



(51) International Patent Classification:

G06T 7/50 (2017.01) G06T 7/70 (2017.01)
G06N 3/04 (2006.01) H04N 13/122 (2006.01)
G06T 1/40 (2006.01) H04N 13/388 (2018.01)

(21) International Application Number:

PCT/IB2020/052936

(22) International Filing Date:

27 March 2020 (27.03.2020)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

3,046,612 14 June 2019 (14.06.2019) CA

(71) Applicant: WRNCH INC. [CA/CA]; 1001 Lenoir, Suite B113, Montreal, Québec H4Z 2Z6 (CA).

(72) Inventors: BROWN, Colin; 1001 Lenoir, Suite B113, Montreal, Québec H4Z 2Z6 (CA). HARBOUR, Louis; 1001 Lenoir, Suite B113, Montreal, Québec H4Z 2Z6 (CA).

(74) Agent: MAR, Richard et al.; 100 King St. W., Suite 6000, Toronto, Ontario M5X 1E2 (CA).

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,

(54) Title: METHOD AND SYSTEM FOR MONOCULAR DEPTH ESTIMATION OF PERSONS




FIG. 1

(57) Abstract: Systems and methods are provided for estimating the 3D joint location of skeleton joints from an image segment of an object and a 2D joint heatmaps comprising 2D locations of skeleton joints on the image segment. This includes applying the image segment and 2D joint heatmaps to a convolutional neural network containing at least one 3D convolutional layer block, wherein the 2D resolution is reduced at each 3D convolutional layer and the depth resolution is expanded to produce an estimated depth for each joint. Combining the 2D location of each kind of joint with the estimated depth of the kind of joint generates an estimated 3D joint position of the skeleton joint.

TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Published:

- *with international search report (Art. 21(3))*
- *in black and white; the international application as filed
contained color or greyscale and is available for download
from PATENTSCOPE*

METHOD AND SYSTEM FOR MONOCULAR DEPTH ESTIMATION OF PERSONS

FIELD

[0001] This disclosure relates to methods and systems for the estimation of relative 3D positions of persons' joints in image or video data from a monocular camera. In particular, it relates to a machine learning-based approach that estimates relative depth values of one or more persons' joints from a camera.

BACKGROUND

[0002] Monocular estimation of relative 3D positions or depth values, namely the distances from the camera, of human joints, or any anatomical landmarks such as elbows, eyes, toes, etc., has broad ranging applications including motion capture, virtual and augmented reality, sports analytics, personal training and activity recognition.

[0003] In particular, 3D joint estimation from standard RGB image data with no additional depth sensing hardware facilitates the broadest set of use cases. Depth sensing hardware is often expensive and limiting in its use, such as having a limited range and may only be used inside. However, deriving 3D positions without additional hardware is challenging since using input data from a single RGB image excludes the possibility of triangulation.

[0004] Mathematically, determining the depth of a human joint from a 2D position in a single image is ill-posed, but methods that use prior and contextual information to tackle this issue have been proposed. Modern methods typically use machine learning-based models and largely fall into two categories. First, lifting methods accept previously detected 2D joint positions (or skeletons) of a person and generate 3D positions corresponding to the most likely pose, typically using a fully-connected neural network model. Secondly, convolutional neural network (CNN)-based methods use learned 2D convolutions to extract features from an input image containing a person that are fed into a regression model on the 3D positions or depths.

[0005] Lifting approaches are limited by the fact that 2D skeleton data is missing visual cues such as shadows, occlusions, lighting and volume (e.g. apparent scale of different joints) that may be critical for correctly determining the relative depth of a joint. While CNN-based approaches have access to such cues, extracted image features are inherently 2D (since the 2D convolutions operate on 2D maps) in contrast to the 3D nature of the output. Furthermore, CNN-approaches don't typically have a priori knowledge of 2D joint positions, and instead must infer these positions while inferring the associated depths.

[0006] In part due to these limitations, existing approaches often exhibit limited accuracy and robustness, especially in challenging poses, environments and conditions. Thus, there exists the need for a better approach to monocular RGB estimation of human joint depths.

SUMMARY

[0007] Systems and methods are provided for estimating the 3D joint location of skeleton joints from an image segment of an object and a 2D joint heatmaps comprising 2D locations of skeleton joints on the image segment. This includes applying the image segment and 2D joint heatmaps to a convolutional neural network containing at least one 3D convolutional layer block, wherein the 2D resolution is reduced at each 3D convolutional layer and the depth resolution is expanded to produce an estimated depth for each joint. Combining the 2D location of each kind of joint with the estimated depth of the kind of joint generates an estimated 3D joint position of the skeleton joint.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] In drawings which illustrate by way of example only a preferred embodiment of the disclosure,

[0009] Figure 1 is a representation of the high-level architecture of an embodiment of the 3D pose estimation system.

[0010] Figure 2 is a representation of an example of a reduction in spatial dimension sizes and an increase in the depth dimension size via 3D CNN blocks.

[0011] Figure 3 is a schematic representation of an embodiment of a depth heatmap estimator showing one 2D CNN block and four 3D CNN blocks.

[0012] Figure 4 is a schematic representation of an embodiment of a depth heatmap estimator showing four 3D CNN blocks and no 2D CNN blocks.

[0013] Figure 5 is a schematic representation of an embodiment of a depth heatmap estimator with a 2D CNN block of two residual units and five 3D CNN blocks.

[0014] Figure 6 is a schematic representation of an embodiment of a depth heatmap estimator with one 2D CNN block and three 3D CNN blocks, each containing one convolutional layer.

[0015] Figure 7 is a representation of an example input image, a colorized representation of the corresponding input heatmaps and an example of the corresponding 3D joint positions, portrayed as a skeleton.

[0016] Figure 8 is an example of a plurality of output depth heatmaps (depth range is horizontal) corresponding to a plurality of kinds of joints to be localized.

DETAILED DESCRIPTION

[0017] With reference to Figure 1, the human joint 3D estimation system 110 may comprise a capture device 120 to capture a digital image, a 2D skeleton detector 130, which may detect and crop humans in an image and localize human joint positions in the image, a 2D joint heatmap generator 140, which may generate a positional heatmap for each kind of joint to be localized, a depth heatmap estimator 150, which may generate 1D depth heatmaps for each kind of localized joint, and a 3d joint constructor 160, which may select a depth value for each kind of joint from the depth heatmaps and combine this information with the 2D joint positions to produce 3D joint positions.

[0018] A capture device 120 may represent any device or method for acquiring a digital image or video frame, such as a digital camera, webcam, mobile device, video camera, streaming source (e.g. videos or images streamed from the web) or stored image content such as an image or video database. The image format may be a 3-channel RGB or be encoded in some other color space (such as sRGB, CMYK). The image may or may not include additional image channels, such as from depth sensors or infra-red sensors.

[0019] A 2D skeleton detector 130 may detect one or more persons in an image. For each detected person, the 2D skeleton detector may output a cropped image centered around that person. Each output cropped image may be resized to a specific resolution (e.g. 256x256 pixels), which may be independent of the input resolution or size of the detected person.

[0020] For each detected person, a 2D skeleton detector may also output the detected positions, such as the 2D co-ordinates, of a pre-defined plurality of kinds of human joints. The joints may be anatomical landmarks. For instance, the 2D joint detector may localize the positions of all visible instances of key joints. This may include 22 kinds of joints, including left wrist, right wrist, left knee, base of neck. The 2D joint detector may provide these detected joints as a plurality of co-ordinates, grouped by detected person. The co-ordinates of the joints for each detected person may be provided in one array of numbers.

[0021] A 2D skeleton detector may be embodied by an existing software module such as from OpenPose, or wrnchAI or some other software or hardware modules that can detect persons and localize human joints in an input image.

[0022] A 2D joint heatmap generator 140 may accept a plurality of detected joint positions for one person and produce a heatmap image for each kind of joint. A heatmap image for one kind of joint may convey the locations of a plurality of detected joints of that kind as a corresponding plurality of markers in the image localized at those 2D joint positions. A marker may be a Gaussian kernel. For example, the location may include a standard deviation of 8 pixels or some other size. The location may be in the form of a flat disk or any other localized kernel image.

[0023] A 2D joint heatmap generator 140 may instead output 2D joint heatmaps generated directly by the 2D skeleton detector 130, depending on the available outputs of that module. For example some 2D skeleton detectors may output 2D joint heatmaps directly. In this case the 2D joint heatmap generator may be a part or submodule of the 2D skeleton detector.

[0024] Each heatmap image output by the 2D joint heatmap generator, corresponding to one kind of joint, may be encoded in a single channel. The value of each pixel of that channel may approximate the likelihood of a joint of a given kind being detected at that pixel. The resolution of the heatmap images may be the same resolution as the cropped image of a single person detected by the 2D skeleton detector, such as 256x256 pixels.

[0025] A depth heatmap estimator 150, may accept as input an image centered on the likeness of one person along with 2D heatmap images representing the locations of a predefined plurality of kinds of joints of that person and output a plurality of 1D depth heatmaps. The number of depth heatmaps may correspond to the number of input 2D heatmaps. This may also correspond to the number of kinds of joints detected by the 2D skeleton detector. Each depth heatmap may represent a range of depth values of one kind of joint relative to some fixed point. This fixed point may be the center of the hips of the person, the furthest joint from the camera, or some other point of reference. With reference to Figure 8, a depth heatmap may be stored as 1D array of pixels (e.g. 256 pixels), with each pixel corresponding to a relative depth value. The value of each pixel of each depth heatmap may represent the likelihood of the corresponding joint being at that relative depth.

[0026] A depth heatmap estimator may be embodied by a trained machine learning model. This model may be a convolutional neural network (CNN), or alternatively a RNN, random forest, or deep neural network. The architecture of the CNN model may contain 3D convolutional layers that may extract 3D features (such as object boundaries, occlusions, spatial orderings, texture and colour patterns, lighting features) from the input data. The CNN model may also contain 2D convolution layers, residual units, restricted linear activation units (ReLU, or other activations such as sigmoid units), batch normalization layers, max pooling layers and reshape layers. These

layers may be organized into logical blocks such as 2D CNN blocks and 3D CNN blocks, chained end-to-end.

[0027] The CNN may be organized in a pyramidal fashion such that 2D resolution is reduced at each block and depth resolution is expanded correspondingly, extracting 3D features with increasing fidelity while maintaining an efficient and bounded use of memory.

[0028] A 2D CNN block may extract low-level image features and expand the receptive field by performing an initial reduction in spatial resolution (i.e. pooling). A 2D CNN block may comprise a 2D convolutional layer, a residual unit, a max pooling layer (for example with a 2x2x1 kernel size) and a final 2D convolutional layer. Batch normalization layers may be inserted after each convolutional layer to impose batch normalization. The convolution layer may for example have a 7x7 kernel size and stride 2. The residual unit may have 3 convolutional layers with a skip connection.

[0029] The importance of the inclusion of a 2D CNN block may depend on memory and computational considerations as the reduction in spatial resolution performed by this block can reduce the memory and computational cost of subsequent layers, for example 3D convolution layers that may be relatively computationally expensive.

[0030] A 3D CNN block may comprise a 3D convolution, and ReLU, followed by another 3D convolution and ReLU, followed by a max pooling layer and a reshape layer. Batch normalization layers may be inserted after each convolutional layer to impose batch normalization. Each 3D convolution may have a kernel size of 3x3x3 to reduce the computational cost and number of learnable parameters required for each subsequent layer.

[0031] Each max pooling layer of the 3D CNN block may reduce the size of the input spatial dimensions, H and W, for example by a factor of 2 each, but maintain the depth dimension, D. Each reshape layer may reduce the number of channels produced by the pooling layer and increase the size of the depth dimension accordingly, for example by a factor of 4. Thus, with reference to Figure 2, each 3D CNN block may effectively reduce the H and W dimensions while expanding the D dimension proportionally.

[0032] With reference to Figure 3, the depth heatmap estimator may comprise a 2D CNN block, one or more 3D CNN blocks, for example 4 blocks, followed by one (or more) final 3D convolutional layers and a final max pooling layer. The final 3D convolutional layers may output a number of channels equal to the number of kinds of joints to be estimated, for example 22. The

final max pooling layer may reduce the original spatial dimensions, W and H down to 1 each, leaving only an expanded depth dimension for each joint channel.

[0033] The architecture of a depth heatmap estimator may vary based on specific requirements of the task including memory constraints, runtime constraints, training database size, and other similar factors. For example, with reference to Figure 4, the depth heatmap estimator may contain no 2D blocks. With reference to Figure 5, 2D CNN blocks may comprise more than one residual unit and the depth heatmap estimator may contain more than four 3D blocks. With reference to Figure 6, the depth heatmap estimator may contain fewer than four 3D blocks and each 3D block may contain different numbers of stacked 3D convolutional layers.

[0034] More 3D convolutional layers in each block may enable the extraction of more complex 3D features but at the expense of more required computations and a larger memory footprint during training. Thus the specific number of convolutional layers per block in a given embodiment may depend on requirements of the application.

[0035] For ease of reference, in Figure 3, Figure 4, Figure 5 and Figure 6, ReLU and batch normalization layers are not shown although may be present.

[0036] Learnable parameters of the depth heatmap estimator may be trained with an appropriate training set. The training may use back propagation, or another appropriate learning algorithm. A solver method such as RMSProp, or gradient descent, ADAM or similar may be used. A training framework such as TensorFlow, Caffe, or PyTorch or similar framework may be used. The training set may comprise a set, such as thousands, of images of persons in different poses, settings and lighting conditions, each associated with labelled 3D joint positions. As an example the COCO dataset may be used.

[0037] The 3D joint constructor 160 may accept 2D joint positions for one detected person, for example from the 2D skeleton detector 130, and corresponding joint depth heatmaps, produced by the depth heatmap estimator 150, and output 3D joint positions. A depth component of a 3D joint position may be computed from a depth heatmap via the argmax function. Alternatively, a weighted sum function or a peak detection technique may be used. Input 2D joint positions and computed depth values may be combined to produce 3D joint positions.

[0038] An example of a plurality of 3D joint positions generated from one embodiment of the human joint 3D estimation system is portrayed in Figure 7. An input image representing a human figure 710 is provided along with a 2D joint heatmap 720 such as generated by a 2D joint heatmap generator 140. Although not represented in the heatmap 720, each kind of joint, such as

left wrist, would have a separate heatmap. The heatmaps for all detected joints are overlaid in the colorized representation 720 with right joints colorized in red and left joints colorized in blue.

[0039] An example of depth heatmaps produced by a depth heatmap estimator is portrayed in Figure 8. The depth estimator generates a separate depth heatmap for each kind of joint. The depth range for each joint is indicated in the horizontal axis. The value (i.e. darkness) of each point along a given heatmap may represent the confidence of the corresponding joint being at that depth. The confidence values of all depths may be zero for a given joint, which may correspond to a joint that is not visible. An estimated depth value may be derived by aggregating the confidence values at different depths.

[0040] By combining the 2D joint heatmap with the depth estimation for each kind of joint, a 3D location for each joint can be estimated. A rendering 740 of a skeleton based on these 3D locations is shown.

[0041] These described system modules may be separate software modules, separate hardware units or portions of one or more software or hardware components. For example, the software modules may consist of instructions written in a computer language such as the Python programming language, C++ or C# with suitable modules, such as created using software from Caffe, TensorFlow, or Torch, and run on computer hardware, such as a CPU, GPU or implemented on an FPGA. The system may be run on desktop, mobile phone or other platforms such as part of an embedded systems that includes suitable memory for holding the software and activity data. The system may be integrated with or connect to the capture device.

[0042] Various embodiments of the present disclosure having been thus described in detail by way of example, it will be apparent to those skilled in the art that variations and modifications may be made without departing from the disclosure. The disclosure includes all such variations and modifications as fall within the scope of the appended claims.

CLAIMS:

1. A method for estimating the 3D joint location of skeleton joints from an image segment of an object and a 2D joint heatmaps comprising 2D locations of skeleton joints on the image segment comprising:

applying the image segment and 2D joint heatmaps to a convolutional neural network containing at least one 3D convolutional layer block, wherein the 2D resolution is reduced at each 3D convolutional layer and the depth resolution is expanded to produce an estimated depth for each joint;

combining the 2D location of each kind of joint with the estimated depth of the kind of joint to generate an estimated 3D joint position of the skeleton joint.

2. The method of claim 1 further comprising applying the image segment and 2D joint heatmap to at least one 2D convolutional neural network block prior to the at least one 3D convolutional layers.

3. The method of claim 1 or 2 further comprising a max pooling layer that receives the depths from the at least one 3D convolutional layers and the max pooling layer reduces the 2D resolution to one and leaves only an expanded depth dimension for each joint prior to combining the 2D location of each kind of joint with the estimated depth of the kind of joint.

4. The method of any one of claims 1 to 3 wherein each of the 3D convolutional layer blocks comprises a first 3D convolution, a first ReLU, a second 3D convolution, a second ReLU, a max pooling layer, and a reshape layer.

5. The method of claim 4 wherein each 3D convolution uses a kernel no larger than 3x3x3.

6. The method of claim 4 or 5 wherein the max pooling layer of the 3D convolution layer block reduces the size of the input spatial dimensions by a factor but maintains the

depth dimension, and the reshape layer increases the size of the depth dimension the factor.

7. The method of any one of claims 1 to 6 wherein the convolutional neural network is trained using back propagation using image segments of object and 3D joint positions.

8. The method of any one of claims 1 to 7 wherein the object is a human.

9. A system for determining the 3D position of joints on a skeleton from an 2D image comprising

- a. A 2D skeleton detector that identifies one or more 2D joint skeletons from the 2D image and an image segment for the 2D joint skeleton;
- b. A 2D joint heatmap generator that produces a 2D joint heatmap for each joint of the one or more 2D joint skeletons;
- c. A depth heatmap estimator that receives the 2D joint heatmaps and the image segment, comprising at least one 3D convolution layer block which is configured to estimate the depth of each joint;
- d. A 3D joint constructor that combines the depth of each joint with the 2D joint skeleton to construct a 3D joint skeleton.

10. The system of claim 9 wherein the 2D joint heatmap generator uses 2D joint heatmaps from the 2D skeleton detector.

11. The system of claim 9 or 10 wherein the depth heatmap estimator further comprises comprising at least one 2D convolutional neural network block prior to the at least one 3D convolution layer block.

12. The system of claim 9 or 11 wherein the depth heatmap estimator further comprises a max pooling layer that receives the depths from the at least one 3D convolutional layers and the max pooling layer reduces the 2D resolution to one and

leaves only an expanded depth dimension for each joint prior to combining the 2D location of each kind of joint with the estimated depth of the kind of joint.

13. The system of any one of claims 9 to 12 wherein each of the 3D convolutional layer blocks comprises a first 3D convolution, a first ReLU, a second 3D convolution, a second ReLU, a max pooling layer, and a reshape layer.

14. The system of claim 13 wherein each 3D convolution uses a kernel no larger than 3x3x3.

15. The system of claim 13 or 14 wherein the max pooling layer of the 3D convolution layer block reduces the size of the input spatial dimensions by a factor but maintains the depth dimension, and the reshape layer increases the size of the depth dimension the factor.

16. The system of any one of claims 9 to 15 wherein the convolutional neural network is trained using back propagation using image segments of object and 3D joint positions.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5





FIG. 6






FIG. 8

INTERNATIONAL SEARCH REPORT

International application No.

PCT/IB2020/052936

A. CLASSIFICATION OF SUBJECT MATTER IPC: <i>G06T 7/50</i> (2017.01), <i>G06N 3/04</i> (2006.01), <i>G06T 1/40</i> (2006.01), <i>G06T 7/70</i> (2017.01), <i>H04N 13/122</i> (2018.01), <i>H04N 13/388</i> (2018.01) According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols) G06T (Keywords*)		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic database(s) consulted during the international search (name of database(s) and, where practicable, search terms used) Canadian Patent Database, Esp@cenet, Questel Orbit: convolutional neural network, 3D, skeleton, joint, depth, heatmap		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 2017/133009 (CHEN et al.) 10-08-207 (10 August 2017) Whole document.	1-16
A	CN 108549876 (GAO et al.) 18-09-2018 (18 September 2018) Whole document	1-16
☐ Further documents are listed in the continuation of Box C. ☒ See patent family annex.		
* Special categories of cited documents: "A" document defining the general state of the art which is not considered to be of particular relevance "D" document cited by the applicant in the international application "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family	
Date of the actual completion of the international search 22 June 2020 (22-06-2020)		Date of mailing of the international search report 23 June 2020 (23-06-2020)
Name and mailing address of the ISA/CA Canadian Intellectual Property Office Place du Portage I, C114 - 1st Floor, Box PCT 50 Victoria Street Gatineau, Quebec K1A 0C9 Facsimile No.: 819-953-2476		Authorized officer Ann Wen (819) 635-5307

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.

PCT/IB2020/052936

Patent Document Cited in Search Report	Publication Date	Patent Family Member(s)	Publication Date
WO2017133009A1	10 August 2017 (10-08-2017)	CN105787439A CN105787439B	20 July 2016 (20-07-2016) 05 April 2019 (05-04-2019)
CN108549876A	18 September 2018 (18-09-2018)	None	