, such as by including a library file for that filter package, it likely does so because those filters are to be frequently used by a user of the application. Further, filters in a filter package may be used in close succession, such as with run, jump, strafe, crouch and discharge firearm filters in a first-person shooter package. To this end, where a filter package has been identified as being used by an application, a system processing filters, such as the base filter engine described above, can likely reduce the processing resources required to process image data corresponding to user input by first processing the data for those filters comprising the selected filter package.

[0117] Operation 912 depicts the optional operation of receiving an indication of assigning a third value to the second parameter, and assigning the third value to the second parameter. A relationship between two parameters need not be bilateral. It may be that a change in the value affects the second value, but that a change in the second value does not affect the first value.

[0118] Operation 914 depicts the optional operation of assigning a fourth value to the parameter, the fourth value determined using the third value. It may also be that the relationship between the two parameters are bilateral. In this embodiment, a change in a value of the second parameter results in changing a value of the first parameter, as determined using the new value of the second parameter.

[0119] Operation 916 depicts the optional operation of receiving data comprising an image of a person; determining that the data matches the first filter of the package; and parsing the data for each other filter in the package that can be indicated by the user

simultaneously with the first filter before parsing the data for a filter that cannot be indicated by the user simultaneously with the filter. In an embodiment, filters in a filter package may be used simultaneously, such as a simultaneous run filter and discharge firearm filter in a first-person shooter package. To this end, where a filter package has been identified as being used by an application, a system processing filters, such as the base recognizer engine discussed above, can likely reduce the processing resources required to process image data corresponding to user input by first processing the data for those filters comprising the filter package.

Conclusion

[0120] While the present disclosure has been described in connection with the preferred aspects, as illustrated in the various figures, it is understood that other similar aspects may be used or modifications and additions may be made to the described aspects for performing the same function of the present disclosure without deviating there from. Therefore, the present disclosure should not be limited to any single aspect, but rather construed in breadth and scope in accordance with the appended claims. For example, the various procedures described herein may be implemented with hardware or software, or a combination of both. Thus, the methods and apparatus of the disclosed embodiments, or certain aspects or portions thereof, may take the form of program code (i.e., instructions) embodied in tangible media, such as floppy diskettes, CD-ROMs, hard drives, or any other machine-readable storage medium. When the program code is loaded into and executed by a machine, such as a computer, the machine becomes an apparatus configured for practicing the disclosed embodiments. In addition to the specific implementations explicitly set forth herein, other aspects and implementations will be apparent to those skilled in the art from consideration of the specification disclosed herein. It is intended that the specification and illustrated implementations be considered as examples only.

What is Claimed:

1. A method for providing a package of complementary gesture filters to an application, comprising:

providing a package comprising a plurality of filters, each filter comprising
5 information about a gesture, at least one filter being complementary with at least one other filter in the package (902);

receiving an indication of assigning a first value to a parameter of a first filter
(904);

assigning the value to the parameter (906); and

10 assigning a second value to a second parameter of a second filter, the second value determined using the first value (908).

2. The method of claim 1, wherein the parameter of the first filter represents a body part, a volume of space, a velocity, a direction of movement, an angle, a two-dimensional (2D) plane, or a place where a movement occurs.

15 3. The method of claim 1, wherein the gesture comprises a crouch, a jump, a lean, an arm throw, a toss, a swing, a dodge, a kick, or a block.

4. The method of claim 1, wherein a filter is complementary with at least one other filter in the package when (i) that filter has at least one parameter that is determined based on a parameter of the at least one other filter in the package, (ii) that filter represents a
20 gesture that is commonly made by a user within a short time period of a gesture represented by the at least one other filter in the package, or (iii) the gesture represented by that filter is capable of being made simultaneously with a gesture represented by the at least one other filter in the package.

5. The method of claim 1, wherein the indication of assigning the value to the
25 parameter is received as a result of the user making the gesture.

6. The method of claim 1, wherein the value is a fixed value, a range, or a value with a tolerance.

7. The method of claim 1, wherein the second value is determined using the first value based on a proportional relationship, an inversely proportional relationship, a linear
30 relationship, an exponential relationship, or a function that takes the first value as an input.

8. The method of claim 1, wherein the package contains complementary gesture filters for a particular genre, and wherein the genre is one of a first-person shooter, action, driving, or sports genre.

9. The method of claim 1, wherein the application is a video game or an operating system.

10. The method of claim 1, further comprising:

receiving an indication of assigning a third value to the second parameter; and
5 assigning the third value to the second parameter.

11. A system for providing a package of complementary filters to an application, comprising:

a filter library comprising at least one package of filters, each filter comprising information about a gesture, at least one filter of the at least one package being
10 complementary with at least one other filter in the package (802); and

circuitry that provides a filter package to the application, receives data corresponding to at least one image of at least part of a user, determines that the data corresponds to at least two matched filters, determines how a first filter of the matched filters operates with at least one other filter of the matched filters, and sends an output
15 corresponding to each matched filter to the application (12).

12. The system of claim 11, wherein the output comprises a confidence level, and the circuitry determines how a first matched filter operates with at least one other matched filter by altering the confidence level of the first filter based on the confidence level of the at least one other matched filter.

20 13. The system of claim 11, wherein the circuitry further:

parses the data for each filter in the package to determine if the data indicates a match with one or more of the filters of the package, before parsing the data for a filter not in the package.

14. The system of claim 11, wherein the circuitry further:

25 parses the data for each other filter in the package that can be parsed simultaneously with the at least one filter that corresponds to the data before parsing the data for a filter that cannot be parsed simultaneously with the at least one filter that corresponds to the data

15. A computer readable storage medium, comprising computer readable instructions
30 that when executed on a processor, cause the processor to perform the operations of:

providing a package comprising a plurality of filters, each filter comprising information about a gesture, each filter being complementary with at least one other filter in the package, a gesture being input by a user making a motion or pose associated with that gesture that is captured by a distance camera (902);

receiving an indication of assigning a first value to a parameter of a first filter
(904);

assigning the value to the parameter (906); and

5 assigning a second value to a second parameter of a second filter, the second value
determined using the first value (908).




FIG. 1A

2/12




FIG. 1B




FIG. 2

4/12




FIG. 3A

Computing Environment 220





FIG. 3B

6/12

**FIG. 4A**




FIG. 4B

8/12




FIG. 5A









FIG. 5B






FIG. 7A

FIG. 7B

FIG. 7C

FIG. 7D

FIG. 7E

11/12




FIG. 8

12/12




FIG. 9