



(19) **United States**

(12) **Patent Application Publication**
Berenzweig et al.

(10) **Pub. No.: US 2020/0069211 A1**
(43) **Pub. Date: Mar. 5, 2020**

(54) **CAMERA-GUIDED INTERPRETATION OF NEUROMUSCULAR SIGNALS**

A61B 5/00 (2006.01)
A61B 90/00 (2006.01)
A61B 5/04 (2006.01)

(71) Applicant: **CTRL-labs Corporation**, New York, NY (US)

(52) **U.S. Cl.**
CPC *A61B 5/0488* (2013.01); *G06F 3/011* (2013.01); *A61B 5/6824* (2013.01); *A61B 90/361* (2016.02); *G06T 19/006* (2013.01); *A61B 5/04004* (2013.01); *A61B 5/04012* (2013.01); *A61B 5/681* (2013.01); *G06F 3/017* (2013.01); *A61B 5/04001* (2013.01)

(72) Inventors: **Adam Berenzweig**, Brooklyn, NY (US); **Thomas Reardon**, New York, NY (US); **Christopher Osborn**, New York, NY (US); **Patrick Kaifosh**, New York, NY (US); **Brett Jurman**, New York, NY (US); **Daniel Wetmore**, Brooklyn, NY (US)

(73) Assignee: **CTRL-labs Corporation**, New York, NY (US)

(21) Appl. No.: **16/557,427**

(22) Filed: **Aug. 30, 2019**

Related U.S. Application Data

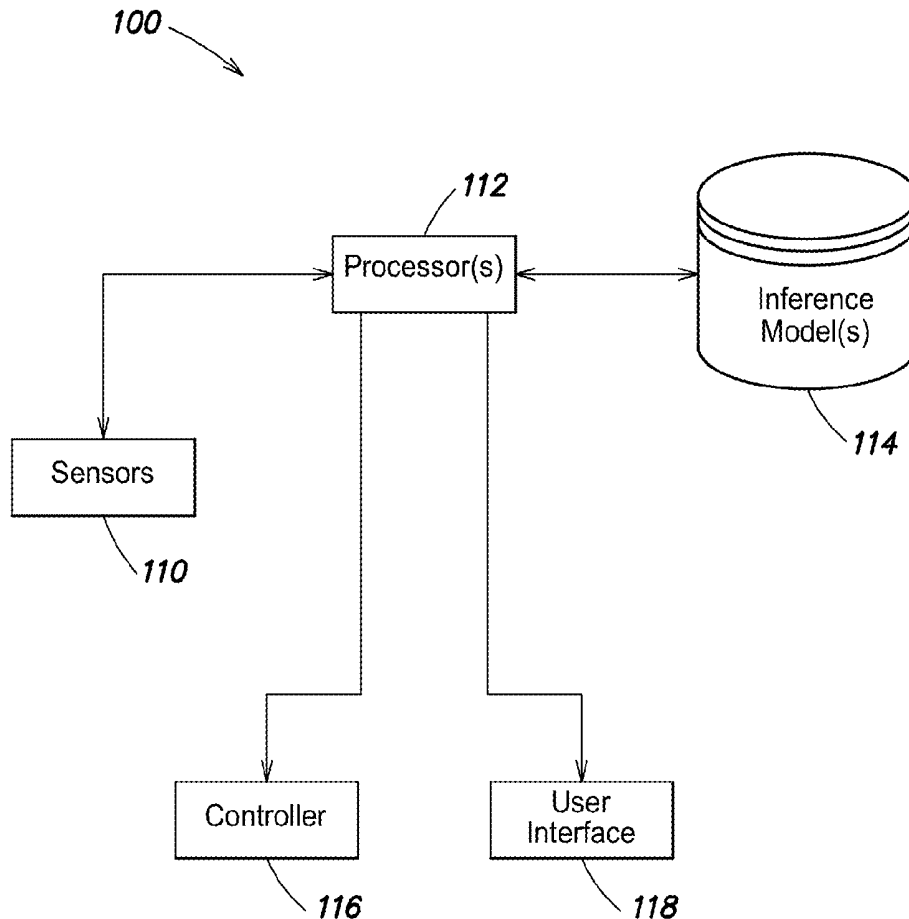
(60) Provisional application No. 62/726,159, filed on Aug. 31, 2018.

Publication Classification

(51) **Int. Cl.**
A61B 5/0488 (2006.01)
G06F 3/01 (2006.01)

(57) **ABSTRACT**

Computerized systems, methods, and computer-readable storage media storing code for implementing the methods are provided for training an inference model used to generate a musculoskeletal representation. One such system includes a processor programmed to: determine, based on information obtained from at least one image, position information describing a spatial relationship between two or more connected musculoskeletal segments of a user; determine, based on a plurality of neuromuscular signals, force information; associate the position information with the force information; train an inference model to output a musculoskeletal representation consistent with the position information and/or the force information when neuromuscular input signals provided to the inference model have at least one predetermined characteristic, to produce an updated inference model; and cause the updated inference model to be stored in a memory.



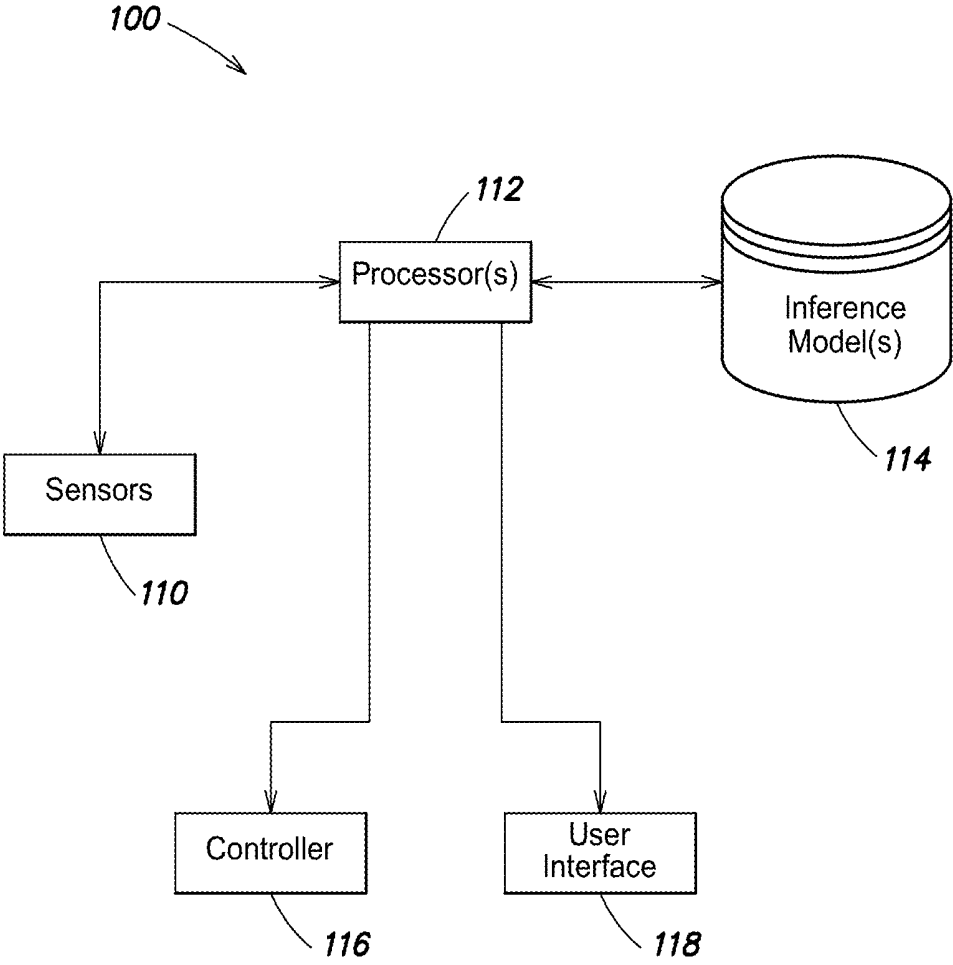


FIG. 1




FIG. 2

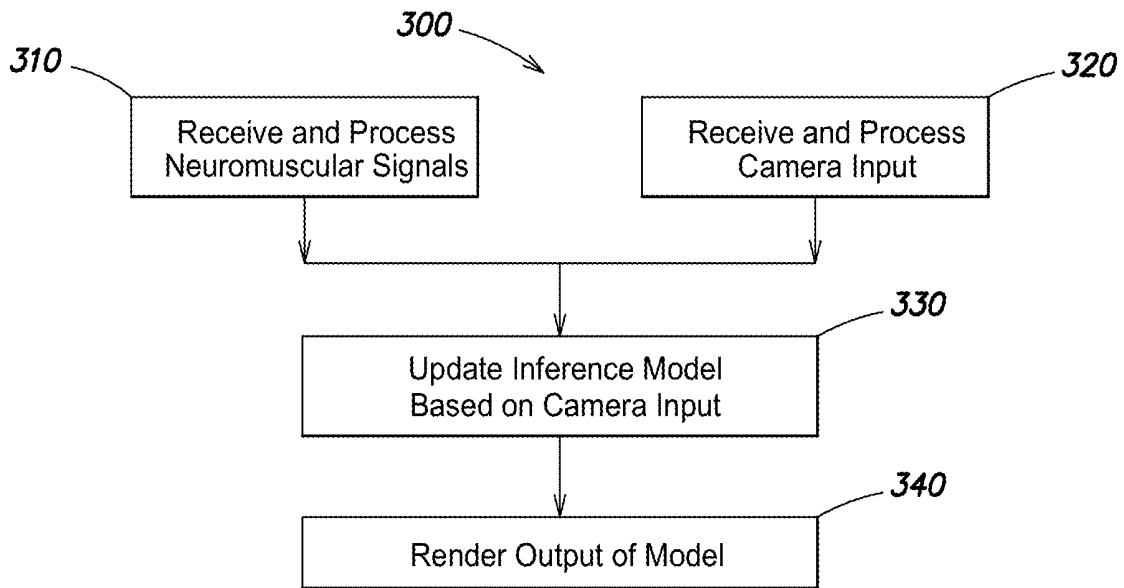


FIG. 3

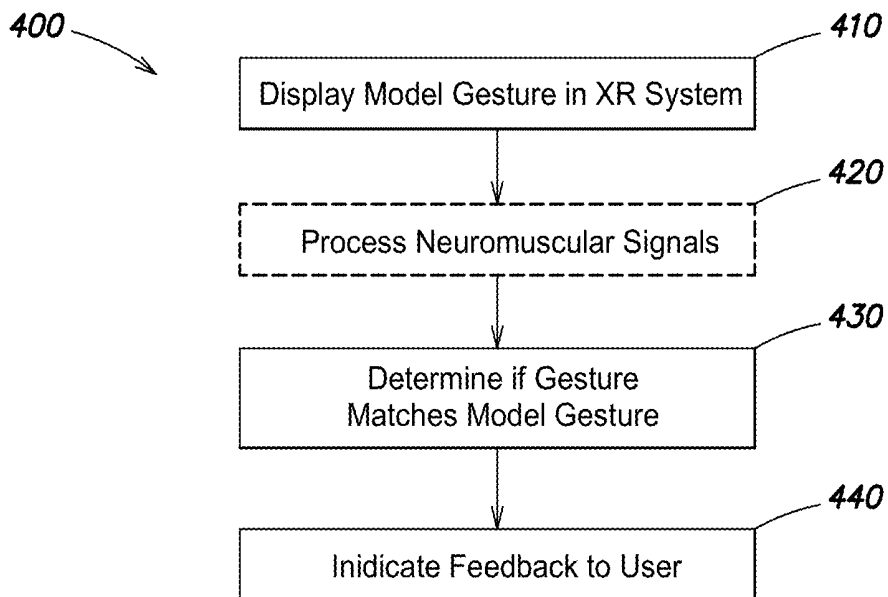


FIG. 4

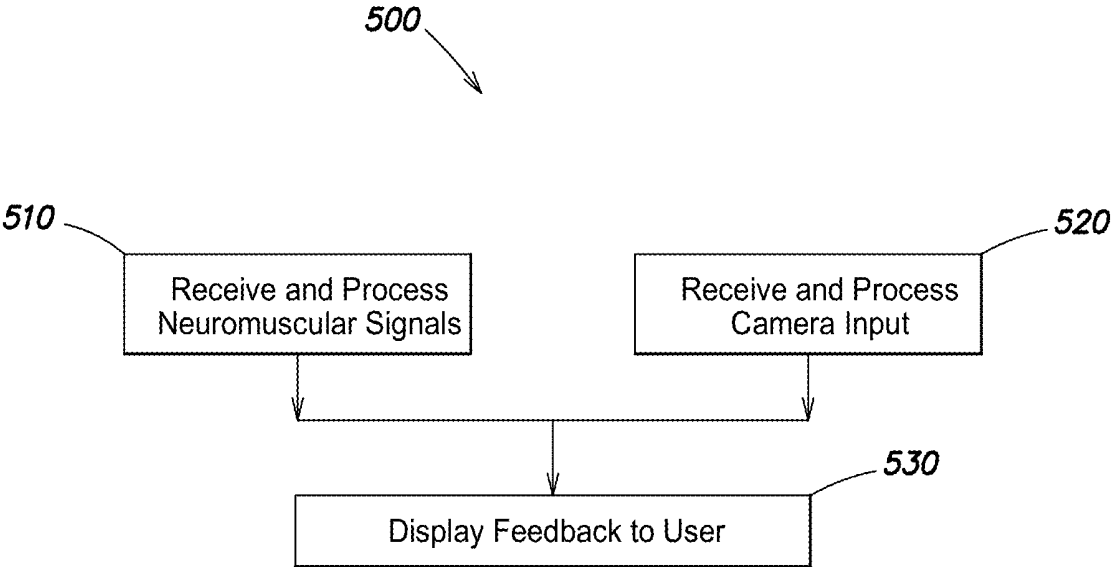


FIG. 5

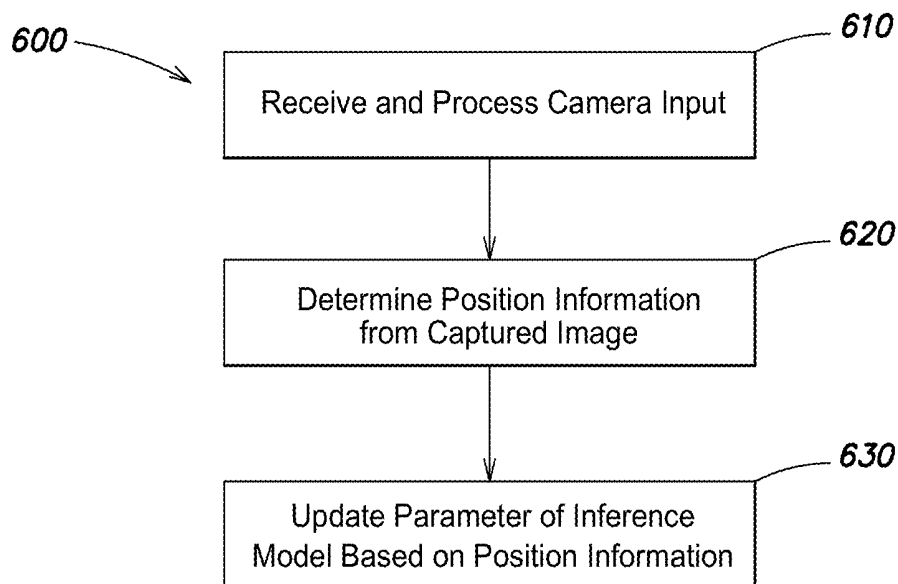


FIG. 6

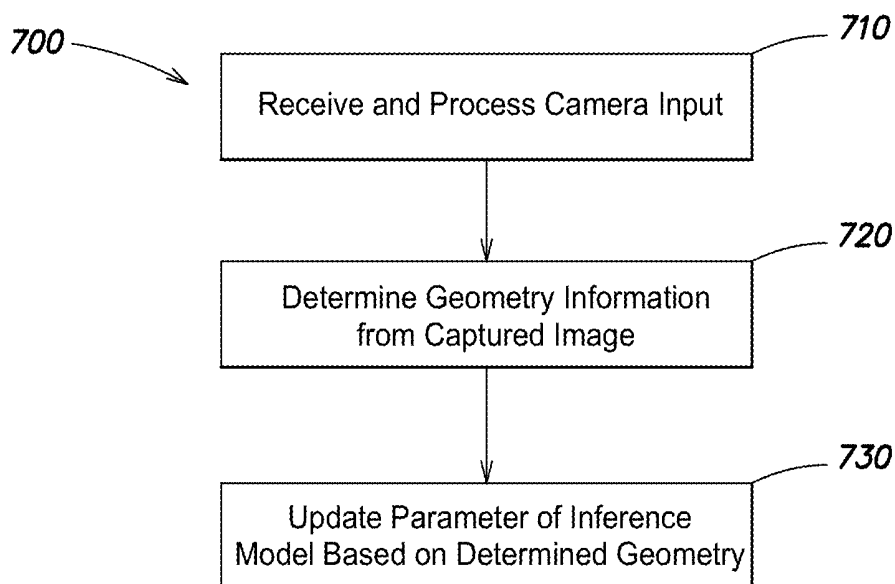


FIG. 7

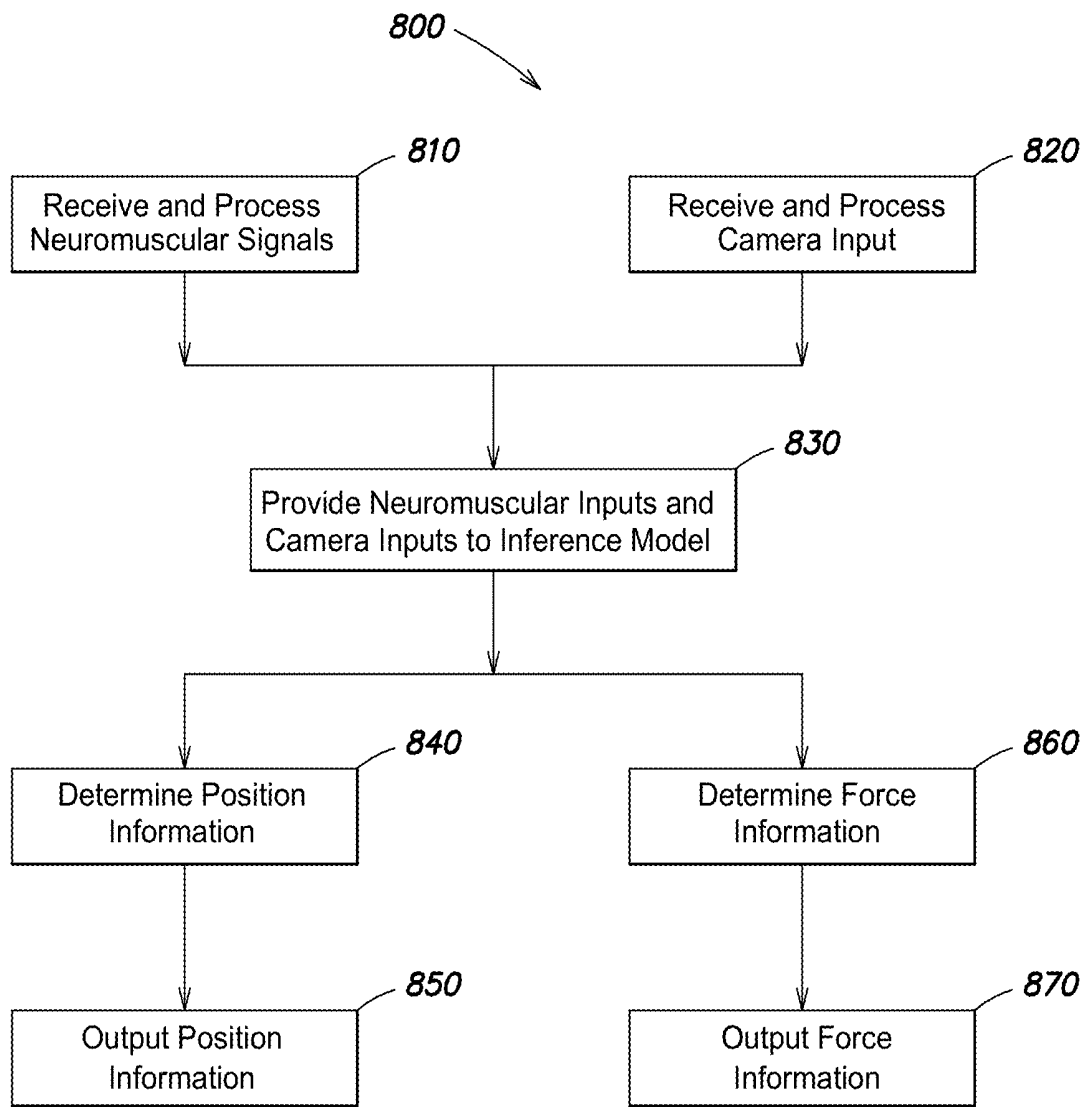


FIG. 8

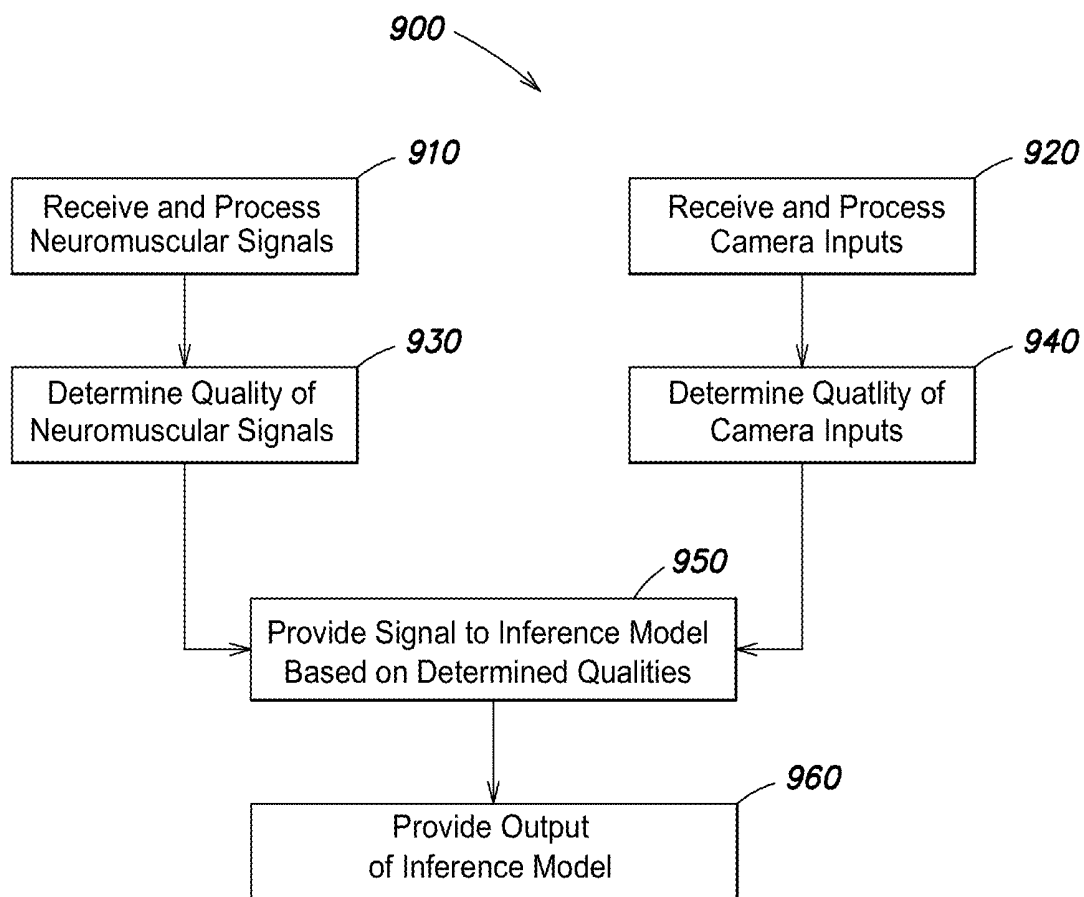


FIG. 9




FIG. 10




FIG. 11




FIG. 12




FIG. 13A




FIG. 13B




FIG. 14A




FIG. 14B

CAMERA-GUIDED INTERPRETATION OF NEUROMUSCULAR SIGNALS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit under 35 U.S.C. § 119(e) of U.S. Provisional Patent Application Ser. No. 62/726,159, filed Aug. 31, 2018, entitled “CAMERA-GUIDED INTERPRETATION OF NEUROMUSCULAR SIGNALS”, the entire contents of which are incorporated by reference herein.

FIELD OF THE INVENTION

[0002] The present technology relates to systems and methods that detect and interpret neuromuscular signals for use in performing functions in an augmented reality environment or a virtual reality environment, and that calibrate models used to interpret neuromuscular signals.

BACKGROUND

[0003] In some computer applications that generate musculoskeletal representations of a human body, it is desirable for the application to know the spatial positioning, orientation, and movement of a user’s body to provide a realistic representation of body movement. For example, in an augmented reality (AR) environment or a virtual reality (VR) environment, tracking the spatial position of the user’s hand enables the application to represent hand motion in the AR or VR environment, which allows the user to interact with (e.g., by grasping or manipulating) virtual objects within the AR or VR environment. Techniques have been developed for tracking movements of a user’s body using wearable sensors, such as Inertial Measurement Units (IMUs), affixed to different parts of the user’s body, to obtain position information and/or orientation information about the different parts of the user’s body.

SUMMARY

[0004] According aspects of the technology described herein, a computerized system for using camera information to calibrate one or more inference models used to generate a musculoskeletal representation is provided. The system may comprise at least one camera configured to capture at least one image; a plurality of neuromuscular sensors configured to sense and record a plurality of neuromuscular signals from a user, wherein the plurality of neuromuscular sensors are arranged on one or more wearable devices structured to be worn by the user to obtain the plurality of neuromuscular signals; and at least one computer processor. The at least one processor may be programmed to calibrate one or more inference models by updating at least one parameter associated with the one or more inference models based, at least in part, on the plurality of neuromuscular signals and the at least one image.

[0005] In an aspect, the updating of the at least one parameter may include any one or any combination of: an updating of a routine for pre-processing the plurality of neuromuscular signals to calibrate the one or more inference models; an updating of a routine for pre-processing an image signal corresponding to the at least one image before the image signal is used to calibrate the one or more inference models; an updating of an architecture or architectures of the one or more inference models; an updating of a routine for

post-processing an output of the one or more inference models; an updating of a selection routine for selecting the one or more inference models from a plurality of inference models; and an updating of a weight used for at least one of the one or more inference models.

[0006] In another aspect, the one or more inference models may be usable to generate a representation of any one or any combination of: a handstate, a still pose, and a dynamic gesture.

[0007] In an aspect, the at least one camera may comprise a camera disposed at a fixed location relative to the user.

[0008] In an aspect, the at least one camera may comprise a camera configured to be mounted on the user.

[0009] In an aspect, the at least one image may include any one or any combination of: an image produced by visible light, an image produced by infrared light, an image produced by light of a predetermined range of wavelengths, and an image produced by light of two or more different predetermined ranges of wavelengths.

[0010] In an aspect, the updating of the at least one parameter associated with the one or more inference models may comprise training the one or more inference models to produce an output that is determined based, at least in part, on the at least one image when the plurality of neuromuscular signals or information derived from the plurality of neuromuscular signals, or both the plurality of neuromuscular signals and the information derived from the plurality of neuromuscular signals, is or are provided as input to the one or more inference models.

[0011] In an aspect, the computerized system may further comprise an extended reality (XR) system that generates an XR environment and displays a visual representation. The at least one computer processor may be further programmed to: detect, based at least in part on the at least one image, whether a gesture performed by the user matches a particular gesture; and, if the gesture performed by the user is detected to match the particular gesture, update the one or more inference models to incorporate neuromuscular data corresponding to the gesture performed by the user.

[0012] In a variation of this aspect, the at least one camera may comprise a camera attached to a head-mounted-display of the XR system.

[0013] In another variation of this aspect, the at least one computer processor may be further programmed to instruct the XR system to provide a visible prompt for the user to perform the particular gesture.

[0014] In another variation of this aspect, the at least one computer processor may be further programmed to instruct the XR system to provide an audible prompt for the user to perform the particular gesture.

[0015] In another variation of this aspect, the gesture may be performed by the user upon an instruction from an external source.

[0016] In another variation of this aspect, the gesture performed by the user may be detected to match the particular gesture when a similarity between the gesture performed by the user and the particular gesture is above a predetermined threshold.

[0017] In another variation of this aspect, the visual representation may be displayed on a display screen viewable by the user, and may comprise a visual representation of a hand.

[0018] In another variation of this aspect, the visual representation may be displayed on a display screen viewable by the user, and may comprise instructions to the user.

[0019] In another variation of this aspect, the prompt for the user to perform the particular gesture may be a visual representation of the particular gesture provided on a display screen viewable by the user.

[0020] In various variations of this aspect, the visual representation may be within the XR environment.

[0021] In a variation of this aspect, the particular gesture may be one of: a motion or dynamic gesture, and a pose or static gesture.

[0022] In a variation of this aspect, the at least one image may be captured by the at least one camera during performance of the gesture by the user, and the at least one computer processor may be further programmed to cause feedback to be provided to the user based, at least in part, on the at least one image captured during the performance of the gesture by the user. In one example, the feedback to be provided to the user may indicate whether a body part of the user necessary for the particular gesture to be performed is fully included in the at least one image captured during the performance of the gesture by the user. In another example, the feedback to be provided to the user may indicate whether a body part of the user necessary for the particular gesture to be performed is fully or partially occluded in the at least one image captured during the performance of the gesture by the user. In another example, the at least one computer processor may instruct the XR system to provide the feedback within the XR environment. In another example, the feedback may comprise one or more of: audible feedback, visible feedback, haptic feedback, and electrical stimulation feedback. For instance, the visible feedback may be a moving image or a still image showing a way to perform the particular gesture.

[0023] In an aspect, the at least one computer processor may be further programmed to determine whether the at least one image includes position information for at least two segments of a musculoskeletal representation. The updating of the at least one parameter associated with the one or more inference models may comprise updating the one or more inference models to incorporate neuromuscular data corresponding to the plurality of neuromuscular signals, when it is determined that the at least one image includes position information for at least two segments of the musculoskeletal representation.

[0024] In an aspect, the at least one computer processor may be further programmed to: detect based, at least in part, on the at least one image, whether a gesture performed by the user to perform a task within an extended reality (XR) environment generated by an XR system matches a stored gesture for performing the task; and, if the gesture performed by the user matches the stored gesture for performing the task, update the one or more inference models to incorporate neuromuscular data corresponding to the gesture performed by the user.

[0025] In a variation of this aspect, the gesture performed by the user may be detected to match the stored gesture when a similarity between the gesture performed by the user and the stored gesture is above a predetermined threshold.

[0026] In another variation of this aspect, the one or more inference models may be usable to generate a representation of any one or any combination of: a handstate, a still pose, and a dynamic gesture.

[0027] In an aspect, the at least one computer processor may be further programmed to determine, based on the at least one image, a user-specific skeletal geometry corresponding to the user. The updating of the at least one parameter associated with the one or more inference models may comprise updating the at least one parameter of the one or more inference models based on the determined user-specific skeletal geometry.

[0028] In a variation of this aspect, the user-specific skeletal geometry may comprise a length of at least one finger of the user.

[0029] In an aspect, the at least one camera may be configured to capture a plurality of images in a time series. The updating of the at least one parameter of the one or more inference models may be further based, at least in part, on the plurality of images captured in a time series.

[0030] According aspects of the technology described herein, a method of a computerized system for using camera information to calibrate one or more inference models used to generate a musculoskeletal representation is provided. The method may comprise: receiving, by at least one computer processor, one or both of: at least one image captured by at least one camera and information derived from the at least one image; receiving, by the at least one computer processor, one or both of: a plurality of neuromuscular signals from a user and information derived from the plurality of neuromuscular signals, the plurality of neuromuscular signals being sensed and recorded by plurality of neuromuscular sensors arranged on one or more wearable devices worn by the user; and calibrating, by the at least one processor, one or more inference models by updating at least one parameter associated with the one or more inference models based, at least in part, on the plurality of neuromuscular signals and the at least one image.

[0031] In an aspect, the updating of the at least one parameter may include any one or any combination of: an updating of a routine for pre-processing the plurality of neuromuscular signals to calibrate the one or more inference models; an updating of a routine for pre-processing an image signal corresponding to the at least one image before the image signal is used to calibrate the one or more inference models; an updating of an architecture or architectures of the one or more inference models; an updating of a routine for post-processing an output of the one or more inference models; an updating of a selection routine for selecting the one or more inference models from a plurality of inference models; and an updating of a weight used for at least one of the one or more inference models.

[0032] In an aspect, the one or more inference models may be usable to generate a representation of any one or any combination of: a handstate, a still pose, and a dynamic gesture.

[0033] In an aspect, the at least one camera may comprise a camera disposed at a fixed location relative to the user.

[0034] In an aspect, the at least one camera may comprise a camera configured to be mounted on the user.

[0035] In an aspect, the at least one image may include any one or any combination of: an image produced by visible light, an image produced by infrared light, an image produced by light of a predetermined range of wavelengths, and an image produced by light of two or more different predetermined ranges of wavelengths.

[0036] In an aspect, the updating of the at least one parameter associated with the one or more inference models

may comprise training the one or more inference models to produce an output that is determined based, at least in part, on the at least one image when the plurality of neuromuscular signals or information derived from the plurality of neuromuscular signals, or both the plurality of neuromuscular signals and the information derived from the plurality of neuromuscular signals, is or are provided as input to the one or more inference models.

[0037] In an aspect, the method may further comprise: detecting, by the at least one processor, based at least in part on the at least one image, whether a gesture performed by the user matches a particular gesture; and, if the gesture performed by the user is detected to match the particular gesture, updating the one or more inference models to incorporate neuromuscular data corresponding to the gesture performed by the user. The at least one processor may be in communication with an extended reality (XR) system that generates an XR environment and displays a visual representation.

[0038] In a variation of this aspect, the at least one camera may comprise a camera attached to a head-mounted-display of the XR system.

[0039] In another variation of this aspect, the method may further comprise instructing, by the at least one processor, the XR system to provide a visible prompt for the user to perform the particular gesture.

[0040] In another variation of this aspect, the method may further comprise instructing, by the at least one processor, the XR system to provide an audible prompt for the user to perform the particular gesture.

[0041] In a variation of some of the aspects, the gesture may be performed by the user upon an instruction from an external source.

[0042] In a variation of some of the aspects, the gesture performed by the user may be detected to match the particular gesture when a similarity between the gesture performed by the user and the particular gesture is above a predetermined threshold.

[0043] In a variation of some of the aspects, the visual representation may be displayed on a display screen viewable by the user, and may comprise a visual representation of a hand.

[0044] In a variation of some of the aspects, the visual representation may be displayed on a display screen viewable by the user, and may comprise instructions to the user.

[0045] In a variation of some of the aspects, the prompt for the user to perform the particular gesture may be a visual representation of the particular gesture provided on a display screen viewable by the user.

[0046] In a variation of some of the aspects, the visual representation may be within the XR environment.

[0047] In a variation of some of the aspects, the particular gesture may be one of: a motion or dynamic gesture, and a pose or static gesture.

[0048] In a variation of some of the aspects, the at least one image may be captured by the at least one camera during performance of the gesture by the user. The method may further comprise the at least one processor causing feedback to be provided to the user based, at least in part, on the at least one image captured during the performance of the gesture by the user.

[0049] In a variation of some of the aspects, the feedback to be provided to the user may indicate whether a body part of the user necessary for the particular gesture to be per-

formed is fully included in the at least one image captured during the performance of the gesture by the user.

[0050] In a variation of some of the aspects, the feedback to be provided to the user may indicate whether a body part of the user necessary for the particular gesture to be performed is fully or partially occluded in the at least one image captured during the performance of the gesture by the user.

[0051] In a variation of some of the aspects, the at least one computer processor may cause the XR system to provide the feedback within the XR environment.

[0052] In a variation of some of the aspects, the feedback may comprise one or more of: audible feedback, visible feedback, haptic feedback, and electrical stimulation feedback. In a further variation, the visible feedback may be a moving image or a still image showing a way to perform the particular gesture.

[0053] In an aspect, the method may further comprise determining, by the at least one computer processor, whether the at least one image includes position information for at least two segments of a musculoskeletal representation. The updating of the at least one parameter associated with the one or more inference models may comprise updating the one or more inference models to incorporate neuromuscular data corresponding to the plurality of neuromuscular signals, when it is determined that the at least one image includes position information for at least two segments of the musculoskeletal representation.

[0054] In an aspect, the method may further comprise: detecting, by the at least one processor, based at least in part on the at least one image, whether a gesture performed by the user to perform a task within an extended reality (XR) environment generated by an XR system matches a stored gesture for performing the task; and, if the gesture performed by the user matches the stored gesture for performing the task, updating, by the at least one processor, the one or more inference models to incorporate neuromuscular data corresponding to the gesture performed by the user.

[0055] In a variation of this aspect, the gesture performed by the user may be detected to match the stored gesture when a similarity between the gesture performed by the user and the stored gesture is above a predetermined threshold.

[0056] In another variation of this aspect, the one or more inference models may be usable to generate a representation of any one or any combination of: a handstate, a still pose, and a dynamic gesture.

[0057] In an aspect, the method may further comprise: determining, by the at least one computer processor, based on the at least one image, a user-specific skeletal geometry corresponding to the user. The updating of the at least one parameter associated with the one or more inference models may comprise updating the at least one parameter of the one or more inference models based on the determined user-specific skeletal geometry.

[0058] In a variation of this aspect, the user-specific skeletal geometry may comprise a length of at least one finger of the user.

[0059] In an aspect, the at least one camera may be configured to capture a plurality of images in a time series. The updating of the at least one parameter of the one or more inference models may be further based, at least in part, on the plurality of images captured in a time series.

[0060] According aspects of the technology described herein, a non-transitory computer-readable storage medium storing code that, when executed by at least one computer,

causes the at least one computer to perform a method for using camera information to calibrate one or more inference models used to generate a musculoskeletal representation is provided. The method may comprise: receiving one or both of: at least one image captured by at least one camera and information derived from the at least one image; receiving one or both of: a plurality of neuromuscular signals from a user and information derived from the plurality of neuromuscular signals, the plurality of neuromuscular signals being sensed and recorded by plurality of neuromuscular sensors arranged on one or more wearable devices worn by the user; and calibrating one or more inference models by updating at least one parameter associated with the one or more inference models based, at least in part, on the plurality of neuromuscular signals and the at least one image.

[0061] In an aspect, the updating of the at least one parameter includes any one or any combination of: an updating of a routine for pre-processing the plurality of neuromuscular signals to calibrate the one or more inference models, an updating of a routine for pre-processing an image signal corresponding to the at least one image before the image signal is used to calibrate the one or more inference models, an updating of an architecture or architectures of the one or more inference models, an updating of a routine for post-processing an output of the one or more inference models, an updating of a selection routine for selecting the one or more inference models from a plurality of inference models, and an updating of a weight used for at least one of the one or more inference models.

[0062] In an aspect, the one or more inference models may be usable to generate a representation of any one or any combination of: a handstate, a still pose, and a dynamic gesture.

[0063] In an aspect, the at least one camera may comprise a camera disposed at a fixed location relative to the user.

[0064] In an aspect, the at least one camera may comprise a camera configured to be mounted on the user.

[0065] In an aspect, the at least one image may include any one or any combination of: an image produced by visible light, an image produced by infrared light, an image produced by light of a predetermined range of wavelengths, and an image produced by light of two or more different predetermined ranges of wavelengths.

[0066] In an aspect, the updating of the at least one parameter associated with the one or more inference models may comprise training the one or more inference models to produce an output that is determined based, at least in part, on the at least one image when the plurality of neuromuscular signals or information derived from the plurality of neuromuscular signals, or both the plurality of neuromuscular signals and the information derived from the plurality of neuromuscular signals, is or are provided as input to the one or more inference models.

[0067] In an aspect, the at least one computer may be in communication with an extended reality (XR) system that generates an XR environment and displays a visual representation. The method may further comprise: detecting, based at least in part on the at least one image, whether a gesture performed by the user matches a particular gesture; and, if the gesture performed by the user is detected to match the particular gesture, updating the one or more inference models to incorporate neuromuscular data corresponding to the gesture performed by the user.

[0068] In a variation of this aspect, the at least one camera may comprise a camera attached to a head-mounted-display of the XR system.

[0069] In another variation of this aspect, the method may further comprise instructing the XR system to provide a visible prompt for the user to perform the particular gesture.

[0070] In another variation of this aspect, the method may further comprise instructing the XR system to provide an audible prompt for the user to perform the particular gesture.

[0071] In another variation of this aspect, the gesture may be performed by the user upon an instruction from an external source.

[0072] In another variation of this aspect, the gesture performed by the user may be detected to match the particular gesture when a similarity between the gesture performed by the user and the particular gesture is above a predetermined threshold.

[0073] In another variation of this aspect, the visual representation may be displayed on a display screen viewable by the user, and may comprise a visual representation of a hand.

[0074] In another variation of this aspect, the visual representation may be displayed on a display screen viewable by the user, and may comprise instructions to the user.

[0075] In another variation of this aspect, the prompt for the user to perform the particular gesture may be a visual representation of the particular gesture provided on a display screen viewable by the user.

[0076] In a variation of some of the aspects, the visual representation may be within the XR environment.

[0077] In a variation of some of the aspects, the particular gesture may be one of: a motion or dynamic gesture, and a pose or static gesture.

[0078] In a variation of some of the aspects, the at least one image may be captured by the at least one camera during performance of the gesture by the user. The method may further comprise causing feedback to be provided to the user based, at least in part, on the at least one image captured during the performance of the gesture by the user.

[0079] In a variation of some of the aspects, the feedback to be provided to the user may indicate whether a body part of the user necessary for the particular gesture to be performed is fully included in the at least one image captured during the performance of the gesture by the user.

[0080] In a variation of some of the aspects, the feedback to be provided to the user may indicate whether a body part of the user necessary for the particular gesture to be performed is fully or partially occluded in the at least one image captured during the performance of the gesture by the user.

[0081] In a variation of some of the aspects, the feedback may be caused to be provided within the XR environment.

[0082] In a variation of some of the aspects, the feedback may comprise one or more of: audible feedback, visible feedback, haptic feedback, and electrical stimulation feedback. In a further variation, the visible feedback may be a moving image or a still image showing a way to perform the particular gesture.

[0083] In an aspect, the method may further comprise determining whether the at least one image includes position information for at least two segments of a musculoskeletal representation. The updating of the at least one parameter associated with the one or more inference models may comprise updating the one or more inference models to incorporate neuromuscular data corresponding to the plu-

rality of neuromuscular signals, when it is determined that the at least one image includes position information for at least two segments of the musculo skeletal representation.

[0084] In an aspect, the method may further comprise: detecting based at least in part on the at least one image, whether a gesture performed by the user to perform a task within an extended reality (XR) environment generated by an XR system matches a stored gesture for performing the task; and, if the gesture performed by the user matches the stored gesture for performing the task, updating, by the at least one processor, the one or more inference models to incorporate neuromuscular data corresponding to the gesture performed by the user.

[0085] In a variation of the aspect, the gesture performed by the user may be detected to match the stored gesture when a similarity between the gesture performed by the user and the stored gesture is above a predetermined threshold.

[0086] In a variation of the some of the aspects, the one or more inference models may be usable to generate a representation of any one or any combination of: a handstate, a still pose, and a dynamic gesture.

[0087] In an aspect, the method may further comprise determining, based on the at least one image, a user-specific skeletal geometry corresponding to the user. The updating of the at least one parameter associated with the one or more inference models may comprise updating the at least one parameter of the one or more inference models based on the determined user-specific skeletal geometry.

[0088] In a variation of this aspect, the user-specific skeletal geometry may comprise a length of at least one finger of the user.

[0089] In an aspect, the at least one camera may be configured to capture a plurality of images in a time series. The updating of the at least one parameter of the one or more inference models may be further based, at least in part, on the plurality of images captured in a time series.

[0090] According aspects of the technology described herein, a computerized system for providing dynamically-updated musculoskeletal information is provided. The system may comprise: at least one camera configured to capture at least one image; a plurality of neuromuscular sensors configured to sense and record a plurality of neuromuscular signals from a user, wherein the plurality of neuromuscular sensors are arranged on one or more wearable devices structured to be worn by the user to obtain the plurality of neuromuscular signals; and at least one computer processor. The at least one processor may be programmed to: provide, as an input to a trained inference model, information based on the plurality of neuromuscular signals and information based on the at least one image; determine, based on an output of the trained inference model, position information describing a spatial relationship between two or more connected musculoskeletal segments of the user, or force information describing a force exerted by at least one musculoskeletal segment of the user, or both the position information and the force information; and output the position information, or the force information, or both the position information and the force information.

[0091] In an aspect, the position information may be musculoskeletal position information that may include any one or any combination of: a handstate, a likelihood of a handstate, a still pose, a likelihood of a still pose, a dynamic gesture, and a likelihood of a dynamic gesture.

[0092] In an aspect, the at least one computer processor may be further programmed to: determine a first quality of the plurality of neuromuscular signals, or a second quality of the at least one image, or both the first quality and the second quality, and weight the plurality of neuromuscular signals, or the at least one image, or both the plurality of neuromuscular signals and the at least one image based on one or both of the first quality and the second quality, to determine the information based on the plurality of neuromuscular signals and the at least one image.

[0093] In a variation of this aspect, the first quality of the plurality of neuromuscular signals may be determined by determining whether at least one of the plurality of neuromuscular signals includes at least one signal artifact.

[0094] In another variation of this aspect, the second quality of the at least one image may be determined by determining whether a hand of the user is fully included in a field of view of the at least one camera, or whether the hand is fully occluded from the field of view of the at least one camera, or whether at least a portion of the hand of the user is occluded from the field of view of the at least one camera.

[0095] In another variation of this aspect, the at least one computer processor may be further programmed to determine the first quality in response to a determination that the second quality is greater than a threshold value.

[0096] In an aspect, the two or more connected musculoskeletal segments may include a forearm rigid segment connected via a wrist joint to a plurality of rigid segments for a hand. The plurality of neuromuscular sensors may include a plurality of electromyography (EMG) sensors. The at least one computer processor may be further programmed to: determine the position information for the plurality of rigid segments in the hand based, at least in part, on the plurality of neuromuscular signals output from the plurality of EMG sensors; and determine the position information for the forearm rigid segment based, at least in part, on the at least one image captured by the at least one camera.

[0097] In an aspect, a wrist of the user, a hand connected to the wrist, and fingers of the hand may be captured by the at least one image. The two or more connected musculoskeletal segments may include the wrist, the hand, and the fingers of the user. The at least one computer processor may be programmed to determine at least one of the position information and the force information for the plurality of rigid segments in the hand based, at least in part, on the plurality of neuromuscular signals and the at least one image captured by the at least one camera.

[0098] In a variation of this aspect, the at least one camera may capture the at least one image while the plurality of neuromuscular signals are sensed and recorded. The one or more wearable devices may be worn on a wrist of the user or a forearm of the user when the plurality of neuromuscular signals are sensed and recorded. The at least one computer processor may be programmed to determine any one or any combination of: a handstate, a dynamic gesture, and a still pose of the user based, at least in part, on the plurality of neuromuscular signals and the at least one image captured by the at least one camera.

[0099] In a variation of some of the aspects, the computerized system may further comprise at least one inertial measurement unit (IMU) sensor. The at least one computer processor may be further programmed to determine the position information for the forearm rigid segment further based, at least in part, on IMU signals output from the at

least one IMU sensor. In a further variation, the at least one computer processor may be programmed to: determine an initial position of the forearm rigid segment based, at least in part, on the IMU signals output from the at least one IMU sensor; and adjust the initial position information for the forearm rigid segment based, at least in part, on the at least one image captured by the at least one camera.

[0100] In an aspect, the one or more wearable devices may include at least one positional marker included thereon. The position information may be determined based, at least in part, on the at least one positional marker captured in the at least one image.

[0101] In an aspect, the trained inference model may be trained: to determine the position information based, at least in part, on the at least one image; and to determine the force information based on the plurality of neuromuscular signals.

[0102] In a variation of this aspect, the trained inference model may be trained to determine the position information based only on the at least one image.

[0103] In another variation of this aspect, the trained inference model may determine the position information by determining, based on the at least one image, that at least two of the connected musculoskeletal segments of the user are touching. The trained inference model may determine the force information by determining, based on the plurality of neuromuscular signals, a force applied between the at least two of the connected musculoskeletal segments of the user that are touching.

[0104] In various further variations, the at least two of the connected musculoskeletal segments of the user may comprise a thumb of the user and at least one other finger of the user. The trained inference model may determine, based on the at least one image, that the thumb of the user and the at least one other finger of the user are touching. The trained inference model may determine a force applied between the thumb of the user and the at least one other finger of the user. The trained inference model may further determine, based on the at least one image, that tips of the thumb of the user and the at least one other finger of the user are touching.

[0105] In an aspect, the at least one computer processor may be further programmed to: determine, based on the at least one image, a position of a hand of the user relative to a physical object or a virtual object; and determine the position information, or the force information, or both the position information and the force information based, at least in part, on the position of the hand of the user relative to the physical object or the virtual object.

[0106] In a variation of this aspect, the at least one computer processor may be programmed to: determine from the at least one image that the user is grasping the physical object or the virtual object; and determine the force information solely based on the plurality of neuromuscular signals.

[0107] In another variation of this aspect, the physical object may have a pliable surface. The at least one computer processor may be programmed to: determine the user is grasping the physical object based on a deformation in the pliable surface captured in the at least one image; and determine the force information based on one or both of: the plurality of neuromuscular signals and the deformation in the pliable surface captured in the at least one image.

[0108] In a variation of some of the aspects, the at least one image may include any one or any combination of: an image produced by visible light, an image produced by

infrared light, an image produced by light of a predetermined range of wavelengths, and an image produced by light of two or more different predetermined ranges of wavelengths.

[0109] In a variation of some of the aspects, the information based on the plurality of neuromuscular signals may comprise one or both of: the plurality of neuromuscular signals and information derived from the plurality of neuromuscular signals.

[0110] In a variation of some of the aspects, the information based on the at least one image may comprise one or both of: an image signal of the at least one image and information derived from the image signal of the at least one image.

[0111] In a variation of some of the aspects, the physical object may be a surface. The at least one computer processor may be further programmed to determine the force information based, at least in part, on a position of at least a portion of the hand relative to the surface.

[0112] In a variation of some of the aspects, wherein the virtual object may be a virtual object displayed in an extended reality (XR) environment generated by an XR system. The at least one computer processor may be further programmed to determine the position information based on whether the hand of the user is touching the virtual object in the XR environment.

[0113] According aspects of the technology described herein, a method of a computerized system for providing dynamically-updated musculoskeletal information is provided. The method may comprise: receiving, by at least one processor, at least one image captured by at least one camera; receiving, by the at least one processor, a plurality of neuromuscular signals sensed and recorded by a plurality of neuromuscular sensors configured to sense and record a plurality of neuromuscular signals from a user, wherein the plurality of neuromuscular sensors are arranged on one or more wearable devices structured to be worn by the user to obtain the plurality of neuromuscular signals; providing, by the at least one processor, as an input to a trained inference model, information based on the plurality of neuromuscular signals and information based on the at least one image; determining, by the at least one processor, based on an output of the trained inference model, position information describing a spatial relationship between two or more connected musculoskeletal segments of the user, or force information describing a force exerted by at least one musculoskeletal segment of the user, or both the position information and the force information; and outputting, by the at least one processor, the position information, or the force information, or both the position information and the force information.

[0114] In an aspect, the position information may be musculoskeletal position information that includes any one or any combination of: a handstate, a likelihood of a handstate, a still pose, a likelihood of a still pose, a dynamic gesture, and a likelihood of a dynamic gesture.

[0115] In an aspect, the method may further comprise: determining, by the at least one processor, a first quality of the plurality of neuromuscular signals, or a second quality of the at least one image, or both the first quality and the second quality; and weighting, by the at least one processor, the plurality of neuromuscular signals, or the at least one image, or both the plurality of neuromuscular signals and the at least one image based on one or both of the first quality and the

second quality, to determine the information based on the plurality of neuromuscular signals and the at least one image.

[0116] In a variation of this aspect, the first quality of the plurality of neuromuscular signals may be determined by determining whether at least one of the plurality of neuromuscular signals includes at least one signal artifact.

[0117] In another variation of this aspect, the second quality of the at least one image may be determined by determining whether a hand of the user is fully included in a field of view of the at least one camera, or whether the hand is fully occluded from the field of view of the at least one camera, or whether at least a portion of the hand of the user is occluded from the field of view of the at least one camera.

[0118] In another variation of this aspect, the method may further comprise determining, by the at least one processor, the first quality in response to a determination that the second quality is greater than a threshold value.

[0119] In an aspect, the two or more connected musculoskeletal segments may include a forearm rigid segment connected via a wrist joint to a plurality of rigid segments for a hand. The plurality of neuromuscular sensors may include a plurality of electromyography (EMG) sensors. The method may further comprise: determining, by the at least one processor, the position information for the plurality of rigid segments in the hand based, at least in part, on the plurality of neuromuscular signals output from the plurality of EMG sensors; and determining the position information for the forearm rigid segment based, at least in part, on the at least one image captured by the at least one camera.

[0120] In an aspect, a wrist of the user, a hand connected to the wrist, and fingers of the hand may be captured by the at least one image. The two or more connected musculoskeletal segments may include the wrist, the hand, and the fingers of the user. The method may further comprise determining, by the at least one processor, at least one of the position information and the force information for the plurality of rigid segments in the hand based, at least in part, on the plurality of neuromuscular signals and the at least one image captured by the at least one camera.

[0121] In a variation of this aspect, the at least one camera may capture the at least one image while the plurality of neuromuscular signals are sensed and recorded. The one or more wearable devices may be worn on a wrist of the user or a forearm of the user when the plurality of neuromuscular signals are sensed and recorded. The method may further comprise determining, by the at least one processor, any one or any combination of: a handstate, a dynamic gesture, and a still pose of the user based, at least in part, on the plurality of neuromuscular signals and the at least one image captured by the at least one camera.

[0122] In a variation of some of the aspect, the method may further comprise: receiving, by the at least one processor, inertial measurement unit (IMU) signals sensed and recorded by at least one IMU sensor; and determining, by the at least one processor, the position information for the forearm rigid segment further based, at least in part, on the IMU signals. In a further variation, the method may comprise: determining, by the at least one processor, an initial position of the forearm rigid segment based, at least in part, on the IMU signals; and adjusting the initial position information for the forearm rigid segment based, at least in part, on the at least one image captured by the at least one camera.

[0123] In an aspect, the one or more wearable devices may include at least one positional marker included thereon. The position information may be determined based, at least in part, on the at least one positional marker captured in the at least one image.

[0124] In an aspect, the trained inference model may be trained: to determine the position information based, at least in part, on the at least one image; and to determine the force information based on the plurality of neuromuscular signals.

[0125] In a variation of this aspect, the trained inference model may be trained to determine the position information based only on the at least one image.

[0126] In another variation of this aspect, the trained inference model may determine the position information by determining, based on the at least one image, that at least two of the connected musculoskeletal segments of the user are touching. The trained inference model may determine the force information by determining, based on the plurality of neuromuscular signals, a force applied between the at least two of the connected musculoskeletal segments of the user that are touching.

[0127] In a variation of some of the aspects, the at least two of the connected musculoskeletal segments of the user may comprise a thumb of the user and at least one other finger of the user. The trained inference model may determine, based on the at least one image, that the thumb of the user and the at least one other finger of the user are touching. The trained inference model may determine a force applied between the thumb of the user and the at least one other finger of the user. In a further variation, the trained inference model may determine, based on the at least one image, that tips of the thumb of the user and the at least one other finger of the user are touching.

[0128] In an aspect, the method may further comprise: determining, by the at least one processor based on the at least one image, a position of a hand of the user relative to a physical object or a virtual object; and determining, by the at least one processor, the position information, or the force information, or both the position information and the force information based, at least in part, on the position of the hand of the user relative to the physical object or the virtual object.

[0129] In a variation of this aspect, the method may further comprise: determining, by the at least one processor, based on the at least one image, that the user is grasping the physical object or the virtual object; and determining, by the at least one processor, the force information solely based on the plurality of neuromuscular signals.

[0130] In another variation of this aspect, the physical object may have a pliable surface. The method may further comprise: determining, by the at least one processor, whether the user is grasping the physical object based on a deformation in the pliable surface captured in the at least one image; and determining, by the at least one processor, the force information based on one or both of: the plurality of neuromuscular signals and the deformation in the pliable surface captured in the at least one image.

[0131] In an aspect, the at least one image may include any one or any combination of: an image produced by visible light, an image produced by infrared light, an image produced by light of a predetermined range of wavelengths, and an image produced by light of two or more different predetermined ranges of wavelengths.

[0132] In an aspect, the information based on the plurality of neuromuscular signals may comprise one or both of: the plurality of neuromuscular signals and information derived from the plurality of neuromuscular signals.

[0133] In an aspect, the information based on the at least one image may comprise one or both of: an image signal of the at least one image and information derived from the image signal of the at least one image.

[0134] In a variation of some of the aspects, the physical object may be a surface. The method may further comprise: determining, by the at least one processor, the force information based, at least in part, on a position of at least a portion of the hand relative to the surface.

[0135] In a variation of some of the aspects, the virtual object may be a virtual object displayed in an extended reality (XR) environment generated by an XR system. The method may further comprise determining, by the at least one processor, the position information based on whether the hand of the user is touching the virtual object in the XR environment.

[0136] According aspects of the technology described herein, a non-transitory computer-readable storage medium storing code that, when executed by at least one computer, causes the at least one computer to perform a method for providing dynamically-updated musculoskeletal information is provided. The method may comprise: receiving at least one image captured by at least one camera; receiving a plurality of neuromuscular signals sensed and recorded by a plurality of neuromuscular sensors configured to sense and record a plurality of neuromuscular signals from a user, wherein the plurality of neuromuscular sensors are arranged on one or more wearable devices structured to be worn by the user to obtain the plurality of neuromuscular signals; providing, as an input to a trained inference model, information based on the plurality of neuromuscular signals and information based on the at least one image; determining, based on an output of the trained inference model, position information describing a spatial relationship between two or more connected musculoskeletal segments of the user, or force information describing a force exerted by at least one musculoskeletal segment of the user, or both the position information and the force information; and outputting the position information, or the force information, or both the position information and the force information.

[0137] In an aspect, the position information may be musculoskeletal position information that may include any one or any combination of: a handstate, a likelihood of a handstate, a still pose, a likelihood of a still pose, a dynamic gesture, and a likelihood of a dynamic gesture.

[0138] In an aspect, the method may further comprise: determining a first quality of the plurality of neuromuscular signals, or a second quality of the at least one image, or both the first quality and the second quality; and weighting the plurality of neuromuscular signals, or the at least one image, or both the plurality of neuromuscular signals and the at least one image based on one or both of the first quality and the second quality, to determine the information based on the plurality of neuromuscular signals and the at least one image.

[0139] In a variation of this aspect, the first quality of the plurality of neuromuscular signals may be determined by determining whether at least one of the plurality of neuromuscular signals includes at least one signal artifact.

[0140] In another variation of this aspect, the second quality of the at least one image may be determined by determining whether a hand of the user is fully included in a field of view of the at least one camera, or whether the hand is fully occluded from the field of view of the at least one camera, or whether at least a portion of the hand of the user is occluded from the field of view of the at least one camera.

[0141] In another variation of this aspect, the method may further comprise determining the first quality in response to a determination that the second quality is greater than a threshold value.

[0142] In an aspect, the two or more connected musculoskeletal segments may include a forearm rigid segment connected via a wrist joint to a plurality of rigid segments for a hand. The plurality of neuromuscular sensors may include a plurality of electromyography (EMG) sensors. The method may further comprise: determining the position information for the plurality of rigid segments in the hand based, at least in part, on the plurality of neuromuscular signals output from the plurality of EMG sensors; and determining the position information for the forearm rigid segment based, at least in part, on the at least one image captured by the at least one camera.

[0143] In an aspect, a wrist of the user, a hand connected to the wrist, and fingers of the hand may be captured by the at least one image. The two or more connected musculoskeletal segments may include the wrist, the hand, and the fingers of the user. The method may further comprise determining at least one of the position information and the force information for the plurality of rigid segments in the hand based, at least in part, on the plurality of neuromuscular signals and the at least one image captured by the at least one camera.

[0144] In a variation of this aspect, the at least one camera may capture the at least one image while the plurality of neuromuscular signals are sensed and recorded. The one or more wearable devices may be worn on a wrist of the user or a forearm of the user when the plurality of neuromuscular signals are sensed and recorded. The method may further comprise determining any one or any combination of: a handstate, a dynamic gesture, and a still pose of the user based, at least in part, on the plurality of neuromuscular signals and the at least one image captured by the at least one camera.

[0145] In a variation of some of the aspects, the method may further comprise: receiving inertial measurement unit (IMU) signals sensed and recorded by at least one IMU sensor; and determining the position information for the forearm rigid segment further based, at least in part, on the IMU signals. In a further variation, the method may comprise: determining an initial position of the forearm rigid segment based, at least in part, on the IMU signals; and adjusting the initial position information for the forearm rigid segment based, at least in part, on the at least one image captured by the at least one camera.

[0146] In an aspect, the one or more wearable devices may include at least one positional marker included thereon. The position information may be determined based, at least in part, on the at least one positional marker captured in the at least one image.

[0147] In an aspect, the trained inference model may be trained: to determine the position information based, at least in part, on the at least one image; and to determine the force information based on the plurality of neuromuscular signals.

[0148] In a variation of this aspect, the trained inference model may be trained to determine the position information based only on the at least one image.

[0149] In another variation of this aspect, the trained inference model may determine the position information by determining, based on the at least one image, that at least two of the connected musculoskeletal segments of the user are touching. The trained inference model may determine the force information by determining, based on the plurality of neuromuscular signals, a force applied between the at least two of the connected musculoskeletal segments of the user that are touching.

[0150] In a variation of some of the aspects, the at least two of the connected musculoskeletal segments of the user may comprise a thumb of the user and at least one other finger of the user. The trained inference model may determine, based on the at least one image, that the thumb of the user and the at least one other finger of the user are touching. The trained inference model may determine a force applied between the thumb of the user and the at least one other finger of the user. In a further variation, the trained inference model may determine, based on the at least one image, that tips of the thumb of the user and the at least one other finger of the user are touching.

[0151] In an aspect, the method may further comprise: determining, based on the at least one image, a position of a hand of the user relative to a physical object or a virtual object; and determining the position information, or the force information, or both the position information and the force information based, at least in part, on the position of the hand of the user relative to the physical object or the virtual object.

[0152] In a variation of this aspect, the method may further comprise: determining, based on the at least one image, that the user is grasping the physical object or the virtual object; and determining the force information solely based on the plurality of neuromuscular signals.

[0153] In another variation of this aspect, the physical object may have a pliable surface. The method may further comprise: determining whether the user is grasping the physical object based on a deformation in the pliable surface captured in the at least one image; and determining the force information based on one or both of: the plurality of neuromuscular signals and the deformation in the pliable surface captured in the at least one image.

[0154] In an aspect, the at least one image may include any one or any combination of: an image produced by visible light, an image produced by infrared light, an image produced by light of a predetermined range of wavelengths, and an image produced by light of two or more different predetermined ranges of wavelengths.

[0155] In an aspect, the information based on the plurality of neuromuscular signals may comprise one or both of: the plurality of neuromuscular signals and information derived from the plurality of neuromuscular signals.

[0156] In an aspect, the information based on the at least one image may comprise one or both of: an image signal of the at least one image and information derived from the image signal of the at least one image.

[0157] In a variation of some of the aspects, the physical object may be a surface. The method may further comprise determining the force information based, at least in part, on a position of at least a portion of the hand relative to the surface.

[0158] In a variation of some of the aspects, the virtual object may be a virtual object displayed in an extended reality (XR) environment generated by an XR system. The method may further comprise determining the position information based on whether the hand of the user is touching the virtual object in the XR environment.

[0159] According to aspects of the technology described herein, a computerized system for training an inference model used to generate a musculoskeletal representation is provided. The system may comprise a plurality of neuromuscular sensors configured to sense and record a plurality of neuromuscular signals from a user, wherein the plurality of neuromuscular sensors are arranged on one or more wearable devices structured to be worn by the user to obtain the plurality of neuromuscular signals; at least one camera configured to capture at least one image while the plurality of neuromuscular signals are sensed and recorded; and at least one computer processor. The at least one computer processor may be programmed to: determine, based on information obtained from the at least one image, position information describing a spatial relationship between two or more connected musculoskeletal segments of the user; determine, based on the plurality of neuromuscular signals, force information; associate the position information with the force information; train an inference model to output a musculoskeletal representation consistent with the position information, or the force information, or the position information and the force information when neuromuscular input signals provided to the inference model have at least one predetermined characteristic, to produce an updated inference model; and cause the updated inference model to be stored in a memory.

[0160] In an aspect, the at least one computer processor may determine the position information based on the information obtained from the at least one image and information obtained from the plurality of neuromuscular signals.

[0161] In an aspect, the at least one image may comprise a series of time-sequenced images forming a video.

[0162] In an aspect, the musculoskeletal representation may correspond to any one or any combination of: a handstate of the user, a still pose of the user, and a dynamic gesture of the user.

[0163] In an aspect, the at least one image may include any one or any combination of: an image produced by visible light, an image produced by infrared light, an image produced by light of a predetermined range of wavelengths, and an image produced by light of two or more different predetermined ranges of wavelengths.

[0164] In an aspect, the at least one computer processor may be programmed to: determine a first quality of the plurality of neuromuscular signals, or a second quality of the at least one image, or both the first quality and the second quality; and weight the plurality of neuromuscular signals, or the at least one image, or both the plurality of neuromuscular signals and the at least one image based on one or both of the first quality and the second quality, to train the inference model.

[0165] In a variation of this aspect, the first quality of the plurality of neuromuscular signals may be determined by determining whether at least one of the plurality of neuromuscular signals includes at least one signal artifact.

[0166] In another variation of this aspect, the second quality of the at least one image may be determined by determining whether a hand of the user is fully included in

a field of view of the at least one camera, or whether the hand is fully occluded from the field of view of the at least one camera, or whether at least a portion of the hand of the user is occluded from the field of view of the at least one camera.

[0167] In another variation of this aspect, the at least one computer processor may be further programmed to determine the first quality in response to a determination that the second quality is greater than a threshold value.

[0168] In an aspect, the at least one predetermined characteristic may comprise any one or any combination of: a presence of an artifact of the neuromuscular input signals, a presence of a plurality of artifacts of the neuromuscular input signals, a relative position of a plurality of artifacts of the neuromuscular input signals, a range of amplitudes of the neuromuscular input signals, and a periodic frequency of artifacts of the neuromuscular input signals.

[0169] In an aspect, the computerized system may further comprise at least one inertial measurement unit (IMU) sensor. The at least one computer processor may be further programmed to determine the position information based, at least in part, on IMU signals output from the at least one IMU sensor.

[0170] In an aspect, the one or more wearable devices may include at least one positional marker included thereon. The position information may be determined based, at least in part, on the at least one positional marker captured in the at least one image.

[0171] According aspects of the technology described herein, method of a computerized system for training an inference model used to generate a musculoskeletal representation is provided. The method may comprise: receiving, by at least one processor, a plurality of neuromuscular signals of a user, the plurality of neuromuscular signals being sensed and recorded by a plurality of neuromuscular sensors arranged on one or more wearable devices worn by the user; receiving, by the at least one processor, at least one image captured by at least one camera while the plurality of neuromuscular signals are sensed and recorded; determining, by the at least one processor based on information obtained from the at least one image, position information describing a spatial relationship between two or more connected musculoskeletal segments of the user; determining, by the at least one processor based on the plurality of neuromuscular signals, force information; associating, by the at least one processor, the position information with the force information; training, by the at least one processor, an inference model to output a musculoskeletal representation consistent with the position information, or the force information, or the position information and the force information when neuromuscular input signals provided to the inference model have at least one predetermined characteristic, to produce an updated inference model; and causing, by the at least one processor, the updated inference model to be stored in a memory.

[0172] In an aspect, the method may further comprise determining, by the at least one processor, the position information based on the information obtained from the at least one image and information obtained from the plurality of neuromuscular signals.

[0173] In an aspect, the at least one image may comprise a series of time-sequenced images forming a video.

[0174] In an aspect, the musculoskeletal representation may correspond to any one or any combination of: a handstate of the user, a still pose of the user, and a dynamic gesture of the user.

[0175] In an aspect, the at least one image may include any one or any combination of: an image produced by visible light, an image produced by infrared light, an image produced by light of a predetermined range of wavelengths, and an image produced by light of two or more different predetermined ranges of wavelengths.

[0176] In an aspect, the method may further comprise: determining, by the at least one processor, a first quality of the plurality of neuromuscular signals, or a second quality of the at least one image, or both the first quality and the second quality; and weighting, by the at least one processor, the plurality of neuromuscular signals, or the at least one image, or both the plurality of neuromuscular signals and the at least one image based on one or both of the first quality and the second quality, to train the inference model.

[0177] In a variation of this aspect, the first quality of the plurality of neuromuscular signals may be determined by determining whether at least one of the plurality of neuromuscular signals includes at least one signal artifact.

[0178] In another variation of this aspect, the second quality of the at least one image may be determined by determining whether a hand of the user is fully included in a field of view of the at least one camera, or whether the hand is fully occluded from the field of view of the at least one camera, or whether at least a portion of the hand of the user is occluded from the field of view of the at least one camera.

[0179] In another variation of this aspect, the method may further comprise determining, by the at least one processor, the first quality in response to a determination that the second quality is greater than a threshold value.

[0180] In an aspect, the at least one predetermined characteristic may comprise any one or any combination of: a presence of an artifact of the neuromuscular input signals, a presence of a plurality of artifacts of the neuromuscular input signals, a relative position of a plurality of artifacts of the neuromuscular input signals, a range of amplitudes of the neuromuscular input signals, and a periodic frequency of artifacts of the neuromuscular input signals.

[0181] In an aspect, the method may further comprise: receiving, by the at least one processor, inertial measurement unit (IMU) signals output from at least one IMU sensor; and determining, by the at least one processor, the position information based, at least in part, on the IMU signals output from the at least one IMU sensor.

[0182] In an aspect, the one or more wearable devices may include at least one positional marker included thereon. The position information may be determined based, at least in part, on the at least one positional marker captured in the at least one image.

[0183] According aspects of the technology described herein, a non-transitory computer-readable storage medium storing code that, when executed by at least one computer, causes the at least one computer to perform a method for training an inference model used to generate a musculoskeletal representation is provided. The method may comprise: receiving a plurality of neuromuscular signals of a user, the plurality of neuromuscular signals being sensed and recorded by a plurality of neuromuscular sensors arranged on one or more wearable devices worn by the user; receiving at least one image captured by at least one camera while the

plurality of neuromuscular signals are sensed and recorded; determining, based on information obtained from the at least one image, position information describing a spatial relationship between two or more connected musculoskeletal segments of the user; determining, based on the plurality of neuromuscular signals, force information; associating the position information with the force information; training an inference model to output a musculoskeletal representation consistent with the position information, or the force information, or the position information and the force information when neuromuscular input signals provided to the inference model have at least one predetermined characteristic, to produce an updated inference model; and causing the updated inference model to be stored in a memory.

[0184] In an aspect, the method may further comprise determining the position information based on the information obtained from the at least one image and information obtained from the plurality of neuromuscular signals.

[0185] In an aspect, the at least one image may comprise a series of time-sequenced images forming a video.

[0186] In an aspect, the musculoskeletal representation may correspond to any one or any combination of: a handstate of the user, a still pose of the user, and a dynamic gesture of the user.

[0187] In an aspect, the at least one image may include any one or any combination of: an image produced by visible light, an image produced by infrared light, an image produced by light of a predetermined range of wavelengths, and an image produced by light of two or more different predetermined ranges of wavelengths.

[0188] In an aspect, the method may further comprise: determining a first quality of the plurality of neuromuscular signals, or a second quality of the at least one image, or both the first quality and the second quality; and weighting the plurality of neuromuscular signals, or the at least one image, or both the plurality of neuromuscular signals and the at least one image based on one or both of the first quality and the second quality, to train the inference model.

[0189] In a variation of this aspect, the first quality of the plurality of neuromuscular signals may be determined by determining whether at least one of the plurality of neuromuscular signals includes at least one signal artifact.

[0190] In another variation of this aspect, the second quality of the at least one image may be determined by determining whether a hand of the user is fully included in a field of view of the at least one camera, or whether the hand is fully occluded from the field of view of the at least one camera, or whether at least a portion of the hand of the user is occluded from the field of view of the at least one camera.

[0191] In another variation of this aspect, the method may further comprise determining the first quality in response to a determination that the second quality is greater than a threshold value.

[0192] In an aspect, the at least one predetermined characteristic may comprise any one or any combination of: a presence of an artifact of the neuromuscular input signals, a presence of a plurality of artifacts of the neuromuscular input signals, a relative position of a plurality of artifacts of the neuromuscular input signals, a range of amplitudes of the neuromuscular input signals, and a periodic frequency of artifacts of the neuromuscular input signals.

[0193] In an aspect, the method may further comprise: receiving inertial measurement unit (IMU) signals output from at least one IMU sensor; and determining the position

information based, at least in part, on the IMU signals output from the at least one IMU sensor.

[0194] In an aspect, the one or more wearable devices may include at least one positional marker included thereon. The position information may be determined based, at least in part, on the at least one positional marker captured in the at least one image.

[0195] It should be appreciated that all combinations of the foregoing concepts and additional concepts discussed in greater detail below (provided such concepts are not mutually inconsistent) are contemplated as being part of the inventive subject matter disclosed herein. In particular, all combinations of claimed subject matter appearing at the end of this disclosure are contemplated as being part of the inventive subject matter disclosed herein.

BRIEF DESCRIPTION OF DRAWINGS

[0196] Various non-limiting embodiments of the technology will be described with reference to the following figures. It should be appreciated that the figures are not necessarily drawn to scale.

[0197] FIG. 1 is a schematic diagram of a computer-based system for processing sensor data, such as signals obtained from neuromuscular sensors, in accordance with some embodiments of the technology described herein;

[0198] FIG. 2 is a schematic diagram of a distributed computer-based system that integrates an augmented reality (AR) system with a neuromuscular activity system, in accordance with some embodiments of the technology described herein;

[0199] FIG. 3 is a flowchart of a process for processing neuromuscular signals and camera data in accordance with some embodiments of the technology described herein;

[0200] FIG. 4 is a flowchart of a process for processing gesture information in an XR system in accordance with some embodiments of the technology described herein;

[0201] FIG. 5 is a flowchart of a process for integrating neuromuscular signals and camera data and providing feedback to a user in accordance with some embodiments of the technology described herein;

[0202] FIG. 6 is a flowchart of a process for updating an inference model based on camera data in accordance with some embodiments of the technology described herein;

[0203] FIG. 7 is a flowchart of a process for updating an inference model based on camera data in accordance with some embodiments of the technology described herein;

[0204] FIG. 8 is a flowchart of a process for determining position information and force information, in accordance with some embodiments of the technology described herein;

[0205] FIG. 9 is a flowchart of a process for determining qualities of input signals and performing model functions based on those qualities, in accordance with some embodiments of the technology described herein;

[0206] FIG. 10 is a diagram showing an example implementation using EMG sensors and a camera, in accordance with some embodiments of the technology described herein;

[0207] FIG. 11 is a diagram showing a trained inference model with representative inputs and outputs, in accordance with some embodiments of the technology described herein;

[0208] FIG. 12 illustrates a wristband having EMG sensors arranged circumferentially thereon, in accordance with some embodiments of the technology described herein;

[0209] FIG. 13A illustrates a wearable system with sixteen EMG sensors arranged circumferentially around a band

configured to be worn around a user's lower arm or wrist, in accordance with some embodiments of the technology described herein;

[0210] FIG. 13B is a cross-sectional view through one of the sixteen EMG sensors illustrated in FIG. 13A;

[0211] FIGS. 14A and 14B schematically illustrate components of a computer-based system in which some embodiments of the technology described herein are implemented. FIG. 14A illustrates a wearable portion of the computer-based system, and FIG. 14B illustrates a dongle portion connected to a computer, wherein the dongle portion is configured to communicate with the wearable portion.

DETAILED DESCRIPTION

[0212] Some embodiments are directed to coupling a system that senses neuromuscular signals via neuromuscular sensors with a system that performs extended reality (XR) functions. As will be appreciated, XR functions may include augmented reality (AR) functions, virtual reality (VR) functions, mixed reality (MR) functions, and the like. In particular, a system that senses neuromuscular signals for the purpose of determining a position of a body part (e.g., a hand, an arm, etc.) may be used in conjunction with an XR system to provide an improved XR experience for a user. For instance, information gained within both systems may be used to improve the overall XR experience. In embodiments where a musculoskeletal representation associated with the body part is generated based on sensor data, a camera in an XR system may capture data that is used to improve the accuracy of a model of the musculoskeletal representation and/or may be used to calibrate the model. Further, in another implementation, sensor data may provide muscle activation information, which may be visualized and displayed to a user in an XR environment. In yet another implementation, display information in the XR environment may be used as feedback to the user to permit the user to more accurately control his/her musculoskeletal input (e.g., movement input) to the system. Further, control features may be provided that permit neuromuscular signals to control XR system elements, including operation of the XR system itself.

[0213] The inventors recognize that neither cameras nor neuromuscular sensors are by themselves ideal input systems. Cameras such as those provided by an XR system may provide good positional information (relative both to other skeletal segments and to external objects) when joint segments are clearly within view, but may be limited by field of view restrictions and occlusion, and may be ill-suited for measuring forces. At the same time, measurements of signals by neuromuscular sensors (e.g., electromyography (EMG) signals or another modality of neuromuscular signal as described herein) may, on their own, be insufficient for distinguishing between forces that a user is applying against himself/herself versus forces that he/she applies to an external object, and such signals may not provide sufficiently accurate information about skeletal geometry, for example finger lengths. According to some embodiments, it is appreciated that it would be beneficial to increase the accuracy of XR systems and neuromuscular-sensor-based systems to provide more accurate and more realistic user experiences.

[0214] As will be appreciated, an inference model may involve a generalized skeletal geometry for a type of user (e.g., a typical adult male, a typical child, a typical adult

female) or may involve a user-specific skeletal geometry for a particular user (e.g., Jane Miller, Henry Smith).

[0215] According to some embodiments, a system is configured to use camera information to improve interpretation of neuromuscular signals and their relationship to movement and force generation. For example, inside-out cameras and/or other camera systems associated with XR systems may be used in association with neuromuscular signals to more accurately represent the position of segments of a user's body, movement of a user's body, and/or representations of force exerted by segments of a user's body. For example, camera information, such as images, video, time series of images, etc., may be used to calibrate neuromuscular systems by providing ground truth labels for data from neuromuscular signals. In one implementation, a system may perform a calibration operation using prompts (e.g., hand gestures, verbal information provided visually or audibly (e.g., words (e.g., "fist") or phrases (e.g., "make a thumbs up gesture"))) provided to a user through an XR display or other screen (which may include a smartphone, smartwatch, tablet computer, laptop, desktop computer, AR display, VR display, etc.), where the user is asked to match his/her hand posture to that of a projected hand on the XR display, with the camera optionally assisting with detection of when a match occurs. Further, other types of camera data may be used to calibrate a neuromuscular system, such as calibrating a geometric model of skeletal geometry using camera data. For instance, finger lengths of a geometric model may be verified and/or corrected using camera data. In another example, camera information may be used to determine that two or more parts of a user's body are touching (e.g. the tips of the thumb and index finger touching in a 'pinch' pose). In such a case, images may be processed to automatically identify one or more parts of a user's body (i.e. hand, fingers, and/or wrist) and determine the relative position of the parts of a user's body. For example, the processed camera information may be translated to quaternions or another mathematical framework for representing the segments and joints of a part of a user's body.

[0216] In some embodiments, neuromuscular signals, camera data, or both may be used to provide a real-time determination of musculoskeletal representations or gestures. For instance, as neuromuscular (e.g., EMG) and IMU information may be used to determine a more accurate musculoskeletal representation, other data such as camera data may be used to create an even more accurate and consistent representation. Further, it is appreciated that multiple signals can be used, including, but not limited to, signals from one or more cameras, neuromuscular signals from one or more devices, among other types of sensor data, to determine real-time musculoskeletal representations. Other data, such as IMU data and/or camera data, may be used to train and improve an inference model for musculoskeletal representation as well as improve the real-time representation of skeletal position. As will be appreciated, an inference model may be a model that utilizes statistical inference based on a probability distribution to deduce a result; in this regard, an inference model may comprise a statistical model.

[0217] As used herein, the term "gestures" refers to a static or dynamic configuration of one or more body parts including the position of the one or more body parts and forces associated with the configuration. For example, gestures include discrete gestures, such as placing or pressing the

palm of a hand down on a solid surface or grasping a ball, continuous gestures, such as waving a finger back and forth, grasping and throwing a ball, or a combination of discrete and continuous gestures. Gestures may include covert gestures that are imperceptible to another person, such as slightly tensing a joint by co-contracting opposing muscles or using sub-muscular activations. Gestures may be defined by an application configured to prompt a user to perform the gestures or, alternatively, gestures may be arbitrarily defined by a user. The gestures performed by the user may include symbolic gestures (e.g., gestures mapped to other gestures, interactions, or commands, for example, based on a gesture vocabulary that specifies the mapping). In some cases, hand and arm gestures may be symbolic and used to communicate according to cultural standards.

[0218] In some embodiments, a system may combine neuromuscular signals and camera data (e.g., camera signals or image signals from a camera) to infer or reconstruct the position of segments (i.e., skeletal segments) of the user's body. The system may be adapted to adjust a level of influence given to each signal based on the quality of that signal. In the case of a camera signal (e.g., an image signal from a camera), there may be field of view or occlusion restrictions that cause the signal to be unreliable or inaccurate. In the case of neuromuscular signals such as a signal from an EMG sensor, there may be EMG artifacts produced that cause the EMG signal to be unreliable (or other artifacts present in neuromuscular signals derived from an alternative modality of neuromuscular signal, as described below). In such cases, the system may be configured to assess a quality level of each of the signals to determine whether either or both should be used (e.g., to determine a handstate or gesture). The system may also use weighting or another combining method to adjust the degree to which a signal is used between the different types of signal sources (e.g., the different sensor types). Also, when confidence in a particular signal source is high, the signal from that source may be used to train and/or correct another source or model. For instance, a quality of the neuromuscular signals may be assessed in cases where the hand is in clear view of the camera, and a retraining of the handstate model may be performed.

[0219] In some embodiments, a system may include a first inference model for generating a musculoskeletal representation based on neuromuscular signals and a second model for generating a musculoskeletal representation based on camera input (i.e., input obtained from a camera). The system may be configured to transition between the use of the first inference model and the use of the second inference model for representing a user's handstate based, at least in part, on information associated with the camera input, such as whether all or a portion of the user's hand is within the camera's field of view. For example, when the user's hand (or another portion of the user's body) is within the camera's field of view, the second inference model may be used to determine the position of the segments of the user's hand, whereas when the user's hand is completely or partially outside the camera's field of view (including cases wherein the user's body is completely or partially occluded), the first inference model may be used to determine the position of the segments of the user's hand representing the handstate. As an alternative implementation, a single inference model that receives both neuromuscular signals and camera input may be used, and the contribution of the inputs may be weighted,

as described above. In instances in which the user's hand is out of the field of view of the camera, the camera-input weight may be set to zero or some other small value to reflect the unreliability of the camera input for estimating position information when the user's hand is out of the camera's field of view.

[0220] In some embodiments, data from one or more cameras may be used to determine the position of an arm, a hand, a forearm, or another part of the user's body. Also, camera data may be used to combat drift in an IMU-based estimate of forearm position, with the IMU information being used to measure forearm orientation and the neuromuscular signals being used to determine hand and wrist configuration and forces. In this embodiment, positional tracking reference marks on a band of neuromuscular sensors (e.g., a band of EMG sensors) may be used, especially when the camera is used to refine the IMU-based system for tracking one or more position(s) of articulated rigid bodies. As will be appreciated, data from a single camera may be used or data from two or more cameras may be used.

[0221] According to some embodiments, camera data may be used for determining whether an object (e.g., a hand, finger, or another physical object) is subjected to a force. For instance, camera data may be used to distinguish between whether someone is moving freely or pressing against object(s) and/or surface(s) (which may include another part of the user's body), determine which object(s) and/or surfaces is or are being interacted with, which position(s) on the surface(s) and/or object(s) are being touched, and can assist with estimating a skeletal configuration, position, and/or force. It is appreciated that although camera data can be used to determine whether a force is being applied, camera data is not particularly suited to determining a magnitude of the force(s) applied. To this end, other input signals (e.g., neuromuscular signals) can be used to determine an amount of force applied and also assist with determining the skeletal configuration and/or position.

[0222] When an XR user is applying self force(s) against the user's own body (e.g., the user is pinching his or her own arm), the camera can assist with determining the skeletal configuration and/or position (e.g. which joint segments are involved in touching the arm), and the neuromuscular sensors can be used to determine the intensity of the force. In this way, a more accurate representation of the arm position (s) and/or the hand position(s) and force(s) can be constructed. Further, it is more generally appreciated that the physical context as determined by the cameras could inform priors for estimating the musculoskeletal representation from neuromuscular signals, for example if users are more likely to perform certain movements in specific environments.

[0223] Some embodiments are directed to predicting information about the position and movements of portions of a user's arm and/or hand, which may be represented as a multi-segment articulated rigid body system with joints connecting the multiple segments of the rigid body system. For example, in the case of a hand movement, signals sensed and recorded by wearable neuromuscular sensors placed at locations on the user's body (e.g., the user's arm and/or wrist) may be provided as input to an inference model trained to predict estimates of the position (e.g., absolute position, relative position, orientation) and the force(s) associated with a plurality of rigid segments in a computer-based musculoskeletal representation associated with a hand when

the user performs one or more hand movements. The combination of position information and force information associated with segments of a musculoskeletal representation associated with a hand is colloquially referred to herein as a “handstate” of the musculoskeletal representation. As a user performs different movements, a trained inference model interprets neuromuscular signals recorded by the wearable neuromuscular sensors into position and force estimates (handstate information) that are used to update the musculoskeletal representation. Because the neuromuscular signals are continuously sensed and recorded, the musculoskeletal representation is updated in real time and a visual representation of a hand (e.g., within an extended reality environment) may be rendered based on the current handstate estimates. As will be appreciated, an estimate of a user’s handstate may be used to determine a gesture being performed by the user and/or to predict a gesture that the user will perform.

[0224] According to some embodiments, musculoskeletal representations (e.g., hand-rendering) may include actual visual representations of biomimetic (realistic) hands, synthetic (robotic) hands, as well as abstract “internal representations” that serve as input for gesture control (e.g., to other applications, systems, etc.). That is, the hand’s position and/or force may be provided to downstream algorithms (e.g., control algorithms in an XR system) but may not be directly rendered. In some embodiments, camera data may be used to assist in creating actual visual representations (e.g., improving an XR version of a user’s hands based on the camera data).

[0225] According to some embodiments, information received in an EMG context (e.g., force) may be used to control modes of a computer system (e.g., an XR system). For instance, a detection of force (e.g., beyond a threshold amount, such as applied by a finger or a hand) may be used to control the XR system, such as to, for example, open a help dialog interface, open a global system properties menu, or perform a mode switching function. Such dual-mode inputs may also include, for example, submitting a “wake” signal to an XR system or other type of system having heavy computing costs (e.g., high usage of electrical power) because the system wakes or is responsive to the user engaging the system with a position-, force-, or gesture-related event (e.g., the user clenching his or her first results in a wake signal being sent to the XR system). In another implementation, any number of combinations of neuromuscular signals may be used in conjunction with camera data to control a computer system. For instance, in a scenario where a user is selecting an object in an XR environment, the camera data may be used to determine the user’s posture, and an EMG signal may be used to inform the XR system that the object has been selected for action (e.g., by detecting a force applied through the EMG signal). It should be appreciated that any combination of modes of control using different types of input data may be performed.

[0226] Some embodiments are directed to coupling a system that senses neuromuscular signals with a system that performs XR functions (e.g., AR functions, VR functions, etc.). In particular, a system that senses neuromuscular signals for the purpose of determining a position of a body part (e.g., a hand, an arm, etc.) may be used in conjunction with an XR system to provide an improved XR experience for a user. For instance, information gained within both systems may be used to improve the overall XR experience.

In one example, a camera in an AR system may capture data that is used to improve the accuracy of a model of a musculoskeletal representation, used to calibrate the model and/or to control the system(s) using any combination of camera and neuromuscular signals. Further, in another implementation, muscle activation data may be visualized and displayed to a user in an XR environment. In yet another example, display information in the XR environment may be used as feedback to the user to permit the user to more accurately control their musculoskeletal input to the system. Further, control features may be provided that permit neuromuscular signals to control XR system elements.

[0227] As discussed above, according to some embodiments of the present technology, a system is configured to use camera information to improve interpretation of neuromuscular signals and their relationship to movement, position, and force generation. As will be appreciated, the camera information may be, for example, an image signal corresponding to at least one image captured by a camera; thus, as used herein, an image from a camera may be understood to refer to an image signal from a camera. The camera may be a still camera, a video camera, an infrared camera, and the like, which is able to capture or record an image of a user. One or more filters may be used on the camera, so that the camera may capture images only within a particular range of wavelengths of light. As will be appreciated, the image may be a still image, a sequence of still images (or image sequence), a moving image (or video sequence), and the like, which may be captured and recorded as a signal. The terms “camera information,” “camera data,” and “camera signal,” may be used herein to represent information about the user that may be captured by a camera. It should be understood that although various embodiments may refer to “a” camera or “the” camera, such embodiments may utilize two or more cameras instead of one camera. Further, the camera information may relate to any one or any combination of: an image produced by visible light, an image produced by non-visible (e.g., infrared) light, an image produced by light of a predetermined range of wavelengths, and an image produced by light of two or more different predetermined ranges of wavelengths. For example, non-visible light may be used to capture an image that shows heat distribution in the user’s body, which may provide an indication of blood flow within the user, which in turn may be used to infer a condition of the user (e.g., a force being exerted by a finger of the user may have a different blood-flow pattern than a finger that is not exerting force).

[0228] A camera may be mounted on the user (e.g., on a head-mounted display worn by the user, or on a glove worn on the user’s hand) or may be mounted external to the user to capture the user and/or the user’s environment. When a camera is mounted on the user, the camera may be used to capture the user’s environment and/or portions of the user’s body (e.g., a hand-mounted camera may be used to capture an image of the user’s other hand).

[0229] Signals sensed and recorded by wearable sensors placed at locations on the user’s body may be provided as input to an inference model trained to generate spatial information for rigid segments of a multi-segment articulated rigid body model of a human body. The spatial information may include, for example, position information of one or more segments, orientation information of one or more segments, joint angles between segments, and the like.

As a result of the training, the inference model may implicitly represent inferred motion of the articulated rigid body under defined movement constraints. The trained inference model may output data useable for applications such as rendering a representation of the user's body in a virtual environment, in which the user interacts with physical or virtual objects, and monitoring the user's movements as the user performs a physical activity to assess, for example, whether the user is performing the physical activity in a desired manner. As will be appreciated, the output data from the trained inference model may be used for applications other than those specifically identified herein.

[0230] For instance, movement data obtained by a single movement sensor positioned on a user (e.g., on a user's wrist or arm) may be provided as input data to a trained inference model. Corresponding output data generated by the trained inference model may be used to determine spatial information for one or more segments of a multi-segment articulated rigid body model for the user. For example, the output data may be used to determine the position and/or the orientation of one or more segments in the multi-segment articulated rigid body model. In another example, the output data may be used to determine angles between connected segments in the multi-segment articulated rigid body model.

[0231] Different types of sensors may be used to provide input data to a trained inference model, as discussed below.

[0232] As described herein, in some embodiments, various muscular activation states may be identified directly from recorded sensor data. In other embodiments, hand-states, gestures, postures, and the like (collectively referred to as muscular activation states) may be identified based, at least in part, on the output of a trained inference model. In some embodiments, the trained inference model may output motor unit or muscle activations and/or position, orientation, and/or force estimates for segments of a computer-generated musculoskeletal model. In one example, all or portions of a human musculoskeletal system can be modeled as a multi-segment articulated rigid body system, with joints forming the interfaces between the different segments, and with joint angles defining the spatial relationships between connected segments in the model.

[0233] Constraints on the movement at a joint are governed by the type of joint connecting the segments and the biological structures (e.g., muscles, tendons, ligaments) that may restrict the range of movement at the joint. For example, a shoulder joint connecting an upper arm to a torso of a body of a human subject, and a hip joint connecting an upper leg to the torso, are ball and socket joints that permit extension and flexion movements as well as rotational movements. By contrast, an elbow joint connecting an upper arm and a lower arm (or forearm), and a knee joint connecting the upper leg and a lower leg of the human subject, allow for a more limited range of motion. In this example, a multi-segment articulated rigid body system may be used to model portions of the human musculoskeletal system. However, it should be appreciated that although some segments of the human musculoskeletal system (e.g., the forearm) may be approximated as a rigid body in the articulated rigid body system, such segments may each include multiple rigid structures (e.g., the forearm may include ulna and radius bones), which enable for more complex movements within the segment that is not explicitly considered by the rigid body model. Accordingly, a model of an articulated rigid body system for use with some embodiments of the

technology described herein may include segments that represent a combination of body parts that are not strictly rigid bodies. It will be appreciated that physical models other than the multi-segment articulated rigid body system may be used to model portions of the human musculoskeletal system without departing from the scope of this disclosure.

[0234] Continuing with the example above, in kinematics, rigid bodies are objects that exhibit various attributes of motion (e.g., position, orientation, angular velocity, acceleration). Knowing the motion attributes of one segment of a rigid body enables the motion attributes for other segments of the rigid body to be determined based on constraints in how the segments are connected. For example, the hand may be modeled as a multi-segment articulated body, with joints in the wrist and each finger forming interfaces between the multiple segments in the model. In some embodiments, movements of the segments in the rigid body model can be simulated as an articulated rigid body system in which position (e.g., actual position, relative position, or orientation) information of a segment relative to other segments in the model are predicted using a trained inference model, as described in more detail below.

[0235] The portion of the human body approximated by a musculoskeletal representation, as described herein as one non-limiting example, is a hand or a combination of a hand with one or more arm segments. The information used to describe a current state of the positional relationships between segments, force relationships for individual segments or combinations of segments, and muscle and motor-unit activation relationships between segments, in the musculoskeletal representation is referred to herein as the handstate of the musculoskeletal representation. It should be appreciated, however, that the techniques described herein are also applicable to musculoskeletal representations of portions of the body other than the hand, including, but not limited to, an arm, a leg, a foot, a torso, a neck, or any combination of the foregoing.

[0236] In addition to spatial (e.g., position and/or orientation) information, some embodiments enable a prediction of force information associated with one or more segments of the musculoskeletal representation. For example, linear forces or rotational (torque) forces exerted by one or more segments may be estimated. Examples of linear forces include, but are not limited to, the force of a finger or hand pressing on a solid object such as a table, and a force exerted when two segments (e.g., two fingers) are pinched together. Examples of rotational forces include, but are not limited to, rotational forces created when a segment, such as the wrist or a finger, is twisted or flexed relative to another segment. In some embodiments, the force information determined as a portion of a current handstate estimate includes one or more of: pinching force information, grasping force information, and information about co-contraction forces between muscles represented by the musculoskeletal representation.

[0237] Turning now to the figures, FIG. 1 schematically illustrates a system **100**, for example, a neuromuscular activity system, in accordance with some embodiments of the technology described herein. The system **100** includes a plurality of sensors **110** (e.g., neuromuscular sensors) configured to sense and record signals arising from neuromuscular activity in skeletal muscle of a human body. The term "neuromuscular activity" as used herein refers to neural

activation of spinal motor neurons that innervate a muscle, muscle activation, muscle contraction, or any combination of the neural activation, muscle activation, and muscle contraction. Neuromuscular sensors may include one or more electromyography (EMG) sensors, one or more mechanomyography (MMG) sensors, one or more sonomyography (SMG) sensors, a combination of two or more types of EMG sensors, MMG sensors, and SMG sensors, and/or one or more sensors of any suitable type able to detect neuromuscular signals. In some embodiments, the plurality of neuromuscular sensors may be arranged relative to the human body and used to sense muscular activity related to a movement of the part of the body controlled by muscles from which the muscular activity is sensed by the neuromuscular sensors. Spatial information (e.g., position and/or orientation information) and force information describing the movement may be predicted based on the sensed neuromuscular signals as the user moves over time.

[0238] As the tension of a muscle increases during performance of a motor task, the firing rates of active neurons increases and additional neurons may become active, which is a process referred to as motor unit recruitment. The pattern by which neurons become active and increase their firing rate is stereotyped, such that expected motor unit recruitment patterns may define an activity manifold associated with standard or normal movement. Some embodiments may record activation of a single motor unit or a group of motor units that are “off-manifold,” in that the pattern of motor-unit activation is different than an expected or typical motor-unit recruitment pattern. Such off-manifold activation may be referred to herein as, “sub-muscular activation” or “activation of a sub-muscular structure,” where a sub-muscular structure refers to the single motor unit or the group of motor units associated with the off-manifold activation. Examples of off-manifold motor-unit recruitment patterns include, but are not limited to, selectively activating a higher-threshold motor unit without activating a lower-threshold motor unit that would normally be activated earlier in the recruitment order and modulating the firing rate of a motor unit across a substantial range without modulating the activity of other neurons that would normally be co-modulated in typical motor-unit recruitment patterns. In some embodiments, the plurality of neuromuscular sensors may be arranged relative to the human body and used to sense sub-muscular activation without observable movement, i.e., without a corresponding movement of the human body that can be readily observed. Sub-muscular activation may be used, at least in part, to control an augmented reality system in accordance with some embodiments of the technology described herein.

[0239] The sensors **110** may include one or more Inertial Measurement Units (IMUs), which measure a combination of physical aspects of motion, using, for example, an accelerometer, a gyroscope, a magnetometer, or any combination of one or more accelerometers, gyroscopes, and magnetometers. In some embodiments, IMUs may be used to sense information about the movement of the part of the body on which the IMU is attached, and information derived from the sensed data (e.g., position and/or orientation information) may be tracked as the user moves over time. For example, one or more IMUs may be used to track movements of portions (e.g., arms, legs) of a user’s body proximal to the user’s torso relative to the IMU(s) as the user moves over time.

[0240] In embodiments that include at least one IMU and a plurality of neuromuscular sensors, the IMU(s) and the neuromuscular sensors may be arranged to detect movement of different parts of a human body. For example, the IMU(s) may be arranged to detect movements of one or more body segments proximal to the torso (e.g., movements of an upper arm), whereas the neuromuscular sensors may be arranged to detect movements of one or more body segments distal to the torso (e.g., movements of a lower arm (forearm) or a wrist). (The terms “lower arm” and “forearm” may be used interchangeably herein.) It should be appreciated, however, that the sensors (i.e., the IMU(s) and the neuromuscular sensors) may be arranged in any suitable way, and embodiments of the technology described herein are not limited based on the particular sensor arrangement. For example, in some embodiments, at least one IMU and a plurality of neuromuscular sensors may be co-located on a body segment to track movements of the body segment using different types of measurements. In one implementation described in more detail below, an IMU and a plurality of EMG sensors are arranged on a wearable device structured to be worn around the lower arm or the wrist of a user. In such an arrangement, the IMU may be configured to track, over time, movement information (e.g., positioning and/or orientation) associated with one or more arm segments, to determine, for example, whether the user has raised or lowered his/her arm, whereas the EMG sensors may be configured to determine movement information associated with wrist and/or hand segments to determine, for example, whether the user has an open or a closed hand configuration, or to determine sub-muscular information associated with activation of sub-muscular structures in muscles of the wrist and/or the hand.

[0241] Some or all of the sensors **110** may include one or more sensing components configured to sense information about a user. In the case of IMUs, the sensing component(s) of an IMU may include one or more: accelerometers, gyroscopes, magnetometers, or any combination thereof, to measure or sense characteristics of body motion and/or characteristics related to body motion, examples of which include, but are not limited to, acceleration, angular velocity, and a magnetic field around the body during the body motion. In the case of neuromuscular sensors, the sensing component(s) may include, but are not limited to, one or more: electrodes that detect electric potentials on the surface of the body (e.g., for EMG sensors), vibration sensors that measure skin surface vibrations (e.g., for MMG sensors), acoustic sensing components that measure ultrasound signals (e.g., for SMG sensors) arising from muscle activity.

[0242] In some embodiments, at least some of the plurality of sensors **110** may be arranged as a portion of a wearable device structured to be worn on or around a part of a user’s body. For example, in one non-limiting example, an IMU and a plurality of neuromuscular sensors may be arranged circumferentially on an adjustable and/or elastic band, such as a wristband or an armband configured to be worn around a user’s wrist or arm. Alternatively, at least some of the sensors **110** may be arranged on a wearable patch structured to be affixed to a portion of the user’s body. In some embodiments, multiple wearable devices, each having one or more IMUs and/or neuromuscular sensors included thereon may be used to generate control information based on activation from sub-muscular structures and/or based on movement that involves multiple parts of the body.

[0243] In one implementation, the sensors 110 may include sixteen neuromuscular sensors arranged circumferentially around a band (e.g., an elastic band) structured to be worn around a user's lower arm (e.g., encircling the user's forearm). For example, FIG. 12 shows an embodiment of a wearable system in which neuromuscular sensors 1204 (e.g., EMG sensors) are arranged circumferentially around an elastic band 1202. It should be appreciated that any suitable number of neuromuscular sensors may be used and the number and arrangement of neuromuscular sensors used may depend on the particular application for which the wearable system is used. For example, a wearable armband or wristband may be used to generate control information for controlling an augmented reality system, controlling a robot, controlling a vehicle, scrolling through text, controlling a virtual avatar, or any other suitable control task. In some embodiments, the elastic band 1202 may also include one or more IMUs (not shown), configured to sense and record movement information, as discussed above.

[0244] FIGS. 13A-13B and 14A-14B show other embodiments of a wearable system of the present technology. In particular, FIG. 13A illustrates a wearable system with a plurality of sensors 1310 arranged circumferentially around an elastic band 1320 structured to be worn around a user's lower arm or wrist. The sensors 1310 may be neuromuscular sensors (e.g., EMG sensors). As shown, there may be sixteen sensors 1310 arranged circumferentially around the elastic band 1320 at a regular spacing. It should be appreciated that any suitable number of sensors 1310 may be used, and the spacing need not be regular. The number and arrangement of the sensors 1310 may depend on the particular application for which the wearable system is used. For instance, the number and arrangement of the sensors 1310 may differ when the wearable system is to be worn on a wrist in comparison with a thigh. A wearable system (e.g., armband, wristband, thighband, etc.) can be used to generate control information for controlling an augmented reality system, controlling a robot, controlling a vehicle, scrolling through text, controlling a virtual avatar, and/or performing any other suitable control task.

[0245] In some embodiments, the sensors 1310 may include only a set of neuromuscular sensors (e.g., EMG sensors). In other embodiments, the sensors 1310 may include a set of neuromuscular sensors and at least one auxiliary device. The auxiliary device(s) may be configured to continuously sense and record one or a plurality of auxiliary signal(s). Examples of auxiliary devices include, but are not limited to, IMUs, microphones, imaging devices (e.g., cameras), radiation-based sensors for use with a radiation-generation device (e.g., a laser-scanning device), heart-rate monitors, and other types of devices, which may capture a user's condition or other characteristics of the user. As shown in FIG. 13A, the sensors 1310 may be coupled together using flexible electronics 1330 incorporated into the wearable system. FIG. 13B illustrates a cross-sectional view through one of the sensors 1310 of the wearable system shown in FIG. 13A.

[0246] In some embodiments, the output(s) of one or more of sensing component(s) of the sensors 1310 can be optionally processed using hardware signal-processing circuitry (e.g., to perform amplification, filtering, and/or rectification). In other embodiments, at least some signal processing of the output(s) of the sensing component(s) can be performed using software. Thus, signal processing of signals

sampled by the sensors 1310 can be performed by hardware or by software, or by any suitable combination of hardware and software, as aspects of the technology described herein are not limited in this respect. A non-limiting example of a signal-processing procedure used to process recorded data from the sensors 1310 is discussed in more detail below in connection with FIGS. 14A and 14B.

[0247] FIGS. 14A and 14B illustrate a schematic diagram with internal components of a wearable system with sixteen sensors (e.g., EMG sensors), in accordance with some embodiments of the technology described herein. As shown, the wearable system includes a wearable portion 1410 (FIG. 14A) and a dongle portion 1420 (FIG. 14B). Although not illustrated, the dongle portion 1420 is in communication with the wearable portion 1410 (e.g., via Bluetooth or another suitable short range wireless communication technology). As shown in FIG. 14A, the wearable portion 1410 includes the sensors 1310, examples of which are described above in connection with FIGS. 13A and 13B. The sensors 1310 provide output (e.g., recorded signals) to an analog front end 1430, which performs analog processing (e.g., noise reduction, filtering, etc.) on the recorded signals. Processed analog signals produced by the analog front end 1430 are then provided to an analog-to-digital converter 1432, which converts the processed analog signals to digital signals that can be processed by one or more computer processors. An example of a computer processor that may be used in accordance with some embodiments is a microcontroller (MCU) 1434. As shown in FIG. 14A, the MCU 1434 may also receive inputs from other sensors (e.g., an IMU 1440) and from a power and battery module 1442. As will be appreciated, the MCU 1434 may receive data from other devices not specifically shown. A processing output by the MCU 1434 may be provided to an antenna 1450 for transmission to the dongle portion 1420, shown in FIG. 14B.

[0248] The dongle portion 1420 includes an antenna 1452 that communicates with the antenna 1450 of the wearable portion 1410. Communication between the antennas 1450 and 1452 may occur using any suitable wireless technology and protocol, non-limiting examples of which include radiofrequency signaling and Bluetooth. As shown, the signals received by the antenna 1452 of the dongle portion 1420 may be provided to a host computer for further processing, for display, and/or for effecting control of a particular physical or virtual object or objects (e.g., to perform a control operation in an AR or VR environment)

[0249] Although the examples provided with reference to FIGS. 13A, 13B, 14A, and 14B are discussed in the context of interfaces with EMG sensors, it is to be understood that the wearable systems described herein can also be implemented with other types of sensors, including, but not limited to, mechanomyography (MMG) sensors, sonomyography (SMG) sensors, and electrical impedance tomography (EIT) sensors.

[0250] Returning to FIG. 1, in some embodiments, sensor data recorded by the sensors 110 may be optionally processed to compute additional derived measurements, which may then be provided as input to an inference model, as described in more detail below. For example, recorded signals from an IMU may be processed to derive an orientation signal that specifies the orientation of a segment of a rigid body over time. The sensors 110 may implement signal processing using components integrated with the sensing components of the sensors 110, or at least a portion of the

signal processing may be performed by one or more components in communication with, but not directly integrated with, the sensing components of the sensors 110.

[0251] The system 100 also includes one or more computer processors 112 programmed to communicate with sensors 110. For example, signals recorded by one or more of the sensors 110 may be output from the sensor(s) 110 and provided to the processor(s) 112, which may be programmed to execute one or more machine-learning algorithms to process the signals output by the sensor(s) 110. The algorithm(s) may process the signals to train (or retrain) one or more inference models 114, and the resulting trained (or retrained) inference model(s) 114 may be stored for later use in generating control signals and controlling an XR system, as described in more detail below.

[0252] In some embodiments, the inference model(s) 114 may include a neural network and, for example, may be a recurrent neural network. In some embodiments, the recurrent neural network may be a long short-term memory (LSTM) neural network. It should be appreciated, however, that the recurrent neural network is not limited to being an LSTM neural network and may have any other suitable architecture. For example, in some embodiments, the recurrent neural network may be any one or any combination of: a fully recurrent neural network, a gated recurrent neural network, a recursive neural network, a Hopfield neural network, an associative memory neural network, an Elman neural network, a Jordan neural network, an echo state neural network, and/or any other suitable type of recurrent neural network. In other embodiments, neural networks that are not recurrent neural networks may be used. For example, deep neural networks, convolutional neural networks, and/or feedforward neural networks may be used.

[0253] In some embodiments, the inference model(s) 114 may produce discrete outputs. Discrete outputs (e.g., discrete classifications) may be used, for example, when a desired output is to know whether a particular pattern of activation (including individual neural spiking events) is currently being performed by a user. For example, the inference model(s) 114 may be trained to estimate whether the user is activating a particular motor unit, activating a particular motor unit with a particular timing, activating a particular motor unit with a particular firing pattern, or activating a particular combination of motor units. On a shorter time scale, a discrete classification may be used in some embodiments to estimate whether a particular motor unit fired an action potential within a given amount of time. In such a scenario, these estimates may then be accumulated to obtain an estimated firing rate for that motor unit.

[0254] In embodiments in which an inference model is implemented as a neural network configured to output a discrete output, the neural network may include an output layer that is a softmax layer, such that outputs of the inference model add up to one and may be interpreted as probabilities. For instance, the outputs of the softmax layer may be a set of values corresponding to a respective set of control signals, with each value indicating a probability that the user wants to perform a particular control action. As one non-limiting example, the outputs of the softmax layer may be a set of three probabilities (e.g., 0.92, 0.05, and 0.03) indicating the respective probabilities that a detected pattern of activity is one of three known patterns.

[0255] It should be appreciated that when an inference model is a neural network configured to output a discrete output (e.g., a discrete signal), the neural network is not required to produce outputs that add up to one. For example, for some embodiments, instead of a softmax layer, the output layer of the neural network may be a sigmoid layer, which does not restrict the outputs to probabilities that add up to one. In such embodiments, the neural network may be trained with a sigmoid cross-entropy cost. Such an implementation may be advantageous in cases where multiple different control actions may occur within a threshold amount of time and it is not important to distinguish an order in which these control actions occur (e.g., a user may activate two patterns of neural activity within the threshold amount of time). In some embodiments, any other suitable non-probabilistic multi-class classifier may be used, as aspects of the technology described herein are not limited in this respect.

[0256] In some embodiments, an output of the inference model(s) 114 may be a continuous signal rather than a discrete signal. For example, the model(s) 114 may output an estimate of a firing rate of each motor unit, or the model(s) 114 may output a time-series electrical signal corresponding to each motor unit or sub-muscular structure.

[0257] It should be appreciated that aspects of the technology described herein are not limited to using neural networks, as other types of inference models may be employed in some embodiments. For example, in some embodiments, the inference model(s) 114 may comprise a hidden Markov model (HMM), a switching HMM in which switching allows for toggling among different dynamic systems, dynamic Bayesian networks, and/or any other suitable graphical model having a temporal component. Any such inference model may be trained using recorded sensor signals.

[0258] As another example, in some embodiments, the inference model(s) 114 may be a classifier that takes, as input, features derived from the recorded sensor signals. In such embodiments, the classifier may be trained using features extracted from the sensor signals. The classifier may be, e.g., a support vector machine, a Gaussian mixture model, a regression based classifier, a decision tree classifier, a Bayesian classifier, and/or any other suitable classifier, as aspects of the technology described herein are not limited in this respect. Input features to be provided to the classifier may be derived from the sensor signals in any suitable way. For example, the sensor signals may be analyzed as time-series data using wavelet analysis techniques (e.g., continuous wavelet transform, discrete-time wavelet transform, etc.), Fourier-analytic techniques (e.g., short-time Fourier transform, Fourier transform, etc.), and/or any other suitable type of time-frequency analysis technique. As one non-limiting example, the sensor signals may be transformed using a wavelet transform and the resulting wavelet coefficients may be provided as inputs to the classifier.

[0259] In some embodiments, values for parameters of the inference model(s) 114 may be estimated from training data. For example, when the inference model(s) 114 is a neural network, parameters of the neural network (e.g., weights) may be estimated from the training data. In some embodiments, parameters of the inference model(s) 114 may be estimated using gradient descent, stochastic gradient descent, and/or any other suitable iterative optimization technique. In embodiments where the inference model(s)

114 is a recurrent neural network (e.g., an LSTM), the inference model(s) **114** may be trained using stochastic gradient descent and backpropagation through time. The training may employ a squared error or cross-entropy loss function and/or any other suitable loss function, as aspects of the technology described herein are not limited in this respect.

[0260] The system **100** also may optionally include one or more controller(s) **116**. For example, the controller(s) **116** may include a display controller configured to display a visual representation (e.g., a representation of a hand). As discussed in more detail below, the one or more computer processors **112** may implement one or more trained inference models that receive, as input, signals sensed and recorded by the sensors **110** and that provide, as output, information (e.g., predicted handstate information) that may be used to generate control signals and control an augmented reality system.

[0261] The system **100** also may optionally include a user interface **118**. Feedback determined based on the signals recorded by the sensors **110** and processed by the processor (s) **112** may be provided via the user interface **118** to facilitate a user's understanding of how the system **100** is interpreting the user's intended activation. The user interface **118** may be implemented in any suitable way, including, but not limited to, an audio interface, a video interface, a tactile interface, and electrical stimulation interface, or any combination of the foregoing.

[0262] In some embodiments, a computer application that simulates a virtual reality (VR) or augmented reality (AR) environment may be instructed to provide a visual representation by displaying a visual character, such as an avatar (e.g., via the controller(s) **116**). Positioning, movement, and/or forces applied by portions of the visual character within the virtual reality environment may be displayed based on an output of the trained inference model(s) **114**. The visual representation may be dynamically updated as continuous signals are sensed and recorded by the sensors **110** and processed by the trained inference model(s) **114** to provide a computer-generated visual representation of the character's movement that is updated in real-time.

[0263] Information generated in either system (AR camera inputs, sensor inputs) can be used to improve user experience, accuracy, feedback, inference models, calibration functions, and other aspects in the overall system. To this end, in an AR environment for example, the system **100** may include an AR system that includes one or more processors, a camera, and a display (e.g., the user interface **118**, or other interface via AR glasses or other viewing device) that provides AR information within a view of the user. The system **100** may also include system elements that couple the AR system with a computer-based system that generates the musculoskeletal representation based on sensor data. For example, the systems may be coupled via a special-purpose or other type of computer system that receives inputs from the AR and system that generates the computer-based musculoskeletal representation. Such a system may include a gaming system, robotic control system, personal computer, or other system that is capable of interpreting AR and musculoskeletal information. The AR system and the system that generates the computer-based musculoskeletal representation may also be programmed to communicate directly. Such information may be communicated using any number of interfaces, protocols, or media.

[0264] As discussed above, some embodiments are directed to using an inference model for predicting musculoskeletal information based on signals sensed and recorded by wearable sensors. As discussed briefly above in the example where portions of the human musculoskeletal system can be modeled as a multi-segment articulated rigid body system, the types of joints between segments in a multi-segment articulated rigid body model may serve as constraints that constrain movement of the rigid body. Additionally, different human individuals may move in characteristic ways when performing a task that can be captured in statistical patterns that may be generally applicable to individual user behavior. At least some of these constraints on human body movement may be explicitly incorporated into inference models used for prediction of user movement, in accordance with some embodiments. Additionally or alternatively, the constraints may be learned by the inference model through training based on recorded sensor data, as discussed briefly above.

[0265] As discussed above, some embodiments are directed to using an inference model for predicting handstate information to enable generation of a computer-based musculoskeletal representation and/or a real-time update of a computer-based musculoskeletal representation. The inference model may be used to predict the handstate information based on IMU signals, neuromuscular signals (e.g., EMG, MMG, and/or SMG signals), external device signals (e.g., camera or laser-scanning signals), or a combination of IMU signals, neuromuscular signals, and external device signals detected as a user performs one or more movements. For instance, as discussed above, a camera associated with an AR system may be used to capture data of an actual position of a human subject of the computer-based musculoskeletal representation, and such actual-position information may be used to improve the accuracy of the representation. Further, outputs of the inference model may be used to generate a visual representation of the computer-based musculoskeletal representation in an XR environment. For example, a visual representation of muscle groups firing, force being applied, text being entered via movement, or other information produced by the computer-based musculoskeletal representation may be rendered in a visual display of an XR system. In some embodiments, other input/output devices (e.g., auditory inputs/outputs, haptic devices, etc.) may be used to further improve the accuracy of the overall system and/or to improve user experience.

[0266] Some embodiments of the technology described herein are directed to using an inference model, at least in part, to map muscular-activation state information, which is information identified from neuromuscular signals sensed and recorded by neuromuscular sensors, to control signals. The inference model may receive as input IMU signals, neuromuscular signals (e.g., EMG, MMG, and/or SMG signals), external device signals (e.g., camera or laser-scanning signals), or a combination of IMU signals, neuromuscular signals, and external device signals detected as a user performs one or more sub-muscular activations, one or more movements, and/or one or more gestures. The inference model may be used to predict control information without the user having to make perceptible movements.

[0267] FIG. 2 illustrates a schematic diagram of an AR-based system **200**, which may be a distributed computer-based system that integrates an augmented reality (AR) system **201** with a neuromuscular activity system **202**. The

neuromuscular activity system **202** is similar to the system **100** described above with respect to FIG. 1.

[0268] Generally, an augmented reality, such as the AR system **201**, may take the form of a pair of goggles or glasses, or eyewear, or other type of display device that shows display elements to a user that may be superimposed on the user's "reality." This reality in some cases could be the user's view of the environment (e.g., as viewed through the user's eyes), or a captured version (e.g., by camera(s)) of the user's view of the environment. In some embodiments, the AR system **201** may include one or more cameras (e.g., camera(s) **204**), which may be mounted within a device worn by a user, that captures one or more views experienced by the user in the user's environment. The system **201** may have one or more processor(s) **205** operating within the device worn by the user and/or within a peripheral device or computer system, and such processor(s) **205** may be capable of transmitting and receiving video information and other types of data (e.g., sensor data).

[0269] The AR system **201** may also include one or more sensor(s) **207**, such as microphones, GPS elements, accelerometers, infrared detectors, haptic feedback elements, or any other type of sensor, or any combination thereof. In some embodiments, the AR system **201** may be an audio-based or auditory AR system and the one or more sensor(s) **207** may also include one or more headphones or speakers. Further, the AR system **201** may also have one or more display(s) **208** that permit the AR system **201** to overlay and/or display information to the user in addition to providing the user with a view of the user's environment as presented by the AR system **201**. The AR system **201** may also include one or more communication interface(s) **206**, which enable information to be communicated to one or more computer systems (e.g., a gaming system or other system capable of rendering or receiving AR data). AR systems can take many forms and are available from a number of different manufacturers. For example, various embodiments may be implemented in association with one or more types of AR systems, such as HoloLens holographic reality glasses available from the Microsoft Corporation (Redmond, Wash., USA), Lightwear AR headset from Magic Leap (Plantation, Fla., USA), Google Glass AR glasses available from Alphabet (Mountain View, Calif., USA), R-7 Smartglasses System available from Osterhout Design Group (also known as ODG; San Francisco, Calif., USA), or any other type of AR and/or VR device. Although discussed by way of example, it should be appreciated that one or more embodiments may be implemented within a VR system.

[0270] The AR system **201** may be operatively coupled to the neuromuscular activity system **202** through one or more communication schemes or methodologies, including but not limited to, Bluetooth protocol, Wi-Fi, Ethernet-like protocols, or any number of connection types, wireless and/or wired. It should be appreciated that, for example, the systems **201** and **202** may be directly connected or coupled through one or more intermediate computer systems or network elements. The double-headed arrow in FIG. 2 represents the communicative coupling between the systems **201** and **202**.

[0271] As mentioned earlier, the neuromuscular activity system **202** may be similar in structure and function to the system **100** described above with reference to FIG. 1. In particular, the system **202** may include one or more neuro-

muscular sensor(s) **209**, one or more inference model(s) **210**, and may create, maintain, and store a musculoskeletal representation **211**. In an example embodiment, similar to one discussed above, the system **202** may include or may be implemented as a wearable device, such as a band that can be worn by a user, in order to collect and analyze neuromuscular signals from the user. Further, the system **202** may include one or more communication interface(s) **212** that permit the system **202** to communicate with the AR system **201**, such as by Bluetooth, Wi-Fi, or other communication method. Notably, the AR system **201** and the neuromuscular activity system **202** may communicate information that can be used to enhance user experience and/or allow the AR system **201** to function more accurately and effectively.

[0272] While FIG. 2 shows a distributed computer-based system **200** that integrates the AR system **201** with the neuromuscular activity system **202**, it will be understood integration of these systems **201** and **202** may be non-distributed in nature. In some embodiments, the neuromuscular activity system **202** may be integrated into the AR system **201** such that the various components of the neuromuscular activity system **202** may be considered as part of the AR system **201**. For example, inputs from the neuromuscular sensor(s) **209** may be treated as another of the inputs (e.g., from the camera(s) **204**, from the sensor(s) **207**) to the AR system **201**. In addition, processing of the inputs (e.g., sensor signals) obtained from the neuromuscular sensors **209** may be integrated into the AR system **201**.

[0273] FIG. 3 shows a flowchart of a process **300** for processing neuromuscular signals and camera or signals inputs in accordance with some embodiments of the technology described herein. The process **300** may be implemented on a system such as, for example, the AR-based system **200**, the AR system **201**, the neuromuscular activity system **202**, and/or a separate computer system to which the systems **201**, **202** provide signal inputs. In one implementation, the process **300** may be performed by the neuromuscular activity system **202**. In act **310**, sensor signals may be sensed and recorded by one or more sensor(s) (also referred to herein as "raw sensor signals") of the neuromuscular activity system **202**. In some embodiments, the sensor(s) may include a plurality of neuromuscular sensors **209** (e.g., EMG sensors) arranged on a wearable device worn by a user. For example, the sensors **209** may be EMG sensors arranged on an elastic band configured to be worn around a wrist or a forearm of a user to record neuromuscular signals from the user as the user performs various movements or gestures. In some embodiments, the EMG sensors may be the sensors **1204** arranged on the band **1202**, as shown in FIG. 12; in some embodiments, the EMG sensors may be the sensors **1310** arranged on the elastic band **1320**, as shown in FIG. 13A. The gestures performed by the user may include static gestures, such as placing the user's hand palm down on a table; dynamic gestures, such as waving a finger back and forth; and covert gestures that are imperceptible to another person, such as slightly tensing a joint by co-contracting opposing muscles, pressing on an object or surface, or using sub-muscular activations. The gestures performed by the user may include symbolic gestures (e.g., gestures mapped to other gestures, interactions, or commands, for example, based on a gesture vocabulary that specifies the mapping).

[0274] In addition to a plurality of neuromuscular sensors, some embodiments of the technology described herein may

include one or more auxiliary sensor(s) configured to record auxiliary signals that may also be provided as input to the one or more trained inference model(s), as discussed above. Examples of auxiliary sensors include IMUs, imaging devices, radiation detection devices (e.g., laser scanning devices), heart rate monitors, or any other type of biosensor configured to sense and record biophysical information from a user during performance of one or more movements or gestures. According to one embodiment, as shown in the process 300 of FIG. 3 at act 320, the system receives and processes one or more camera inputs. Such inputs may include one or more raw signals, such as signals of images, video, stream, etc., or signals of one or more pre-processed forms, such as signals of a representation of a detected object, a 3D model, etc., and/or other information, such as information reflecting a state of the camera inputs. Further, it should be appreciated that some embodiments may be implemented using camera-based systems that perform skeletal tracking, such as, for example, the Kinect system available from the Microsoft Corporation (Redmond, Wash., USA) and the LeapMotion system available from Leap Motion, Inc. (San Francisco, Calif., USA). It should be appreciated that any combination of hardware and/or software may be used to implement various embodiments described herein.

[0275] The acts 310 and 320 may also include processing acts themselves, where raw sensor signals, which may include the signals sensed and recorded by the one or more sensor(s) (e.g., EMG sensors, IMUs) and/or the camera input signals from the one or more camera(s), are optionally processed. In some embodiments, the raw sensor signals may be processed using hardware signal-processing circuitry (e.g., to perform amplification, filtering, and/or rectification). In other embodiments, at least some signal processing of the raw sensor signals may be performed using software. Accordingly, signal processing of the raw sensor signals, sensed and recorded by the one or more sensor(s) and/or obtained from the one or more camera(s), may be performed using hardware, or software, or any suitable combination of hardware and software. In some implementations, the raw sensor signals may be processed to derive other signal data. For example, accelerometer data recorded by one or more IMU(s) may be integrated and/or filtered to determine derived signal data associated with one or more muscles during activation of a muscle or performance of a gesture.

[0276] The process 300 then proceeds to act 330, where the raw sensor signals (e.g., the neuromuscular signals and the camera input signals) or the processed versions of the raw sensor signals are optionally provided as input to the trained inference model(s), which is or are configured to produce user-movement information, such as handstate information, as described above.

[0277] The process 300 then proceeds to act 340, where outputs of the trained inference model(s) are provided. For instance, in some embodiments, control outputs of the system are provided based on the raw sensor signals, the processed sensor signals, and/or the outputs of the trained inference model(s) (e.g., handstate information and/or other rendered output of the trained inference model(s), etc.). For example, in some embodiments, the AR system 201 may receive a rendered output that the AR system 210 (or other system display) can display as a rendered gesture.

[0278] According to some embodiments, one or more computer processors (e.g., the processor(s) 112 of the system 100, or the processor(s) 205 of the AR-based system 200) may be programmed to identify one or more muscular activation states of a user from raw sensor signals (e.g., signals sensed and recorded by the sensor(s) discussed above and/or camera input signals) and/or information based on these signals. The information based on the signals the raw sensor signals may include information associated with processed sensor signals (e.g., processed EMG signals) and/or information associated with outputs of the trained inference model(s) (e.g., handstate information). The one or more muscular activation states of the user may include a static gesture performed by the user, a dynamic gesture performed by the user, a sub-muscular activation state of the user, and/or a muscular tensing performing by the user. The one or more muscular activation states of the user may be defined by one or more pattern(s) of muscle activity and/or one or more motor unit activation(s) detected in the raw sensor signals and/or information based on the raw sensor signals, associated with various movements or gestures performed by the user.

[0279] FIG. 4 is a flowchart of a process 400 for processing gesture information in an AR-based system in accordance with some embodiments of the technology described herein. In particular, there may be a number of different use scenarios wherein an AR system may be used in association with a neuromuscular activity system for providing a user experience. As shown in FIG. 4, according to the process 400, at act 410, the AR-based system may be configured to display a model gesture to a user within a display of the AR-based system. For example, the AR-based system may be adapted to display a throwing motion or other gesture type that the user would like to mimic or learn. To this end, the AR-based system may display an example of a gesture to the user while the user tries to emulate that particular gesture. At act 420, a neuromuscular activity system processes neuromuscular signals produced by the user while emulating that gesture. At act 430, the neuromuscular activity system determines whether the emulated gesture matches a model of that gesture. Such a determination may be made, for example, using one or more 3D reference model(s) and comparing outputs of an inference model that measures or compares the user's emulated gesture versus a reference model of the gesture. Acceptable matches may occur within certain predefined ranges or error levels between two or more reference models. It should be appreciated that camera data (e.g., input signals from a camera) may also be used (either alone or in combination with the neuromuscular signals) to determine whether a match occurs. For instance, in such cases where neuromuscular signals (e.g., EMG signals) are not available, a match may be detected exclusively using camera data.

[0280] Further, the inference model may also be trained to provide an indication of a match between models. At act 440, feedback may be provided to the user (discussed in more detail below in connection with FIG. 5), such as a display indication within an AR display (e.g., a light or other indicator identifying that the gesture was matched or was not matched); an audio output (e.g., an audio signal for spoken words such as "great job", "match", etc.); haptic feedback; electrical stimulation; and/or other output perceivable by the user. It should be appreciated that one aspect of such a comparison may include a prediction as to whether the

gesture will be completed correctly (or incorrectly). For example, hardware and/or software in the AR-based system may perform a routine for a predictive measure that is capable of determining whether the user is likely to successfully complete the gesture. In one example implementation, an audio output may be presented to the user when the predictive measure indicates that the user will be unsuccessful.

[0281] As will be appreciated, the feedback to the user may relate to a user's position relative to a camera. For example, if the user is wholly or partially occluded from the camera's field of view, the feedback may inform the user when a relevant part of the user's body is in view.

[0282] It should be appreciated that, in some embodiments, some outputs may be delivered to the user before completion of the gesture, based on a prediction of the gesture (e.g., the user may receive an electric stimulus indicating that the user's attempt at performing the gesture will be unsuccessful before the user completes the gesture). Also, it should be appreciated that although a matching operation may be performed for purposes of training an inference model, calibration of the AR-based system, and/or training of the user in an AR context, data that pairs, e.g., EMG signals and joint angles may be used to train and/or correct a musculoskeletal inference model in real time to create more accurate joint-angle predictions. However, for a gesture classifier, it may be desirable to determine specifically whether the user is performing a given gesture. In such a case, one or more camera(s) (e.g., the camera(s) 204) may be used to obtain information usable to provide training and/or classification data, such as gesture labels, by detecting matches or targeting joint angles against which to regress the model.

[0283] FIG. 5 is a flowchart of a process 500 for integrating neuromuscular signals and camera input signals, and providing feedback to a user, in accordance with some embodiments of the technology described herein. The process 500 may be performed by the AR-based system 200 of FIG. 2. At act 510, the neuromuscular activity system 202 receives and processes signals in a similar manner as discussed above with reference to FIG. 3. Further, at act 520, the system 200 receives and processes camera input in a manner similar to that discussed above with reference to FIG. 3 (e.g., at the act 320). At act 530, the system 200 displays feedback to the user (e.g., within one or more of the display(s) 208 of the AR-based system 200). Alternatively, the system 200 may provide feedback in any number of different ways, such as audio feedback, haptic feedback, or other output perceivable by the user. To this end, the neuromuscular activity system 202 may be capable of providing output to other parts of the AR-based system 200 for the purpose of displaying or generating feedback output. In some implementations, the AR-based system 200 may have one or more audio outputs, displays, indicators and/or other type of output device capable of providing or rendering feedback to the user. Alternatively, the system 200 may be configured to provide or render an output signal (e.g., a control signal) to one or more third-party systems (e.g., a manager system, monitor system, supervisor system, or other type of third-party system).

[0284] FIG. 6 is a flowchart of a process 600 for updating an inference model based on camera data in accordance with some embodiments of the technology described herein. The process 600 may be performed by the AR-based system 200

of FIG. 2. At act 610, the system 200 receives and processes camera input (e.g., an input signal of a captured image). In one embodiment, at act 620, the system 200 determines position information from the captured image or from other input (e.g., neuromuscular data). At act 630, the system 200 may update a parameter of the inference model based on the determined position information. For example, in a situation where the neuromuscular activity system 202 determines a hand position and/or an arm position, camera input data (e.g., an image of the hand and/or the arm) may be used to correct position information for that particular inference model in a calibration action, i.e., to calibrate the inference model. In cases where there are two or more sources of data (e.g., camera data and EMG data), the inference model may require that the data from the sources be simultaneously acquired. For instance, there may be error or drift within a signal corresponding to EMG-based position information, and a ground truth position determined from a camera input signal may be used to adjust a parameter of an EMG-based inference model. In this way, accuracy of the EMG-based inference model may be improved by use of both the EMG and camera input signals to calibrate the EMG-based inference model. Such a process may be used as part of a calibration operation where a user performs discrete gestures and/or performs gestures in real time in a "normal" operation, and an inference model for the gestures needs to be adjusted based on ground truth data generated by a signal source other than the source providing the neuromuscular signals for the model; the other signal source may, for example be a camera. In some embodiments, camera data may be used to estimate continuous joint angles (e.g., for a regression model) or to identify discrete gestures (e.g., for a classifier).

[0285] FIG. 7 is a flowchart of a process 700 for updating an inference model based on camera data in accordance with some embodiments of the technology described herein. The process 700 may be performed by the AR-based system 200 of FIG. 2. Similar to the act 610 of FIG. 6, at act 710 the system 200 receives and processes camera input (e.g., an input signal of a captured image). At act 720, the system 200 determines geometry information from the captured image. For instance, when capturing an image of a hand, geometrical information such as segment lengths, joint positions, and other geometrical relations of a hand may be captured. At act 730, such geometrical information may be used to update one or more parameters of an inference model. In one example implementation, the system 200 may correct a geometry used by an EMG model of a hand. It is appreciated, for example, that an image captured by a camera may more accurately portray a geometry of a hand of the user than an image generated using EMG-based signals. For instance, an image of a segment length of a finger as captured by a camera may be used to update an EMG-based inference model of hand geometry. This may be especially pertinent when the EMG-based inference model is corrected or updated for use with that particular user.

[0286] FIG. 8 is a flowchart 800 of a process for determining position information and force information in accordance with some embodiments of the technology described herein. In particular, as discussed above, it may be beneficial to use both neuromuscular signals and camera input signals to more accurately determine position and force information within an AR environment. The process 800 may be performed by the AR-based system 200 of FIG. 2. At act 810,

the neuromuscular activity system **202** receives and processes neuromuscular input signals. Further, at act **820**, the AR system **201** receives and processes at least one camera input signal. At act **830**, the neuromuscular input signals and the camera input signal(s) are provided to an inference model. Based on both inputs, the system **200** may determine position information (e.g., at act **840**), and may determine force information (e.g., at act **860**). The position information and the force information may be used, for example, to render a representation of a user's hand within an AR display. For example, the representation may indicate a more accurate position of the user's hand and finger appendages as well as any forces that may be applied to any of the hand and/or the finger appendages. In one specific example, the system **200** may more accurately render the user's hand position where the user's thumb is pressing on the index finger of the same hand. The camera input may be particularly useful, as it may provide a visual indication that the index finger and the thumb are in contact, while the neuromuscular signals may indicate that forces are being applied by the index finger and the thumb at some relative level (e.g., the index finger and the thumb are being pressed together with a light amount of force). In this way, the system **200** may more accurately represent handstate information of the user's hand. Optionally, the system **200** may provide position information and force information as outputs, at acts **850** and **870**, respectively. Such outputs may be used by one or more other systems to provide indications, render outputs or other representations, and/or control other systems.

[**0287**] In some implementations, the camera input signals may be processed to determine whether one or more fingers are in contact with the user's thumb and, if so, which fingers are in contact with the thumb. In one embodiment, if there is contact, a magnitude of the force between the thumb and the contacting fingers then may be estimated as an affine transformation as a function of the logarithm of a signal power recorded from the neuromuscular sensors that provided the neuromuscular input signals (e.g., EMG signal power) or with another signal processing technique for inferring a magnitude of force. Coefficients of the affine transformation function can be determined by a calibration stage in which the user first lightly touches the fingers and then later produces a maximum voluntary contraction corresponding to a pinch force between the fingers. In this implementation, the inference model may output a force in units of a fraction of the maximum voluntary contraction. In some implementations, the coefficients of the affine transformation function can be specific to the set of fingers contacting the thumb. Similarly, affine transformation functions of EMG signal power can be calibrated for cases where the user is applying forces against objects or surfaces.

[**0288**] FIG. **9** is a flowchart of a process **900** for determining qualities of input signals and performing model functions based on those qualities, in accordance with some embodiments of the technology described herein. The process **900** may be performed by the AR-based system **200** of FIG. **2**. For example, it is appreciated that certain input signals may be more reliable than others in certain situations. For example, a camera input signal may be less reliable for information when an object to be modeled is occluded from view or not within the camera's view at all. Similarly, neuromuscular signals may include errors generated by varying physical conditions, noise, or other interference. Therefore, it may be beneficial for a system to

adaptively determine the qualities of particular input signals so that a model may be properly updated.

[**0289**] At act **910**, the system **200** receives and processes neuromuscular signals (e.g. as received by the neuromuscular activity system **202**). At act **920**, the AR system **201** receives and processes one or more camera input signal(s). At act **930**, the system **200** determines a quality of the neuromuscular signals. For instance, the neuromuscular activity system **202** or another suitable system may determine whether the neuromuscular signals should be used and, if so, to what level. For instance, certain portions of the neuromuscular signals may be accurate (e.g., a force measurement) but other portions of the neuromuscular signals may not (e.g., absolute forearm position). Certain portions of the neuromuscular signals may be reliable within certain time frames, and under varying conditions. For example, EMG signals having a substantial power (e.g., power above a predetermined threshold) at low (e.g., less than 20 Hz) or high (e.g., greater than 500 Hz) frequencies may indicate a low quality of the EMG signals.

[**0290**] At act **940**, the system **200** may determine a quality of the camera input signals. As discussed above, a camera signal may be determined to be of a particular quality level when certain objects are within a field of view, and a camera signal may be determined to be at a lower quality level when objects are occluded or not within the field of view. At act **950**, the system **200** may provide a signal to an inference model based on the determined qualities of the camera input signal(s) and the neuromuscular signals. For example, one or more camera input signal(s) or neuromuscular signals may be filtered or deprecated prior to being input into the inference model. In another implementation, the inference model may be trained under varying conditions and may be trained to be responsive to signals of different quality. In another implementation, separate Bayesian predictions and confidences may be determined from each of the two types of input, and Bayesian approaches may be used to combine the two predictions (e.g., by taking a weighted average). At act **960** the system **200** provides an output of the inference model.

[**0291**] FIG. **10** is a diagram showing an example implementation of a system **1000** that utilizes one or more EMG sensor(s) **1040** and a camera **1060**, in accordance with some embodiments of the technology described herein. For example, FIG. **10** shows a user's arm and an attached hand ("arm/hand") **1010**, which is made up of one or more joints and segments, and which can be depicted as a musculoskeletal representation. More particularly, the user's hand segments **1020** are connected by joints. The arm and hand positions and segment lengths of the arm and the hand can be determined by the system **1000** and positioned within a three-dimensional space of a model musculoskeletal representation. Further, the user's hand may also include an interpolated forearm segment **1030**. As discussed above, a neuromuscular activity system may be used to determine one or more representations of a user's hand/arm positions. To this end, the user may wear a band comprising the one or more EMG sensor(s) **1040**, which sense and record neuromuscular signals that are used to determine a musculoskeletal skeletal representation. Concurrently with the EMG sensor(s) **1040** sensing and recording the neuromuscular signals, a camera **1060** may be used to capture objects within the camera's field of view **1050**. For example, in FIG. **10**, the camera's field of view **1050** include the user's arm/hand

1010. Camera data in addition to the neuromuscular activity signals determined by the EMG sensors **1040** may be used to reconstruct positions, geometries, and forces being applied by the user's arm/hand **1010**. Further, outputs from the system **1000** can be provided that allow the system **1000** to render a representation of the user's arm/hand **1010**, such as within an AR display or other type of system.

[**0292**] FIG. **11** is a diagram showing a process **1100** of processing inputs to a trained inference model in accordance with some embodiments of the technology described herein. For example, it is appreciated that a more accurate representation of a musculoskeletal representation may be obtained by using IMU inputs (**1101**), EMG inputs (**1102**), and camera inputs (**1103**). Each of these inputs **1101**, **1102**, **1103** may be provided to a trained inference model **1110**. The inference model **1110** may be structured to provide one or more outputs, such as position, force, and/or a representation of a musculoskeletal state. Such outputs may be provided to one or more system(s), such as an AR system for indication or display, as discussed above, or for providing feedback to a user. It should be appreciated that any of the inputs **1101**, **1102**, **1103** may be used in any combination with any other input to derive an output from the trained inference model. For instance, forearm positional information may be derived based on a combination of the IMU input **1101** and the camera input **1103**. In one implementation, an estimate of forearm position may be generated based on the IMU input **1101** and adjusted based on ground truth data obtained from the camera input **1103**. Also, forearm position and/or forearm orientation may be derived using the camera input **1103** alone without the IMU input **1101**. In another scenario, the EMG input **1102** (e.g., EMG signals) may be used to derive force-only information to augment posture-only information provided by a camera-model system. Other combinations of inputs may be used to obtain a desired output and are within the scope of various embodiments described herein.

[**0293**] It should also be appreciated that such outputs may be derived with or without generating any musculoskeletal representation. It should also be appreciated that one or more outputs may be used as control inputs to any other system, such as an output of an EMG-based control signal that is used as a control input signal provided to an AR system.

[**0294**] It is appreciated that any embodiment described herein may be used alone or in any combination with any other embodiment described herein. Further, portions of an embodiment described herein may be combined with portions of one or more other embodiments described herein. Additionally, embodiments described herein may be used in whole or in part with embodiments described in U.S. patent application Ser. No. 16/257,979, filed Jan. 25, 2019, entitled "CALIBRATION TECHNIQUES FOR HANDSTATE REPRESENTATION MODELING USING NEUROMUSCULAR SIGNALS," which is incorporated by reference herein, and U.S. patent application Ser. No. 15/659,504, filed Jul. 25, 2017, entitled "SYSTEM AND METHOD FOR MEASURING THE MOVEMENTS OF ARTICULATED RIGID BODIES," which is incorporated by reference herein.

[**0295**] The above-described embodiments can be implemented in any of numerous ways. For example, the embodiments may be implemented using hardware, software or a combination thereof. When implemented using software, code comprising the software can be executed on any

suitable processor or collection of processors, whether provided in a single computer or distributed among multiple computers. It should be appreciated that any component or collection of components that perform the functions described above can be generically considered as one or more controllers that control the above-discussed functions. The one or more controllers can be implemented in numerous ways, such as with dedicated hardware or with one or more processors programmed using microcode or software to perform the functions recited above.

[**0296**] In this respect, it should be appreciated that one implementation of the embodiments of the present invention comprises at least one non-transitory computer-readable storage medium (e.g., a computer memory, a portable memory, a compact disk, etc.) encoded with a computer program (i.e., a plurality of instructions), which, when executed on a processor, performs the above-discussed functions of the embodiments of the technologies described herein. The computer-readable storage medium can be transportable such that the program stored thereon can be loaded onto any computer resource to implement the aspects of the present invention discussed herein. In addition, it should be appreciated that reference to a computer program that, when executed, performs the above-discussed functions, is not limited to an application program running on a host computer. Rather, the term computer program is used herein in a generic sense to reference any type of computer code (e.g., software or microcode) that can be employed to program a processor to implement the above-discussed aspects of the present invention.

[**0297**] Various aspects of the technology presented herein may be used alone, in combination, or in a variety of arrangements not specifically discussed in the embodiments described above and therefore are not limited in their application to the details and arrangements of components set forth in the foregoing description and/or in the drawings.

[**0298**] Also, some of the embodiments described above may be implemented as one or more method(s), of which some examples have been provided. The acts performed as part of the method(s) may be ordered in any suitable way. Accordingly, embodiments may be constructed in which acts are performed in an order different than illustrated or described herein, which may include performing some acts simultaneously, even though shown as sequential acts in illustrative embodiments. The phraseology and terminology used herein is for the purpose of description and should not be regarded as limiting. The use of "including," "comprising," "having," "containing", "involving", and variations thereof, is meant to encompass the items listed thereafter and additional items.

[**0299**] Having described several embodiments of the invention in detail, various modifications and improvements will readily occur to those skilled in the art. Such modifications and improvements are intended to be within the spirit and scope of the invention. Accordingly, the foregoing description is by way of example only, and is not intended as limiting. The invention is limited only as defined by the following claims and the equivalents thereto.

[**0300**] The foregoing features may be used, separately or together in any combination, in any of the embodiments discussed herein.

[**0301**] Further, although advantages of the present invention may be indicated, it should be appreciated that not every embodiment of the invention will include every described

advantage. Some embodiments may not implement any features described as advantageous herein. Accordingly, the foregoing description and attached drawings are by way of example only.

[0302] Variations on the disclosed embodiment are possible. For example, various aspects of the present technology may be used alone, in combination, or in a variety of arrangements not specifically discussed in the embodiments described in the foregoing, and therefore they are not limited in application to the details and arrangements of components set forth in the foregoing description or illustrated in the drawings. Aspects described in one embodiment may be combined in any manner with aspects described in other embodiments.

[0303] Use of ordinal terms such as “first,” “second,” “third,” etc., in the description and/or the claims to modify an element does not by itself connote any priority, precedence, or order of one element over another, or the temporal order in which acts of a method are performed, but are used merely as labels to distinguish one element or act having a certain name from another element or act having a same name (but for use of the ordinal term) to distinguish the elements or acts.

[0304] The indefinite articles “a” and “an,” as used herein in the specification and in the claims, unless clearly indicated to the contrary, should be understood to mean “at least one.”

[0305] Any use of the phrase “at least one,” in reference to a list of one or more elements, should be understood to mean at least one element selected from any one or more of the elements in the list of elements, but not necessarily including at least one of each and every element specifically listed within the list of elements and not excluding any combinations of elements in the list of elements. This definition also allows that elements may optionally be present other than the elements specifically identified within the list of elements to which the phrase “at least one” refers, whether related or unrelated to those elements specifically identified.

[0306] Any use of the phrase “equal” or “the same” in reference to two values (e.g., distances, widths, etc.) means that two values are the same within manufacturing tolerances. Thus, two values being equal, or the same, may mean that the two values are different from one another by $\pm 5\%$.

[0307] The phrase “and/or,” as used herein in the specification and in the claims, should be understood to mean “either or both” of the elements so conjoined, i.e., elements that are conjunctively present in some cases and disjunctively present in other cases. Multiple elements listed with “and/or” should be construed in the same fashion, i.e., “one or more” of the elements so conjoined. Other elements may optionally be present other than the elements specifically identified by the “and/or” clause, whether related or unrelated to those elements specifically identified. Thus, as a non-limiting example, a reference to “A and/or B,” when used in conjunction with open-ended language such as “comprising” can refer, in one embodiment, to A only (optionally including elements other than B); in another embodiment, to B only (optionally including elements other than A); in yet another embodiment, to both A and B (optionally including other elements); etc.

[0308] As used herein in the specification and in the claims, “or” should be understood to have the same meaning as “and/or” as defined above. For example, when separating

items in a list, “or” or “and/or” shall be interpreted as being inclusive, i.e., the inclusion of at least one, but also including more than one, of a number or list of elements, and, optionally, additional unlisted items. Only terms clearly indicated to the contrary, such as “only one of” or “exactly one of,” or, when used in the claims, “consisting of,” will refer to the inclusion of exactly one element of a number or list of elements. In general, the term “or” as used herein shall only be interpreted as indicating exclusive alternatives (i.e. “one or the other but not both”) when preceded by terms of exclusivity, such as “either,” “one of,” “only one of,” or “exactly one of.” “Consisting essentially of,” when used in the claims, shall have its ordinary meaning as used in the field of patent law.

[0309] Also, the phraseology and terminology used herein is for the purpose of description and should not be regarded as limiting. Use of terms such as “including,” “comprising,” “comprised of,” “having,” “containing,” and “involving,” and variations thereof herein, is meant to encompass the items listed thereafter and equivalents thereof as well as additional items.

[0310] The terms “approximately” and “about” if used herein may be construed to mean within $\pm 20\%$ of a target value in some embodiments, within $\pm 10\%$ of a target value in some embodiments, within $\pm 5\%$ of a target value in some embodiments, and within $\pm 2\%$ of a target value in some embodiments. The terms “approximately” and “about” may equal the target value.

[0311] The term “substantially” if used herein may be construed to mean within 95% of a target value in some embodiments, within 98% of a target value in some embodiments, within 99% of a target value in some embodiments, and within 99.5% of a target value in some embodiments. In some embodiments, the term “substantially” may equal 100% of the target value.

What is claimed is:

1. A computerized system for training an inference model used to generate a musculoskeletal representation, the system comprising:

- a plurality of neuromuscular sensors configured to sense and record a plurality of neuromuscular signals from a user, wherein the plurality of neuromuscular sensors are arranged on one or more wearable devices structured to be worn by the user to obtain the plurality of neuromuscular signals;

- at least one camera configured to capture at least one image while the plurality of neuromuscular signals are sensed and recorded; and

- at least one computer processor programmed to:
 - determine, based on information obtained from the at least one image, position information describing a spatial relationship between two or more connected musculoskeletal segments of the user,
 - determine, based on the plurality of neuromuscular signals, force information,
 - associate the position information with the force information,

- train an inference model to output a musculoskeletal representation consistent with the position information, or the force information, or the position information and the force information when neuromuscular input signals provided to the inference model have at least one predetermined characteristic, to produce an updated inference model, and

- cause the updated inference model to be stored in a memory.
2. The computerized system of claim 1, wherein the at least one computer processor determines the position information based on the information obtained from the at least one image and information obtained from the plurality of neuromuscular signals.
 3. The computerized system of claim 1, wherein the at least one image comprises a series of time-sequenced images forming a video.
 4. The computerized system of claim 1, wherein the musculoskeletal representation corresponds to any one or any combination of:
 - a handstate of the user,
 - a still pose of the user, and
 - a dynamic gesture of the user.
 5. The computerized system of claim 1, wherein the at least one image includes any one or any combination of:
 - an image produced by visible light,
 - an image produced by infrared light,
 - an image produced by light of a predetermined range of wavelengths, and
 - an image produced by light of two or more different predetermined ranges of wavelengths.
 6. The computerized system of claim 1, wherein the at least one computer processor is programmed to:
 - determine a first quality of the plurality of neuromuscular signals, or a second quality of the at least one image, or both the first quality and the second quality, and
 - weight the plurality of neuromuscular signals, or the at least one image, or both the plurality of neuromuscular signals and the at least one image based on one or both of the first quality and the second quality, to train the inference model.
 7. The computerized system of claim 6, wherein the first quality of the plurality of neuromuscular signals is determined by determining whether at least one of the plurality of neuromuscular signals includes at least one signal artifact.
 8. The computerized system of claim 6, wherein the second quality of the at least one image is determined by determining whether a hand of the user is fully included in a field of view of the at least one camera, or whether the hand is fully occluded from the field of view of the at least one camera, or whether at least a portion of the hand of the user is occluded from the field of view of the at least one camera.
 9. The computerized system of claim 6, wherein the at least one computer processor is further programmed to determine the first quality in response to a determination that the second quality is greater than a threshold value.
 10. The computerized system of claim 1, wherein the at least one predetermined characteristic comprises any one or any combination of:
 - a presence of an artifact of the neuromuscular input signals,
 - a presence of a plurality of artifacts of the neuromuscular input signals,
 - a relative position of a plurality of artifacts of the neuromuscular input signals,
 - a range of amplitudes of the neuromuscular input signals, and
 - a periodic frequency of artifacts of the neuromuscular input signals.
 11. The computerized system of claim 1, further comprising:
 - at least one inertial measurement unit (IMU) sensor, wherein the at least one computer processor is further programmed to determine the position information based, at least in part, on IMU signals output from the at least one IMU sensor.
 12. The computerized system of claim 1, wherein the one or more wearable devices includes at least one positional marker included thereon, and wherein the position information is determined based, at least in part, on the at least one positional marker captured in the at least one image.
 13. A method of a computerized system for training an inference model used to generate a musculoskeletal representation, the method comprising:
 - receiving, by at least one processor, a plurality of neuromuscular signals of a user, the plurality of neuromuscular signals being sensed and recorded by a plurality of neuromuscular sensors arranged on one or more wearable devices worn by the user;
 - receiving, by the at least one processor, at least one image captured by at least one camera while the plurality of neuromuscular signals are sensed and recorded;
 - determining, by the at least one processor based on information obtained from the at least one image, position information describing a spatial relationship between two or more connected musculoskeletal segments of the user;
 - determining, by the at least one processor based on the plurality of neuromuscular signals, force information;
 - associating, by the at least one processor, the position information with the force information;
 - training, by the at least one processor, an inference model to output a musculoskeletal representation consistent with the position information, or the force information, or the position information and the force information when neuromuscular input signals provided to the inference model have at least one predetermined characteristic, to produce an updated inference model; and
 - causing, by the at least one processor, the updated inference model to be stored in a memory.
 14. The method of claim 13, further comprising:
 - determining, by the at least one processor, the position information based on the information obtained from the at least one image and information obtained from the plurality of neuromuscular signals.
 15. The method of claim 13, wherein the at least one image comprises a series of time-sequenced images forming a video.
 16. The method of claim 13, wherein the musculoskeletal representation corresponds to any one or any combination of:
 - a handstate of the user,
 - a still pose of the user, and
 - a dynamic gesture of the user.
 17. The method of claim 13, wherein the at least one image includes any one or any combination of:
 - an image produced by visible light,
 - an image produced by infrared light,
 - an image produced by light of a predetermined range of wavelengths, and
 - an image produced by light of two or more different predetermined ranges of wavelengths.

- 18.** The method of claim **13**, further comprising:
determining, by the at least one processor, a first quality of the plurality of neuromuscular signals, or a second quality of the at least one image, or both the first quality and the second quality, and
weighting, by the at least one processor, the plurality of neuromuscular signals, or the at least one image, or both the plurality of neuromuscular signals and the at least one image based on one or both of the first quality and the second quality, to train the inference model.
- 19.** The method of claim **18**, wherein the first quality of the plurality of neuromuscular signals is determined by determining whether at least one of the plurality of neuromuscular signals includes at least one signal artifact.
- 20.** The method of claim **18**, wherein the second quality of the at least one image is determined by determining whether a hand of the user is fully included in a field of view of the at least one camera, or whether the hand is fully occluded from the field of view of the at least one camera, or whether at least a portion of the hand of the user is occluded from the field of view of the at least one camera.
- 21.** The method of claim **18**, further comprising:
determining, by the at least one processor, the first quality in response to a determination that the second quality is greater than a threshold value.
- 22.** The method of claim **13**, wherein the at least one predetermined characteristic comprises any one or any combination of:
a presence of an artifact of the neuromuscular input signals,
a presence of a plurality of artifacts of the neuromuscular input signals,
a relative position of a plurality of artifacts of the neuromuscular input signals,
a range of amplitudes of the neuromuscular input signals, and
a periodic frequency of artifacts of the neuromuscular input signals.
- 23.** The method of claim **13**, further comprising:
receiving, by the at least one processor, inertial measurement unit (IMU) signals output from at least one IMU sensor; and
determining, by the at least one processor, the position information based, at least in part, on the IMU signals output from the at least one IMU sensor.
- 24.** The method of claim **13**,
wherein the one or more wearable devices includes at least one positional marker included thereon, and
wherein the position information is determined based, at least in part, on the at least one positional marker captured in the at least one image.
- 25.** A non-transitory computer-readable storage medium storing code that, when executed by at least one computer, causes the at least one computer to perform a method for training an inference model used to generate a musculoskeletal representation, wherein the method comprises:
receiving a plurality of neuromuscular signals of a user, the plurality of neuromuscular signals being sensed and recorded by a plurality of neuromuscular sensors arranged on one or more wearable devices worn by the user;
receiving at least one image captured by at least one camera while the plurality of neuromuscular signals are sensed and recorded;
determining, based on information obtained from the at least one image, position information describing a spatial relationship between two or more connected musculoskeletal segments of the user;
determining, based on the plurality of neuromuscular signals, force information;
associating the position information with the force information;
training an inference model to output a musculoskeletal representation consistent with the position information, or the force information, or the position information and the force information when neuromuscular input signals provided to the inference model have at least one predetermined characteristic, to produce an updated inference model; and
causing the updated inference model to be stored in a memory.
- 26.** The storage medium of claim **25**, wherein the method further comprises:
determining the position information based on the information obtained from the at least one image and information obtained from the plurality of neuromuscular signals.
- 27.** The storage medium of claim **25**, wherein the method further comprises:
determining a first quality of the plurality of neuromuscular signals, or a second quality of the at least one image, or both the first quality and the second quality, and
weighting the plurality of neuromuscular signals, or the at least one image, or both the plurality of neuromuscular signals and the at least one image based on one or both of the first quality and the second quality, to train the inference model.
- 28.** The storage medium of claim **27**, wherein the method further comprises:
determining the first quality in response to a determination that the second quality is greater than a threshold value.
- 29.** The storage medium of claim **25**, wherein the at least one predetermined characteristic comprises any one or any combination of:
a presence of an artifact of the neuromuscular input signals,
a presence of a plurality of artifacts of the neuromuscular input signals,
a relative position of a plurality of artifacts of the neuromuscular input signals,
a range of amplitudes of the neuromuscular input signals, and
a periodic frequency of artifacts of the neuromuscular input signals.
- 30.** The storage medium of claim **25**, wherein the method further comprises:
receiving inertial measurement unit (IMU) signals output from at least one IMU sensor; and
determining the position information based, at least in part, on the IMU signals output from the at least one IMU sensor.

专利名称(译)	相机指导的神经肌肉信号解释		
公开(公告)号	US20200069211A1	公开(公告)日	2020-03-05
申请号	US16/557427	申请日	2019-08-30
[标]发明人	BERENZWEIG ADAM REARDON THOMAS OSBORN CHRISTOPHER KAIFOSH PATRICK WETMORE DANIEL		
发明人	BERENZWEIG, ADAM REARDON, THOMAS OSBORN, CHRISTOPHER KAIFOSH, PATRICK JURMAN, BRETT WETMORE, DANIEL		
IPC分类号	A61B5/0488 G06F3/01 A61B5/00 A61B90/00 A61B5/04		
CPC分类号	A61B5/0488 A61B5/04004 G06F3/011 A61B5/6824 A61B90/361 A61B5/04001 A61B5/04012 G06T19/006 G06F3/017 A61B2562/0219 A61B5/681 A61B5/1107 A61B5/7267 G06F3/015 G06F3/016 G06K9/00355 G06K9/6215 G06N5/04 G06N20/00 G06T7/70 G06T2207/20081 G06T2207/30196		
优先权	62/726159 2018-08-31 US		
外部链接	Espacenet USPTO		

摘要(译)

提供了计算机化的系统，方法和存储用于实现所述方法的代码的计算机可读存储介质，以训练用于生成肌肉骨骼表示的推理模型。一个这样的系统包括处理器，该处理器被编程为：基于从至少一个图像获得的信息，确定描述用户的两个或更多个相连的肌肉骨骼段之间的空间关系的位置信息；以及基于多个神经肌肉信号，确定力信息；将位置信息与力信息相关联；当提供给推理模型的神经肌肉输入信号具有至少一个预定特性时，训练推理模型以输出与位置信息和/或力信息一致的肌肉骨骼表示，以产生更新的推理模型；并使更新后的推理模型存储在内存中。

