

AUTONOMOUS VIDEO CONFERENCING SYSTEM WITH VIRTUAL DIRECTOR ASSISTANCE

BACKGROUND OF THE DISCLOSURE

[0001] The present disclosure relates in general to video conferencing and remote collaboration technologies. Specifically, the present disclosure relates to powering video conferencing and remote collaboration with smart cameras and smart sensors using artificial intelligence. More specifically, the present disclosure relates to autonomous video conferencing systems with virtual director assistance and methods for creating an automated television studio production for a video conferencing space with virtual director assistance.

[0002] Remote collaboration and video conferencing has become a mainstay in the modern business world and societies at large. In traditional video conferencing, the experience for the participants is largely static. Cameras in meeting rooms are generally oblivious of social or space cues, such as speaker-listener dynamics, participants' reactions, body language, direction of attention, and extent of engagement. As a remote participant, therefore, the experience often lacks engagement. Existing hardware and software video conferencing solutions generally rely on a single camera system. Such a system is limited generally to displaying objects in the video conferencing space from one direction or one angle. It is thus challenging for the system to capture facial expressions and subtle changes in gesture and posture among objects variously located in the video conferencing space. As a result, remote participants are presented with a sub-optimal and less engaging experience. In a large video conferencing space, often it may be difficult additionally to capture and frame an object located far away from the camera, and thus difficult if not impossible for remote participants to see or read the object and engage with them.

[0003] More recent remote collaboration and video conferencing solutions have seen improvements that enable a remote user to adjust settings for their screen selection, image

filtering, recording, ambient audio control, file sharing, and privacy control preferences among other things. However, the lack of ability for a remote user to see or read the room and engage fully in real time with all objects in a conferencing space remains a drawback. Due to these limitations, therefore, despite the exponential growth of video conferencing solutions in recent years, in-person meetings are still favored in a variety of contexts if a user is unable or unwilling to compromise their ability to actively engage in a particular meeting or conference program.

[0004] There is therefore a need for smart video conferencing solutions capable of capturing and detecting all objects and their interactions as well as related changes in gesture and posture in the entire conferencing space. More specifically, there is a need to create a cohesive video conferencing production that resembles a television studio production for the entire video conference in real time and stream such production to a remote user thereby improving engagement of the remote user.

SUMMARY OF THE VARIOUS EMBODIMENTS

[0005] It is therefore an object of this disclosure to power video conferencing solutions with subsymbolic artificial intelligence thereby enabling the detection of objects and their interactions as well as related changes in gesture and posture. It is a further object of this disclosure to develop a rule set conforming to television production principles using symbolic artificial intelligence and create an automated video conferencing production for streaming to remote users in real time by applying such rule set.

[0006] Particularly, in accordance with this disclosure, there is provided, in one embodiment, an autonomous video conferencing system comprising a plurality of smart cameras. Each of the plurality comprises an image sensor adapted to capture video images in a video conferencing space and output an overview video stream and a focus video

stream, wherein the focus video stream comprises sub-video images framing detected objects within the overview video stream; an audio device adapted to capture audio signals in the video conferencing space; a vision pipeline unit adapted to receive said overview video stream and said audio signal, and detect objects and postures by means of machine learning, wherein the objects comprises persons and non-person items and said postures comprises positions, orientations, gestures, and directions of detected objects; a virtual director unit connected to the vision pipeline unit and the audio device and adapted to frame objects of interest according to a predetermined rule set thereby allowing the focus video stream to be updated in real time, wherein said predetermined rule set is adapted to create an automated television studio production from the video conferencing space; a stream selector connected to the virtual director unit and adapted to transition the focus video stream to an updated focus video stream; and an application programming interface connected to the virtual director unit and adapted to provide at least one connection to another smart camera in the plurality and one connection to a user computer.

[0007] According to one embodiment, one of the plurality of smart cameras is adapted as a main camera and each of the remaining of the plurality is adapted as a peripheral camera. The application programming interface of each peripheral camera is connected to the application programming interface of the main camera. The stream selector of each peripheral camera is further adapted to transmit its updated focus stream to the stream selector of the main camera, and the stream selector of the main camera is further adapted to select one of the updated focus streams from the plurality of smart cameras as an updated main focus stream and output the updated main focus stream to a user computer.

[0008] In another embodiment, the autonomous video conferencing system further comprises a plurality of smart sensors placed in the video conferencing space and adapted to capture non-image signals in the video conferencing space. Each of the plurality of smart

sensors comprises an application program interface that is adapted to connect to the application program interface of the main camera, thereby providing input to the virtual director unit of the main camera.

[0009] In yet another embodiment, each of the plurality of smart sensors is selected from the group consisting of a touchpad, a microphone, a smartphone, a GPS tracker, an echolocation sensor, a thermometer, a humidity sensor, and a biometric sensor.

[0010] In a further embodiment, the audio device comprises an array of microphones adapted to provide direction of audio with the captured audio signals.

[0011] In another embodiment, the vision pipeline unit comprises at least one hardware accelerated convolutional neural network. In yet another embodiment, the vision pipeline unit is pre-trained with a training set of video and audio data adapted to a special-purpose video conferencing space. In a further embodiment, the special-purpose video conferencing space is selected from the group consisting of a classroom, a workshop, a townhall, a newsroom, a boardroom, a courtroom, an interview studio, and a voting chamber.

[0012] In another embodiment, the predetermined rule set comprises a first rule for evaluating possible framing for each object in the video conferencing space based on a first plurality of parameters thereby determining the best frame. In yet another embodiment, the first plurality of parameters comprises: (i) whether the object is speaking; (ii) the length of speaking time; (iii) the direction of the object's gaze; (iv) the extent of the object's visibility in the frame; (v) the posture of the object; and (vi) what other objects are visible in the frame.

[0013] In a further embodiment, the predetermined rule set further comprises a second rule for detecting changes in the video conferencing space based on a second plurality of parameters to trigger transition of frames. In another embodiment, the second plurality of parameters comprises: (i) an object starts to speak; (ii) an object moves; (iii) an object

stands up; (iv) the direction of an object's gaze changes; (v) an object shows a reaction; (vi) an object displays a new item in the scene; (vii) an object has spoken for a predefined length of time; and (viii) lack of meaningful reactions in other objects for a predefined length of time.

[0014] In yet another embodiment, the predetermined rule set further comprises a third rule for applying suitable shot types to each frame based on a third plurality of parameters consistent with television studio production principles. In a further embodiment, the third plurality of parameters comprises: (i) a total shot to frame substantially all objects and most of the video conferencing space thereby providing an overall context to the video conferencing space; (ii) a medium shot to frame a predefined number of objects and focus on one who is speaking, thereby featuring an active dialog; and (iii) a close shot to frame one object speaking for a predefined length of time, thereby featuring a presenter.

[0015] In another embodiment, the third plurality of parameters further comprises (i) an interest shot to frame an object of interest based on cues of the scene in the video conferencing space, including an object at the center of the gaze from every object within the video conferencing space and an item held up by an object; (ii) a listening shot to frame at least one object who is not speaking, thereby featuring engagement of non-speaking objects in the video conferencing space; and (iii) a presenter shot to frame an object who has been speaking for a predefined longer length of time compared to other objects, thereby featuring the presenter from different camera angles and compositions within the video conferencing space. According to various embodiments, the interest shot is adapted as a close shot, the listening shot is adapted as one of a close and medium shot, and the presenter shot is adapted as one of a close and medium shot.

[0016] In another embodiment, the predetermined rule set further comprises a fourth rule for applying virtual Director's Cut based on a fourth plurality of parameters to the video

conferencing space thereby tailoring it to special-purpose television production scenarios. In yet another embodiment, the fourth plurality of parameters comprises: (i) a classroom production scenario starting with showing a presenter and an audience using total shots, then transitioning to framing the presenter in presenter shots for a predefined length of time, followed by switching between listening shots showing the audience and presenter shots showing the presenter; (ii) a workshop production scenario starting with showing all participating objects using total shots, then transitioning to showing an active object using medium shots, followed by showing an object moving in the video conferencing space using medium shots, and rotating back after a predefined length of time to showing active objects within the video conferencing space; and (iii) a meeting room production scenario starting with total shots creating an understanding of the entire video conferencing space with all visible objects, after a predefined length of time transitioning to framing a group of objects with medium shots in a sub-location of the video conferencing space focusing on an active object, followed by framing an object who is speaking at the sub-location using medium shots that best display the front of the object's face, after another predefined length of time switching to framing other objects in the video conferencing space using listening shots that best display the front of the object's faces, and rotating back to total shots featuring all objects if no object is speaking in the video conferencing space.

[0017] According to various embodiments, the active object comprises an object who is speaking, a whiteboard on which content is being drawn, and an item being held up by an object on demonstration, and the sub-location of the video conferencing space comprises a side of a table, a stage, a bench, a podium, and a row of chairs. In a further embodiment, the meeting room production scenario is further adapted to represent a production scenario selected from the group consisting of a news broadcast or podcast, a bilateral negotiation, a court proceeding, a panel discussion, and a voting assembly.

[0018] In another embodiment, the predetermined rule set further comprises a fifth rule for framing clean shots for objects within the virtual conferencing space. The fifth rule comprises not selecting any shot with objects partially visible, and aligning the eyes of an active object at the top one third of a frame. In yet another embodiment, the fifth rule further comprises adding space padding in the direction of an object's gaze, and using a medium shot to frame together active objects that are in proximity. According to various embodiments, the active objects in proximity comprises an object, a whiteboard, a display monitor, a lectern, a podium, and an item on demonstration.

[0019] In a further embodiment, the plurality of smart cameras further comprises at least one peripheral smart camera placed in a separate video conferencing space. The predetermined rule set in the main camera is further adapted to create an automated television studio production for a combined virtual conferencing space.

[0020] In another embodiment, the autonomous video conferencing system further comprises at least one smart sensor located in the separate video conferencing space. The smart sensor is adapted to capture non-image signals in the separate video conferencing space and provide input to the virtual director unit of the main camera.

[0021] In yet another embodiment, the application programming interface of the main camera is adapted to receive input from the user computer thereby modifying the predetermined rule set for the virtual director unit of the main camera.

[0022] In accordance with this disclosure, there is provided, in another embodiment, a method for creating an automated television studio production for a video conferencing space with virtual director assistance. The method comprises: capturing video images in the video conferencing space using a plurality of image sensors; capturing audio signals in the video conferencing space using a plurality of audio devices; generating an overview video stream and a focus video stream, wherein the focus video stream comprises sub-video

images framing detected objects within the overview video stream; detecting objects and postures from the overview stream and the audio signals by means of machine learning, wherein the objects comprises persons and non-person items and the postures comprises positions, orientations, gestures, and directions of detected objects; implementing a virtual director comprising a predetermined rule set that conforms to television studio production principles; applying the predetermined rule set to the detected objects and postures thereby framing objects of interest and updating the focus video stream in real time; and, outputting the updated focus video stream to a user computer.

[0023] In yet another embodiment, the machine learning is performed on at least one hardware accelerated convolutional neural network. In a further embodiment, the neural network is pre-trained with a training set of video and audio data adapted to a special-purpose video conferencing space. In another embodiment, the special-purpose video conferencing space is selected from the group consisting of a classroom, a workshop, a townhall, a newsroom, a boardroom, a courtroom, an interview studio, and a voting chamber.

[0024] According to various embodiments, the predetermined rule set comprises a first rule for evaluating possible framing for each object in the video conferencing space based on a first plurality of parameters thereby determining the best frame, a second rule for detecting changes in the video conferencing space based on a second plurality of parameters to trigger transition of frames, a third rule for applying suitable shot types to each frame based on a third plurality of parameters consistent with television studio production principles, a fourth rule for applying virtual Director's Cut based on a fourth plurality of parameters to the video conferencing space thereby tailoring it to special-purpose television production scenarios, and a fifth rule for framing clean shots for objects within the virtual conferencing space based on a fifth plurality of parameters.

[0025] In another embodiment, the shot types are selected from the group consisting of a total shot, a medium shot, a close shot, an interest shot, a listening shot, and a presenter shot. In yet another embodiment, the special-purpose television production scenarios comprise a classroom, a workshop, a meeting room, a broadcast, a bilateral negotiation, a court proceeding, a panel discussion, and a voting assembly.

[0026] In a further embodiment, the method for creating an automated television studio production further comprises capturing non-image signals in the video conferencing space using a plurality of smart sensors. Each of the plurality of smart sensors comprises an application program interface connected with the virtual director thereby providing input to the virtual director. In various embodiments, each of the plurality of smart sensors is selected from the group consisting of a touchpad, a microphone, a smartphone, a GPS tracker, an echolocation sensor, a thermometer, a humidity sensor, and a biometric sensor.

BRIEF DESCRIPTION OF THE DRAWINGS

[0027] Figure 1 depicts an autonomous video conferencing system according to one embodiment.

[0028] Figure 2 depicts an overview video stream and a focus video stream according to one embodiment.

[0029] Figure 3 illustrates a total shot, a medium shot, and a close shot according to one of the rules of the predetermined rule set in one embodiment.

[0030] Figure 4 illustrates framing clean shots according to one of the rules of the predetermined rule set in one embodiment.

[0031] Figure 5 illustrates alignment for framing clean shots according to one of the rules of the predetermined rule set in one embodiment.

[0032] Figure 6 shows two examples of a small video conferencing space for a special-purpose scenario in which an autonomous video conferencing system is deployed according to one embodiment.

[0033] Figure 7 shows two examples of a medium video conferencing space for a special-purpose scenario in which an autonomous video conferencing system is deployed according to one embodiment.

[0034] Figure 8 shows two examples of a large video conferencing space for a special-purpose scenario in which an autonomous video conferencing system is deployed according to one embodiment.

DETAILED DESCRIPTION OF THE VARIOUS EMBODIMENTS

[0035] Video conferencing systems and methods of this disclosure are powered by multiple smart cameras and smart sensors using subsymbolic and symbolic artificial intelligence. In various embodiments, the autonomous video conferencing systems of this disclosure include a plurality of smart cameras optionally coupled with a plurality of smart sensors located in a video conferencing space. Each smart camera is equipped with a vision pipeline supported by machine learning to detect objects and postures and discern speaker-listener dynamics, and a virtual director adapted to apply a predetermined rule set consistent with television studio production principles. Methods are meanwhile provided for creating an automated television studio production for a variety of conferencing spaces and special-purpose scenarios with virtual director assistance.

Pluralities of Smart Cameras & Sensors

[0036] Referring to Fig. 1, the autonomous video conferencing (“AVC”) system according to one embodiment comprises a plurality of smart cameras (100, 101) and optionally one or more smart sensors (201). One of the plurality of smart cameras is adapted as the main

camera (100) and the remaining of the plurality of smart cameras is adapted as the peripheral cameras (101). Each of the plurality of smart cameras comprises an image sensor (201, 301), an audio device (204, 304), a vision pipeline unit or vision pipeline (202, 302), a virtual director unit or virtual director (203, 303), a stream selector (205, 305), and an application programming interface ("API") (206, 306). The one or more smart sensors each comprises an API (401), capable of connecting to the API of a smart camera. According to one embodiment, the API of the main camera (100) is connected to the API of each peripheral camera (101) and the API of each smart sensor (102). The API of the main camera is further adapted to provide a connection to a user computer (103).

[0037] The plurality of smart cameras and one or more smart sensors in the AVC system are connected via Ethernet, other local area networks, or wireless networks in various embodiments. The main cameras and peripheral cameras are variously located in a video conferencing space to provide effective coverage of the video conferencing space according to one embodiment. In another embodiment, a plurality of smart sensors is activated and strategically located in a video conferencing space to capture non-image signals and provide input to the main camera of the AVC system. The smart sensors of this disclosure include touchpads, microphones, smartphones, GPS trackers, echolocation sensors, thermometers, humidity sensors, and biometric sensors in various embodiment.

[0038] In alternative embodiments, one or more peripheral cameras and smart sensors of the AVC system are placed in a separate video conferencing space as a secondary space for a video conference. These peripheral cameras and smart sensors are networked with the main camera and adapted to provide image and non-image input from the secondary space to the main camera. The AVC system of these alternative embodiments is thus adapted further to produce an automated television studio production for a combined video conferencing space based on input from all cameras and smart sensors in both spaces.

[0039] The smart camera of the AVC system is adapted with different degrees of field of view in various embodiments. Where a video conferencing space is small and the AVC system has a small plurality of cameras, for example, the smart cameras may have a wide field of view, e.g., for approximately 150 degrees. Where a video conferencing space is large and the AVC system has a large plurality of cameras, on the other hand, the smart cameras may have a narrower field of view, e.g., for approximately 90 degrees. The AVC system in another embodiment is equipped with smart cameras with various degrees of field of view, allowing optimal coverage for a video conferencing space. In a further embodiment, the image sensor (201, 301) is adapted to zoom up to 10X, enabling close-up images of objects at the far end of a video conferencing space. In alternative embodiments, one or more smart cameras in the AVC system are adapted to capture content on or about an object that is a non-person item within a video conferencing space, e.g., a whiteboard, a TV display, a poster, and a demonstration bench. These cameras may be smaller and placed differently than the other smart cameras in the AVC system and may be mounted in the ceiling to provide effective coverage of the target content.

[0040] In a further embodiment, the audio device (204, 304) in the smart camera is a microphone array adapted to take in audio signals from different locations around the camera. Using signals from different microphones allows the smart camera to determine the direction of audio ("DOA") and discern if there is silence in a location or direction. This information is in turn made available to the vision pipeline (202, 302) and the virtual director (203, 303) of the AVC system. In a further alternative embodiment, a computational device with high computing power is connected to the AVC system through an ethernet switch and adapted to provide additional computing power to the AVC system. It has one or more high performance CPUs and GPUs and runs parts of the vision pipeline for the main camera and any of the designated peripheral cameras.

[0041] The inclusion and configuration of multiple smart cameras and smart sensors in the AVC system therefore allows effective and robust coverage of a variety of video conferencing spaces and scenarios. By placing multiple smart cameras in a video conferencing space that collaborate to frame objects in the conference from different camera angles and zoom levels, the AVC system of this disclosure creates a more fulsome, natural, and engaging experience for all participants, including remote users.

Vision Pipeline; Overview & Focus Streams

[0042] Referring to Fig. 2, for each smart camera in the AVC system there are internally two video streams, an overview stream and a focus stream. The overview stream sees the entire scene, which is consumed by the vision pipeline (202, 302) as shown in Fig. 1. The focus stream is a high resolution stream framing objects of interest depending on the activities in a video conferencing space over time. This is where video settings are applied under the direction of the virtual director (203, 303), thereby transitioning the focus stream to an updated focus stream as discussed in detail below.

[0043] The vision pipeline of this disclosure is adapted to process the incoming overview stream and audio signals and detect objects and postures by means of machine learning. Objects include persons and non-person items, and postures include positions, orientations, gestures, and directions in various embodiments. The vision pipeline in one embodiment includes one or more hardware accelerated programmable convolutional neural networks, which adopts pretrained weights to allow detection of certain properties of objects in view. For example, the vision pipeline detects where in the view of a smart camera objects are, the extent of their visibility in the view, if they are speaking or not, their facial expressions, their body positions, and their head poses. The vision pipeline also tracks each object over time and determines where an object was previously in the field of view, if they are in motion or not, and in what direction they are heading.

[0044] One advantage of the vision pipeline utilizing subsymbolic artificial intelligence to detect objects and their activities and interactions according to various embodiments, is that these convolutional neural networks are trained to not be biased on properties such as gender, age, race, scene, light, and size. This enables the AVC system to create a more accurate and natural video stream production for the entire scene and all objects in a video conferencing space.

[0045] The vision pipeline of various embodiments is adapted to run on GPUs or other special purpose chipset with hardware accelerators for relevant mathematical operations in its convolutional neural networks. In an alternative embodiment, the vision pipeline runs on available CPU capacities in the AVC system. Further optimization of the vision pipeline in certain embodiments to fit its hardware chipset is made by replacing mathematical operations in its convolutional neural network architecture with equivalent mathematical operations supported by the chipset. The specialized hardware support of the vision pipeline therefore allows it to perform detections of objects and postures with high frequency and fast processing time, which in turn enables the AVC system to react to changes in the view of its smart cameras promptly and on demand.

[0046] In certain embodiments, the vision pipeline is pre-trained by running thousands of images and videos relevant to the scene and the detection objectives for a video conferencing space. During the training, the vision pipeline is evaluated with a loss function that measures how well it performs particular detections. The feedback from the loss function is then used to adjust the weights and parameters of the vision pipeline until it performs the particular detections with a predefined level of satisfaction. In one embodiment, the vision pipeline is further pre-trained and finetuned with a training set of video and audio data adapted to a special-purpose video conferencing space for a special-purpose scenario. For example, the vision pipeline may be finetuned for a classroom, a

workshop, a townhall, a newsroom, a boardroom, a courtroom, an interview studio, or a voting chamber to support special-purpose scenarios such as a lecture, an interview, a news broadcast or podcast, a court proceeding, a workshop, a bilateral negotiation, or a voting assembly.

[0047] The vision pipeline is further adapted to aggregate and process audio signals from microphones or other audio devices in the AVC system in various embodiments. It is capable of distinguishing voices, including if voices are raised or lowered by objects depending on what is happening in a video conferencing space. In one embodiment, the vision pipeline is adapted to classify the topic of a conversation based on audio signals. Voices that do not belong to an object is classified as artificial sound and may be attributed to other sources such as loudspeakers. The voice classification and characteristics are combined by the vision pipeline with other information and knowledge it detects and gleans from image data for relevant objects and their activities and interactions, thereby generating a comprehensive understanding of the video conferencing space and all detected objects. The vision pipeline makes available this corpus of a comprehensive understanding to the virtual director, which is in turn responsible for selecting best shots and creating an automated television studio production from the video conferencing space.

Virtual Director; Predetermined Rule Set

[0048] As discussed above, a virtual director (203, 303) of the AVC system is connected to the vision pipeline (202, 302) and the audio device (204, 304) as shown in Fig. 1, and adapted to frame objects of interest according to a predetermined rule set thereby allowing the focus stream to be updated in real time. The predetermined rule set is developed and implemented for the virtual director consistent with television studio production principles, as discussed in detail below. The transition of the focus stream to an updated focus stream

is carried out by the stream selector (205, 305) at the direction the virtual director in each smart camera.

[0049] Multiple smart cameras of the AVC system collaborate seamlessly to produce and update the focus streams for various locations and target objects within a video conferencing space. The main camera's virtual director (203) is connected with each peripheral camera's virtual director (303) through their respective APIs (206, 306) as shown in Fig. 1. The main camera's stream selector (205) is connected with each peripheral camera's stream selector (306) and adapted to consume the focus streams from the main camera and all peripheral cameras. The main camera's stream selector (205) is responsible further for selecting, at the direction of the main camera's virtual director (203), one of the updated focus streams from all cameras in the AVC system as the updated main focus stream. This updated main focus stream is the stream output made available to a user computer as shown in Fig. 1. The main camera's virtual director (203) is thus the brain or command center of the AVC system, and responsible for creating an automated television studio production for the video conferencing space based on input from all cameras and any smart sensors deployed in the AVC system.

[0050] The virtual director of this disclosure is a software component utilizing rule-based symbolic artificial intelligence to optimize decision making in various embodiments. The virtual director takes input from the vision pipeline and determines what frame from which camera in the AVC system should be selected and streamed to a user ultimately. In certain embodiments, the virtual director accomplishes this by evaluating possible framing for each object in the video conferencing space and their activities over time. For each object captured by a particular smart camera in the AVC system, for example, the virtual director evaluates different crops of images where the object is visible in order to find the best frame for the object in the context of their activities and interactions. The virtual director then

determines video settings suitable for transitioning the particular camera's focus stream to the selected best frame. The transitioning and updating of the focus stream using these video settings is then carried out by the stream selector at the direction the virtual director.

[0051] In certain embodiments, the virtual director of the main camera takes input additionally from one or more smart sensors deployed in the AVC system through their respective APIs. The input from smart sensors includes non-image signals or cues in general about the video conferencing space and all objects therein, e.g., objects' locations, movements, and physiological or biometric properties.

[0052] The APIs of the AVC system are adapted to transmit messages between various components over an internal network bus. These messages include information regarding the status of each smart camera, e.g., if it is connected, what type of software it is running, current health status etc. The APIs also communicate what a camera detects from image data, e.g., where in the image an object is detected, where in the conferencing space they are placed, and other information detected by the vision pipeline. The APIs in addition communicate video settings that are being applied to the focus stream, e.g., image properties, color, brightness, etc. Further, the APIs are adapted to communicate the virtual director parameters, which allows the AVC system to automatically set and adjust the virtual director rule set and related parameters for all its component cameras as well as allows users of the AVC system to personalize the virtual director experience by modifying certain parameters and rules for its predetermined rule set.

[0053] As discussed above, the virtual director of this disclosure employs rule-based decision-making utilizing symbolic artificial intelligence. A predetermined rule set is developed in various embodiments conforming to television studio production principles. This allows the AVC system to produce an automated television studio production of a conference experience that resembles a television production directed by real-world

professionals. In a certain embodiment, the predetermined rule set in each camera of the AVC system may be differently adapted for its location and target objects or content within a video conferencing space. In one embodiment, as shown in Fig. 1 the predetermined rule set applied by the virtual director (203) of the main camera includes a larger number of rules than the predetermined rule set applied by a peripheral camera's virtual director (303). In an alternative embodiment, a user is provided with the option to modify or tune certain rules and parameters of the predetermined rule set in an AVC system through the API of the main camera, which is connected to the main camera's virtual director unit (203) as shown in Fig. 1.

[0054] According to one embodiment, the predetermined rule set includes a first rule for evaluating possible framing for each object in the video conferencing space based on a first plurality of parameters. The first rule determines the best frame for each object. The first plurality of parameters in one embodiment includes (i) whether the object is speaking; (ii) the length of speaking time; (iii) the direction of the object's gaze; (iv) the extent of the object's visibility in the frame; (v) the posture of the object; and (vi) what other objects are visible in the frame.

[0055] The predetermined rule set includes a second rule in another embodiment for detecting changes in the video conferencing space based on a second plurality of parameters. The second rule triggers transition of frames. The second plurality of parameters include in one embodiment: (i) an object starts to speak; (ii) an object moves; (iii) an object stands up; (iv) the direction of an object's gaze changes; (v) an object shows a reaction; (vi) an object displays a new item in the scene; (vii) an object has spoken for a predefined length of time; and (viii) lack of meaningful reactions in the other objects for a predefined length of time.

[0056] The predetermined rule set includes a third rule in another embodiment for applying suitable shot types to each frame based on a third plurality of parameters. The third plurality of parameters in one embodiment includes: (i) a total shot to frame substantially all objects and most of the video conferencing space thereby providing an overall context to the video conferencing space, see Fig. 3, top frame; (ii) a medium shot to frame a predefined number of objects and focus on one who is speaking, thereby featuring an active dialog, see Fig. 3, middle frame; and (iii) a close shot to frame one object speaking for a predefined length of time, thereby featuring a presenter, see Fig. 3, bottom frame.

[0057] The third rule plurality of parameters in another embodiment further includes: (i) an interest shot to frame an object of interest based on cues of the scene in the video conferencing space, including an object at the center of the gaze from every object within the video conferencing space and an item held up by an object; (ii) a listening shot to frame at least one object who is not speaking, thereby featuring engagement of non-speaking objects in the video conferencing space; and (iii) a presenter shot to frame an object who has been speaking for the longest length of time compared to other objects, thereby featuring the presenter from different camera angles and compositions within a video conferencing space. In various embodiments, the interest shot is adapted as a close shot, the listening shot is adapted as a close shot or a medium shot, and the presenter shot is adapted as a close shot or a medium shot.

[0058] The predetermined rule set includes a fourth rule in another embodiment for applying virtual Director's Cut based on a fourth plurality of parameters to the video conferencing space. The fourth rule tailors video settings to special-purpose television production scenarios. The fourth plurality of parameters in one embodiment includes a classroom production scenario, a workshop production scenario, and a meeting room production scenario. The meeting room production scenario is further adapted in certain

embodiments for a news broadcast or podcast, a bilateral negotiation, a court proceeding, a panel discussion, and a voting assembly.

[0059] The pre-determined rule set includes a fifth rule in a further embodiment for framing clean shots for objects within the virtual conferencing space. The fifth rule in one embodiment includes not selecting any shot with objects partially visible. Referring to Fig. 4, for example, the left frames are clean shots, and the right frames are not clean shots and not selected in the AVC system under the fifth rule. The fifth rule in another embodiment includes aligning the eyes of an active object at the top one third of a frame. Referring to Fig. 5, for example, each of the images (top, middle, and bottom) is optimally aligned with the eyes of an active object at the top one third of the frame. These shots are clean shots framed and selected in the AVC system under the fifth rule.

[0060] The fifth rule in yet another embodiment includes adding space padding in the direction of an object's gaze, and using a medium shot to frame together active objects that are in proximity. The active objects in proximity include an object, a whiteboard, a display monitor, a lectern, a podium, and an item on demonstration in various embodiments.

Special-Purpose Scenarios

[0061] The AVC system of this disclosure is adapted for a variety of special-purpose video conferencing spaces for special-purpose video conferencing scenarios in various embodiments. The vision pipeline unit of the smart cameras in the AVC system is pre-trained according to certain embodiments with a training set of video and audio data adapted to a special-purpose video conferencing space, including a classroom, a workshop, a townhall, a newsroom, a boardroom, a courtroom, an interview studio, and a voting chamber.

[0062] As discussed above, among the predetermined rule set implemented and applied by the virtual director unit of the AVC system is a fourth rule that makes virtual Director's Cut

based on a fourth plurality of parameters. This fourth rule for virtual Director's Cut is designed to tailor a video conferencing space to special-purpose television production scenarios. The fourth plurality of parameters in one embodiment includes a classroom production scenario, a workshop production scenario, and a meeting room production scenario.

[0063] A classroom scenario in one embodiment starts with showing a presenter and an audience using total shots, then transitions to framing the presenter in presenter shots for a predefined length of time, thereafter turns to switching between listening shots showing the audience and presenter shots showing the presenter.

[0064] A workshop production scenario in one embodiment starts with showing all participating objects using total shots, then transitions to showing an active object using medium shots, thereafter turns to showing an object moving in the video conferencing space using medium shots, and finally rotates back after a predefined length of time to showing active objects within the video conferencing space.

[0065] A meeting room production scenario in one embodiment starts with total shots creating an understanding of the entire video conferencing space with all visible objects, after a predefined length of time transitions to framing a group of objects with medium shots in a sub-location of the video conferencing space focusing on an active object, thereafter turns to framing an object who is speaking at the sub-location using medium shots that best display the front of the object's face, after another predefined length of time then switches to framing other objects in the video conferencing space using listening shots that best display the front of the object's faces, and finally rotates back to total shots featuring all objects if no object is speaking in the video conferencing space. An active object in various embodiments represents an object who is speaking, a whiteboard on which content is being drawn, or an item being held up by an object on demonstration. The sub-location of the

video conferencing space in various embodiments represents a side of a table, a stage, a bench, a podium, or a row of chairs.

[0066] The meeting room production scenario in certain embodiments is further adapted to represent a news broadcast or podcast, an interview, a board meeting, a bilateral negotiation, a court proceeding, a panel discussion, and a voting assembly.

[0067] For various special-purpose scenarios, the main camera and the peripheral cameras are strategically and variously placed within the video conferencing space to provide effective coverage of the space. One or more optional smart sensors are additionally dispensed within the video conferencing space with connection to the main camera and provide input to the virtual director unit of the main camera. Examples of AVC systems of this disclosure deployed for special-purpose scenarios are shown in Figures 6-8. In these drawings each dot represents a smart camera or a smart sensor of an AVC system and each circle represents an object, including a person or a non-person item such as a chair or a demonstration subject. The small rectangles along the sides of each of these drawings represents a TV, a whiteboard, a poster, or a projection display (601, 602, 701, 702, 703, 704, 801, 802). The square or rectangles in the middle of the drawings represent a table.

[0068] Referring to Fig. 6, both the top and bottom depict a small video conferencing space in which an AVC system is deployed. The top configuration represents a small meeting and a panel discussion scenario in one embodiment. The bottom configuration represents a news broadcast or podcast and an interview scenario in one embodiment.

[0069] Referring to Fig. 7, both the top and bottom depict a medium video conferencing space in which an AVC system is deployed. The top configuration represents a board meeting scenario in one embodiment. The bottom configuration represents a bilateral negotiation in one embodiment.

[0070] Referring to Fig. 8, both the top and bottom depict a large video conferencing space in which an AVC system is deployed. The top configuration represents a workshop scenario in one embodiment. The bottom configuration represents a classroom scenario in one embodiment.

[0071] The descriptions of the various embodiments, including the drawings and examples, are to exemplify and not to limit the invention and the various embodiments thereof.

WHAT IS CLAIMED IS:

1. An autonomous video conferencing system comprising a plurality of smart cameras, each of said plurality comprising:
 - an image sensor adapted to capture video images in a video conferencing space and output an overview video stream and a focus video stream, wherein said focus video stream comprises sub-video images framing detected objects within said overview video stream;
 - an audio device adapted to capture audio signals in said video conferencing space;
 - a vision pipeline unit adapted to receive said overview video stream and said audio signal, and detect objects and postures by means of machine learning, wherein said objects comprises persons and non-person items and said postures comprises positions, orientations, gestures, and directions of detected objects;
 - a virtual director unit connected to said vision pipeline unit and said audio device and adapted to frame objects of interest according to a predetermined rule set thereby allowing said focus video stream to be updated in real time, wherein said predetermined rule set is adapted to create an automated television studio production from said video conferencing space;
 - a stream selector connected to said virtual director unit and adapted to transition said focus video stream to an updated focus video stream; and
 - an application programming interface connected to said virtual director unit and adapted to provide at least one connection to another smart camera in said plurality and one connection to a user computer,

wherein one of said plurality of smart cameras is adapted as a main camera and each of the remaining of said plurality is adapted as a peripheral camera, wherein the application programming interface of each peripheral camera is connected to the application programming interface of said main camera, wherein the stream selector of each peripheral camera is further adapted to transmit its updated focus stream to the stream selector of the main camera, and wherein the stream selector of the main camera is further adapted to select one of the updated focus streams from said plurality of smart cameras as an updated main focus stream and output said updated main focus stream to a user computer.

2. The system of claim 1, further comprising a plurality of smart sensors placed in said video conferencing space and adapted to capture non-image signals in said video conferencing space, each of said plurality of smart sensors comprising an application program interface, wherein said application program interface is adapted to connect to the application program interface of the main camera, thereby providing input to the virtual director unit of the main camera.
3. The system of claim 2, wherein each of the plurality of smart sensors is selected from the group consisting of a touchpad, a microphone, a smartphone, a GPS tracker, an echolocation sensor, a thermometer, a humidity sensor, and a biometric sensor.
4. The system of claim 1, wherein the audio device comprises an array of microphones adapted to provide direction of audio with the captured audio signals.

5. The system of claim 1, wherein the vision pipeline unit comprises at least one hardware accelerated convolutional neural network.
6. The system of claim 5, wherein the vision pipeline unit is pre-trained with a training set of video and audio data adapted to a special-purpose video conferencing space.
7. The system of claim 6, wherein the special-purpose video conferencing space is selected from the group consisting of a classroom, a workshop, a townhall, a newsroom, a boardroom, a courtroom, an interview studio, and a voting chamber.
8. The system of claim 1, wherein the predetermined rule set comprises a first rule for evaluating possible framing for each object in the video conferencing space based on a first plurality of parameters thereby determining the best frame.
9. The system of claim 8, wherein the first plurality of parameters comprises: (i) whether the object is speaking; (ii) the length of speaking time; (iii) the direction of the object's gaze; (iv) the extent of the object's visibility in the frame; (v) the posture of the object; and (vi) what other objects are visible in the frame.
10. The system of claim 8, wherein the predetermined rule set further comprises a second rule for detecting changes in the video conferencing space based on a second plurality of parameters to trigger transition of frames.
11. The system of claim 10, wherein the second plurality of parameters comprises:
(i) an object starts to speak; (ii) an object moves; (iii) an object stands up;
(iv) the direction of an object's gaze changes; (v) an object shows a reaction; (vi) an object displays a new item in the scene; (vii) an object has

spoken for a predefined length of time; and (viii) lack of meaningful reactions in other objects for a predefined length of time.

12. The system of claim 10, wherein the predetermined rule set further comprises a third rule for applying suitable shot types to each frame based on a third plurality of parameters consistent with television studio production principles.
13. The system of claim 12, wherein the third plurality of parameters comprises:
 - (i) a total shot to frame substantially all objects and most of the video conferencing space thereby providing an overall context to the video conferencing space; (ii) a medium shot to frame a predefined number of objects and focus on one who is speaking, thereby featuring an active dialog; and (iii) a close shot to frame one object speaking for a predefined length of time, thereby featuring a presenter.
14. The system of claim 13, wherein the third plurality of parameters further comprises (i) an interest shot to frame an object of interest based on cues of the scene in the video conferencing space, including an object at the center of the gaze from every object within the video conferencing space and an item held up by an object; (ii) a listening shot to frame at least one object who is not speaking, thereby featuring engagement of non-speaking objects in the video conferencing space; and (iii) a presenter shot to frame an object who has been speaking for the longest length of time compared to other objects, thereby featuring the presenter from different camera angles and compositions within the video conferencing space, wherein said interest shot is adapted as a close shot, said listening shot is adapted as one

of a close and medium shot, and said presenter shot is adapted as one of a close and medium shot.

15. The system of claim 12, wherein the predetermined rule set further comprises a fourth rule for applying virtual Director's Cut based on a fourth plurality of parameters to the video conferencing space thereby tailoring it to special-purpose television production scenarios.
16. The system of claim 15, wherein the fourth plurality of parameters comprises:
 - (i) a classroom production scenario starting with showing a presenter and an audience using total shots, then transitioning to framing the presenter in presenter shots for a predefined length of time, followed by switching between listening shots showing the audience and presenter shots showing the presenter; (ii) a workshop production scenario starting with showing all participating objects using total shots, then transitioning to showing an active object using medium shots, followed by showing an object moving in the video conferencing space using medium shots, and rotating back after a predefined length of time to showing active objects within the video conferencing space; and (iii) a meeting room production scenario starting with total shots creating an understanding of the entire video conferencing space with all visible objects, after a predefined length of time transitioning to framing a group of objects with medium shots in a sub-location of the video conferencing space focusing on an active object, followed by framing an object who is speaking at the sub-location using medium shots that best display the front of the object's face, after another predefined length of time switching to framing other objects in the video conferencing space using listening shots that best display the front of the object's faces,

and rotating back to total shots featuring all objects if no object is speaking in the video conferencing space,
wherein said active object comprises an object who is speaking, a whiteboard on which content is being drawn, and an item being held up by an object on demonstration, and wherein the sub-location of the video conferencing space comprises a side of a table, a stage, a bench, a podium, and a row of chairs.

17. The system of claim 16, wherein the meeting room production scenario is further adapted to represent a production scenario selected from the group consisting of a news broadcast or podcast, a bilateral negotiation, a court proceeding, a panel discussion, and a voting assembly.
18. The system of claim 15, wherein the predetermined rule set further comprises a fifth rule for framing clean shots for objects within the virtual conferencing space, said fifth rule comprising: not selecting any shot with objects partially visible; and, aligning the eyes of an active object at the top one third of a frame.
19. The system of claim 18, wherein the fifth rule further comprises adding space padding in the direction of an object's gaze; and, using a medium shot to frame together active objects that are in proximity, wherein the active objects in proximity comprises an object, a whiteboard, a display monitor, a lectern, a podium, and an item on demonstration.
20. The system of claim 1 wherein the plurality of smart cameras further comprises at least one peripheral smart camera placed in a separate video conferencing space, wherein the predetermined rule set in the main camera

is further adapted to create an automated television studio production for a combined virtual conferencing space.

21. The system of claim 20, further comprising at least one smart sensor located in said separate video conferencing space, wherein said smart sensor is adapted to capture non-image signals in said separate video conferencing space and provide input to the virtual director unit of the main camera.
22. The system of claim 1, wherein the application programming interface of the main camera is adapted to receive input from the user computer thereby modifying the predetermined rule set for the virtual director unit of the main camera.
23. A method for creating an automated television studio production for a video conferencing space with virtual director assistance, comprising: capturing video images in said video conferencing space using a plurality of image sensors; capturing audio signals in said video conferencing space using a plurality of audio devices; generating an overview video stream and a focus video stream, wherein said focus video stream comprises sub-video images framing detected objects within said overview video stream; detecting objects and postures from said overview stream and said audio signals by means of machine learning, wherein said objects comprises persons and non-person items and said postures comprises positions, orientations, gestures, and directions of detected objects; implementing a virtual director comprising a predetermined rule set that conforms to television studio production principles; applying said predetermined rule set to said detected objects and postures thereby framing objects of interest

and updating said focus video stream in real time; and, outputting said updated focus video stream to a user computer.

24. The method of claim 23, wherein said machine learning is performed on at least one hardware accelerated convolutional neural network.
25. The method of claim 24, wherein said neural network is pre-trained with a training set of video and audio data adapted to a special-purpose video conferencing space.
26. The method of claim 25, wherein said special-purpose video conferencing space is selected from the group consisting of a classroom, a workshop, a townhall, a newsroom, a boardroom, a courtroom, an interview studio, and a voting chamber.
27. The method of claim 23, wherein the predetermined rule set comprises a first rule for evaluating possible framing for each object in the video conferencing space based on a first plurality of parameters thereby determining the best frame.
28. The method of claim 27, wherein the predetermined rule set further comprises a second rule for detecting changes in the video conferencing space based on a second plurality of parameters to trigger transition of frames.
29. The method of claim 28, wherein the predetermined rule set further comprises a third rule for applying suitable shot types to each frame based on a third plurality of parameters consistent with television studio production principles.
30. The method of claim 29, wherein said shot types are selected from the group consisting of a total shot, a medium shot, a close shot, an interest shot, a listening shot, and a presenter shot.

31. The method of claim 30, wherein the predetermined rule set further comprises a fourth rule for applying virtual Director's Cut based on a fourth plurality of parameters to the video conferencing space thereby tailoring it to special-purpose television production scenarios.
32. The method of claim 31, wherein said special-purpose television production scenarios comprise a classroom, a workshop, a meeting room, a broadcast, a bilateral negotiation, a court proceeding, a panel discussion, and a voting assembly.
33. The method of claim 32, wherein the predetermined rule set further comprises a fifth rule for framing clean shots for objects within the virtual conferencing space based on a fifth plurality of parameters.
34. The method of claim 23, further comprising capturing non-image signals in said video conferencing space using a plurality of smart sensors, each of said plurality of smart sensors comprising an application program interface connected with said virtual director thereby providing input to said virtual director.
35. The method of claim 34, wherein each of the plurality of smart sensors is selected from the group consisting of a touchpad, a microphone, a smartphone, a GPS tracker, an echolocation sensor, a thermometer, a humidity sensor, and a biometric sensor.