

Video-based estimation of pain indicators in dogs

Hongyi Zhu
University of Amsterdam
Amsterdam, the Netherlands
h.zhu@uva.nl

Yasemin Salgırlı, Pınar Can, Durmuş Atılğan
Ankara University
Ankara, Turkey

Albert Ali Salah
Utrecht University
Utrecht, the Netherlands
Boğaziçi University
Istanbul, Turkey

Abstract—Dog owners are typically capable of recognizing behavioral cues that reveal subjective states of their dogs, such as pain. But automatic recognition of the pain state is very challenging. This paper proposes a novel video-based, two-stream deep neural network approach for this problem. We extract and preprocess body keypoints, and compute features from both keypoints and the RGB representation over the video. We propose an approach to deal with self-occlusions and missing keypoints. We also present a unique video-based dog behavior dataset, collected by veterinary professionals, and annotated for presence of pain, and report good classification results with the proposed approach. This study is one of the first works on machine learning based estimation of dog pain state.¹

Index Terms—Pain estimation, dogs, pose estimation, LSTM

I. INTRODUCTION

Interpretation of the behavior of animals is critical to understand their well-being. In this paper, we propose an approach for automatic estimation of pain states of dogs, using video modality. This problem is important for several potential applications, including long-term automatic monitoring of dogs, assessment and diagnosis aids for veterinary clinicians, and early warning systems for nonexperts. It is challenging, because dogs come in different breeds that have very different appearances, and exhibit different behavioral coping strategies.

Animal pain estimation from video uses facial expressions of animals, as well as cues extracted from the body and the posture of the animal. While a comprehensive work for automatic dog pain estimation is missing from the literature, there are works on sheep, mice, pigs, and horses [1]. Each animal requires a separate data collection effort, a different analysis approach, and different pain annotation schemes.

In this paper, we make several contributions to the literature on affective computing for animal welfare:

- 1) We propose a pipeline for pain estimation in dogs, incorporating a number of off-the-shelf and in-house developed tools. Our approach is the first in the literature for automatic dog pain estimation from body pose.
- 2) We introduce a database of dog videos, annotated by veterinary experts, to evaluate our approach. Pain estimation in dogs is an under-researched problem, and there are limited resources; we hope that our approach will establish a good baseline, and there will be more research to follow.

¹Unedited author proof. Cite as: Zhu, H., Y. Salgırlı, P. Can, D. Atılğan, A.A. Salah, "Video-based estimation of pain indicators in dogs," Int. Conference on Affective Computing & Intelligent Interaction (ACII), 2023.

- 3) We make all our code and annotations available publicly².

II. RELATED WORK

Automated pain estimation for animals focuses on indicators from the body and the face of the animal. There is one recently introduced dog emotion recognition in the wild dataset [2], which contains images in uncontrolled settings, mostly suitable for end-to-end estimation without pose detection. For controlled video acquisition, pose estimation is an important step. We first discuss pose estimation in this section, and then move to pain estimation.

A. Pose estimation

There are four main approaches for representing the dog, as shown in Figure 1.

- Body key-joints representation: This representation allows easy modeling of pose semantics, and resembles the skeleton models popular in human behavior analysis. Examples of body joints are “shoulders”, “middle joint of knees”, “paw”, and “spine”.
- Outline boundary: In this representation, the boundary of the dog is marked. This method can represent the rigid shape of the animal.
- Instance segmentation: Processing the image pixel-wise and segmenting the dog out of the background is called instance segmentation. This representation is used in visual scene understanding applications, and subsequently, there are several image databases that have this kind of ground truth annotations for dog images.
- Bounding box: This representation is a cheaper-to-annotate version of instance segmentation that uses a rectangle to mark out the dog’s extent in the image. Some pose estimation and instance segmentation algorithms would require bounding box detection as a first step to reduce background interference.

Body key-joints is the main representation method for pose estimation approaches, as they depict the geometrical configuration of multiple body parts. For a pose estimation algorithm, some methods would first use the bounding box to identify the location of a single animal object [3]–[5] or apply instance segmentation to extract the shape of the animal body.

²The dataset can be accessed under special conditions permitted by the informed consent forms.




Fig. 1. Image based representations for the dog. (1) Body key-joints, (2) Outline boundary, (3) Instance segmentation, (4) Bounding box

Combining the advantage of the above four representation methods can help an algorithm keep track of the target animal and exclude noise from the background.

Recent approaches for animal pose estimation like DeepLabCut [6] and LEAP [7] use a convolutional neural network (CNN) to directly extract features from animal images and to convert it to a key-joint representation. To train the CNNs, a human-annotated training set is used. To deal with the lack of large amounts of human-annotated data, transfer learning is the preferred solution. In DeepLabcut, a ResNet model [8] pre-trained on an object recognition task is used to enable learning animal specific estimators from only hundreds of annotated training samples. Additionally, in [9], active learning is used to reduce the number of required annotations. In this approach, the system will suggest which images to annotate in each round, selecting these images actively. On the other hand, semi-supervised learning is used in [10], where a larger set of images without annotations is used jointly with a smaller set of annotated images to perform pose estimation. Apart from these three main approaches that can be combined (i.e. transfer learning, active learning, and semi-supervised learning), data augmentation is also used to increase the amount of available samples.

In this work, we use the Deep High-Resolution Network (HRNet), which is a state-of-the-art framework for 2D pose estimation that illustrated superior performance on a number of case studies [11]. This approach uses an innovative network architecture, described in Section III-A.

The pose estimation approaches we mentioned are generic and can be used with different animals. The DeepPoseKit method, for example, is applied on vinegar fly, locust, and zebra pose estimation problems [9]. However, popular human pose estimators use further information from the human skeleton structure [12]. In a similar way, animal-specific skeleton models can be used to improve the detection of the pose.

A number of approaches have been used or proposed for pose estimation in dogs and other canines. Modern ob-

ject detection algorithms, trained on large sets, such as MS COCO [13], typically contain a “dog” class. These can be leveraged to find a bounding box (or a more detailed segmentation) for detecting dogs in an image. The work by Tsai and Huang [14] introduced a multiple stages pose estimation algorithm, which uses a Mask R-CNN model [15] to generate the contour mask map of the dog. A faster R-CNN [16] performs posture recognition on the contour mask image and key part recognition on the original image. The key points of the skeleton would be obtained by analysing the candidate bounding box and animal pose jointly.

In the RGBD-Dog approach [17] an RGBD dog dataset with landmark ground truth is used, generated from a 3D motion capture system. They applied a stacked hourglass network [18] to predict a set of 2D heatmaps for a given depth image. This can determine the 3D coordinates of the body joints. A Hierarchical Gaussian Process Latent Variable Model (H-GPLVM) [19] is applied to refine the predicted 3D joint positions and to prevent pose ambiguities. However, this method requires depth data, which is not always available.

In this work, we use the HRNet-W32 model³, where W32 denotes the width of the high resolution network in the last three layers. The model is initially trained on the AnimalPose dataset [20], which contains pose annotation on five animal categories (i.e. dog, cat, cow, horse, sheep, respectively) with 6,000+ instances in 4,000+ images, and 20 body keypoints annotated. We will describe the experiments in Section V.

B. Pain estimation

Pain is defined as “an unpleasant sensory and emotional experience associated with, or resembling that associated with, actual or potential tissue damage,” [21]. It is not possible to measure the subjective experience of pain in animals from a simple image, or a video, but it is possible to automatically observe indications that there is pain, and quantify them to some extent to create useful applications, such as early warning or long-term monitoring systems. It is noted that behavioral cues are not very reliable for pain estimation, because the animal can choose not to exhibit behaviors under certain circumstances.

Automatic, computer vision based pain estimation has not been extensively researched for dogs until recently [22], but there is important work for other animals. [23] proposed a two-stream approach based on convolutional-LSTM (see Section III-D2) similar to our work, and shown its effectiveness for a video-based equine pain recognition task. Related pain estimation research has been performed for sheep [24], [25], mice [26], cats [27], rabbits [28], and horses [29]. These works mostly focus on the face of the animal, and use validated clinical scales that rely on a series of observations to score the presence of pain indicators. These observations can include the position of the ears, the visibility of sclera in the eyes, muscle tension in certain areas like the mouth, or presence of

³<https://github.com/open-mmlab/mmpose>

specific behaviors, such as baring the teeth in horses, which are exhibited when the animal is pain.

There is some related work on emotion estimation for dogs, which is relevant as a basis for pain estimation. In [30], three emotional states (i.e. growl, sleep, and smile, respectively) are automatically estimated from sequences of dog images, based on ears, eyes, mouth and head features. In [31], DeepLabCut is used to estimate anger, fear, happiness, and relaxation in dogs, based on a custom dataset of 400 images, evenly distributed over these four classes. Recently, [22] evaluated canine facial expressions for detecting positive anticipation and frustration.

The main challenges for computer vision based estimation of dog pain are the existence of many different dog breeds, which implies different sizes and colors that change the appearance, the presence of subtle cues, such as muscle contractions that are hidden beneath the coat, and the difficulties that plague ordinary face and body analysis, such as pose and illumination variations. Most importantly, there is a lack of data with pain annotations.

III. METHODOLOGY

Our proposed approach for detecting the visual indication of a dog’s pain state is made up of three parts. The off-the-shelf model for pose estimation is discussed in Section III-A. The feature extraction and pre-processing are described in Section III-B, followed by pose normalization. In Section III-D, we propose an architecture based on two-stream deep learning for pain detection. Figure 2 depicts our proposed system, which takes a short, 8-frame video clip (sampled over 4 seconds) as input, estimates the dog’s pose, and produces a pain score in the 0-1 range, which can be treated as a probability that the video contains a dog in pain.

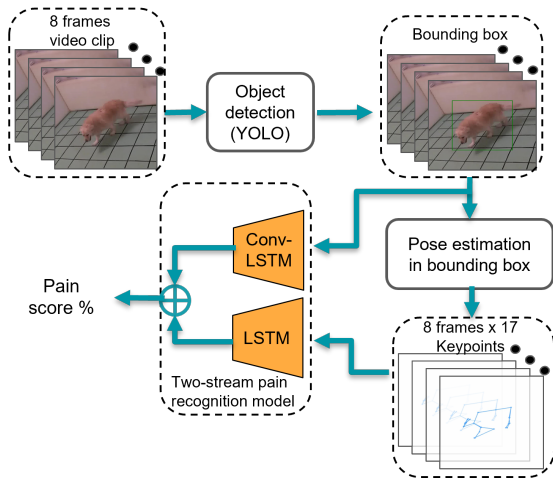


Fig. 2. Overall pipeline of the proposed method.

A. Dog pose estimation

For the initial detection of the dog in the images, we use a YOLOv5 deep neural network model [32], [33], pre-trained on the MSCOCO dataset [13], which includes a “dog” class. The RGB video frames are provided to the network

as input, and the network outputs predicted bounding boxes, a confidence score, and the likelihood of the dog class. We set the confidence threshold to 0.25, and we select only the bounding boxes with a confidence score for containing a dog greater than the threshold.

By combining the feature of YOLOv5 and HRNet, we can construct a data pre-processing pipeline to extract both RGB image and posture information as the model input. (See Figure 4 for the example image pre-processing representation)

As stated before, we use the HRNet [11] approach for pose estimation. The majority of existing pose estimation methods, such as SimpleBaseline [34] and DeepLabCut [35], adapt an encoder-decoder structure composed of a series of convolutional networks to downsample the input image from high to low-resolution feature maps and then apply transposed convolution layers to recover the original high-resolution feature map. HRNet, on the other hand, is composed of a central stem network that maintains high resolution throughout the process, as well as sample branch networks in each layer of the stem, and parallel connections between the multi-resolution branch networks. Additionally, HRNet implements multi-scale fusion between each parallel branch network to facilitate information exchange throughout the process. The stem network’s last layer outputs keypoint estimates (Fig. 3).




Fig. 3. The structure of HRNet. The horizontal and vertical axes correspond to the depth of the network and the scale of the feature maps, respectively (Adopted from [11]).

Since HRNet is a top-down pose estimation algorithm, it requires the approximate location of the target object in the image and preferably uses closely cropped images as inputs. The bounding box output of the YOLOv5 model is extended by 10% to ensure we include the entire dog, and provided to HRNet for further processing. We use the HRNet-W32 model obtained from [36], trained on the AnimalPose dataset [20], at this stage. Figure 4 illustrates the processing of a frame for bounding box localization and keypoints extraction.

B. Feature extraction and pre-processing

We extract 17 keypoints from various parts of the body of the dog, from the head, spine, and leg joints. The head includes two eyes, and the nose tip (3 points), spine includes withers and tail end (2 points) and legs include four elbows, four knees, and four paws (12 points). During training, since self-occlusions of joints is a common problem, we adopt a missing




Fig. 4. Example of bounding box localization and keypoints extraction. (a) We first apply the YOLOv5 algorithm to detect the bounding box of the dog. (b) We then crop the image with the center of the bounding box to exclude noisy background information. The bounding box is extended by 10% to ensure that the majority of the dog’s body is included. (c) We apply the HRNet pose estimation algorithm to detect the body keypoints within the bounding box, and generate the coordinates of the keypoints.

data treatment approach to augment some missing landmarks and to exclude some poorly estimated images. The poses are centered and scaled to be consistent with one another. This ensures that our approach is not affected by the differing size of the dog on the image (due to different distances to the camera).

Problems due to occlusions and self-occlusions can be partially solved by using body symmetry and by inference based on observed keypoints over time. The central concepts are as follows:

- 1) If some keypoints are missing from a frame, they may be discovered and interpolated by using the neighbouring frames in the sequence. Note that this makes would only work in an “offline” mode.
- 2) If -after the first step- there are more than nine keypoints missing, or the spine keypoints are missing, the frame is discarded from further processing.
- 3) Each leg has three keypoints (i.e. elbow, knee, and paw). If one keypoint of a leg is missing, it can be inferred according to the position of the other two keypoints and the spine points.
- 4) If more than one keypoint is missing in one leg, they can be guessed by exploiting the left and right side body symmetry, or front and back side symmetry.

This approach for resolving missing data is simple and effective, and necessary for subsequent processing of the body pose. In comparison to the head and spine keypoints, the leg keypoints convey more detailed information about the pose. Thus, the essence of this approach is to compensate for the missing keypoints induced by self-occlusion.

Our pose estimate step defines a mapping from the RGB image to body keypoints graph as:

$$s_i(t) \xrightarrow{\text{pose detector}} p_i(t) \quad \forall t = 1, \dots, T_i, \quad (1)$$

where $p_i(t)$ denotes the posture for sample i at time frame t and T denotes the total sequence length for sample i . The whole dataset contains N samples. In particular, $p_i(t)$ consists of a sequence of 2D coordinates as:

$$p_i(t) = \{(x_j(t), y_j(t))_i\}_{j \in J}, \quad (2)$$

where j is the index of keypoints, and J is the number of keypoints defined in the pose estimation model.

C. Pose normalization

The pose normalization step aims to transform the keypoint coordinates from an absolute to a local root point-centered reference system. The transformed pose coordinates are defined as:

$$(\bar{x}_j, \bar{y}_j)_i = (x_j, y_j)_i - (x_{root}, y_{root})_i \quad \forall j \in J \quad (3)$$

The dependence of t is omitted here for simplicity. In this paper, we use the middle point between the tail end and the neck as the root point. Thus, the root point centered coordinate is defined as:

$$\bar{p}_i = \{(\bar{x}_j, \bar{y}_j)_i\}_{j \in J} \quad (4)$$

Furthermore, in order to normalize the size of the dog within the image, we re-scale the keypoint coordinates \bar{p}_i :

$$\overline{\bar{p}}_i = \frac{\bar{p}_i}{\|v_{neck, tailend}\|}, \quad (5)$$

where $\bar{v}_{neck, tailend}$ is the vector between the neck and the tail end landmarks.

The transformation and normalization of landmarks has two purposes: The placement of the dog within the image has no effect on the processing, and the distance of the dog to the camera, as well as its size become less important. One may assume that in case much more training data become available in the future, some clustering in this space, followed by cluster-specific models, may improve the modeling.

D. Pain indicator estimation

Based on several other works [23], [37], [38], we propose to use a two-stream model structure to evaluate the temporal information in the detected pose representation. In [37], a two-stream model structure is used to predict human actions, one stream based on the long short term memory (LSTM) architecture [39] and the other with a 1-dimensional convolutional network. Our proposed method resembles this structure (see Figure 5), with some important differences. We use a basic LSTM branch to process the keypoint based pose representation, another, Convolutional LSTM based branch to process the RGB images, as in [23]. Both branches are complemented with an attention module before their outputs are combined with a fully convolutional layer, and the output is a binary classification of the pain state. We describe each component separately.

1) *Basic LSTM*: LSTM is a variant of recurrent neural networks, which use part of the internal information at time step t in the processing of time step $t + 1$. The basic LSTM network maintains a memory vector m to control the outputs at each time step. The LSTM network is able to learn temporal information, and combine stored information from previous time steps into its prediction. It uses three gating functions to control this: a forgot gate, an input gate, and an out gate, whereby it deletes old information or adds new information into the current state. However, this algorithm is not able to




Fig. 5. Our proposed two-stream model for dog pain estimation. The left branch is a four layer LSTM model, which is responsible for processing the keypoint information. The right branch is a four layer convolutional LSTM model, which is responsible for processing RGB images. The yellow and gray blocks are the max pooling and batch normalization layers. Black blocks represent the time attention layers. The green block is the fully connected layer. We fuse our two branches by concatenating the last layers’ feature vectors.

distinguish which time step contains more crucial information, and it would be a challenge for LSTM to learn long-term dependencies from longer sequences. In order to address this problem, an attention mechanism was proposed in [40] to learn long-term dependencies.

2) *Convolutional LSTM*: Traditional versions of LSTM would use the fully connected layers in every state transition and the gate output process. This kind of LSTM can also be called the *fully connected LSTM* (FC-LSTM). The FC-LSTM has the ability to extract temporal information from sequences, which makes it suitable for time series problems [39]–[42]. However, the major drawback of FC-LSTM is that its fully connected layers have too much redundancy for visual inputs. In order to address this problem, Shi et al. [43] introduced Convolutional LSTM (C-LSTM), which has convolutional structures in both the input-to-state and state-to-state transitions. C-LSTM provides the ability to process the spatio-temporal sequence more effectively [43].

3) *Attention LSTM*: In our model, we add attention operation on the top of the LSTM and C-LSTM outputs. The attention mechanism adds a context vector V onto the target sequence y . In [40], the elements of the context v_i depend on a sequence of annotations (b_1, \dots, b_T) with length of T . Each annotation would make a comparison across the whole input sequence and estimates which term of the input sequence should be paid more attention. The weighted sum of annotations b_i is used to compute the context vector as

follows:

$$v_i = \sum_{j=1}^{T_x} \alpha_{ij} b_j \quad (6)$$

The score α_{ij} of each annotation is inferred as:

$$\alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{k=1}^{T_x} \exp(e_{ik})} \quad (7)$$

The alignment model assigns a score α_{ij} to the pair of context vector elements at position j and the annotation at position i . e_{ij} is computed by the vector product between the LSTM hidden state v_{i-1} and the j^{th} annotation. In our model, both context vector and annotation are the hidden state sequences of the last LSTM layer, which forms a self-attention model. In [40], the alignment score α is parameterized by a feed-forward network with a single hidden layer, and this network is jointly trained with the other parts of the model.

4) *Fusion of branches*: The most direct way to combine the two feature vectors created by the two branches is concatenation fusion [44]. $\mathbf{y}^{\text{cat}} = f^{\text{cat}}(\mathbf{x}^a, \mathbf{x}^b)$ means we stack two feature vectors of the feature channels c :

$$y_{c+c'}^{\text{cat}} = x_c^a \quad y_{c'}^{\text{cat}} = x_{c'}^b, \quad (8)$$

where $\mathbf{y}^{\text{cat}} \in \mathbb{R}^{N \times (C+C')}$. This is denoted with FC in Figure 5.

IV. DATASETS

There are a few datasets with dog images available for training pose estimation models, but no datasets are available for assessing pain indicators. The dataset we introduce in this paper, is unique in that respect. We use two other datasets for assessing the pose estimation method that we use in this paper, namely, StanfordExtra [45] and TigDog [46].

StanfordExtra is an image-based dog dataset with annotated 2D keypoints, obtained by labeling part of the existing Stanford Dog Dataset [47], which consists of 20,580 dog images taken “in the wild” from 120 different dog breeds. The dataset contains various poses, with variation on environmental occlusions, interaction with humans or other animals, and partial views. For each image, there are 19 candidate body key points, but only the points that can be observed are labeled. We selected 6 dog breeds out of 120 randomly, and used 60 images for each breed for comparison experiments. The dog breeds are Dandie-Dinmont, Tibetan-terrier, Bluetick, Rhodesian-ridgeback, Brittany-spaniel, and Brabancon-griffon, respectively.



Fig. 6. Example images from the StanfordExtra (top) and TigDog (bottom) datasets.

TigDog is a video base dataset with no annotated keypoints, sourced from the Youtube-Objects dataset [48]. For most videos, the camera is not static, and follows the dog, which is typically centered in the image. Figure 6 shows some example images from these two datasets.

The Ankara University Dog Behavior Dataset (AUDBD) used in our experiment is collected by the Veterinary Medicine Department of Ankara University, in a group that performs research into dog behavior [49], [50]. The complete dataset includes 61 videos, each corresponding to an individual dog. The overall video duration is six hours and 45 minutes. There are 23 videos, totaling three hours and nine minutes, labeled as “pain condition”, and 38 videos with a total of three hours and 36 minutes labeled as “non-pain” condition.

Pain can be associated to wear and tear of tissues, injuries to the bones or a joint, or damage to muscles, ligaments, tendons or other soft tissues (musculoskeletal). These are marked as “orthopedic pain” in the dataset. Other types of pain can be related to irritated or damaged nerves, and these are marked as “neurological pain”. The labels of the dataset are provided by the veterinarian experts according to their medical records. Every sample in the pain set is a real dog patient with a physical problem that could potentially cause pain symptoms, however, not all images in the videos will exhibit all pain indicators.

We extract and process the frames from the videos at two frames per second, and in clips of four seconds of video each. Each such clip is thus composed of eight frames. In order to have enough data to train our models, we use a two frame (i.e. one second) overlap between subsequent clips. This way, the dataset contains a total of 12,435 clips (See Table I).

TABLE I

THE OVERALL STATISTICS OF THE AUDBD DATASET. FRAMES ARE SAMPLED AT TWO FPS, AND CLIPS CONSIST OF EIGHT FRAMES WITH TWO FRAMES OVERLAP. DURATION IS SHOWN IN HH:MM:SS.

	# videos	# clips	Duration
Orthopedic pain	14	2,957	01:19:06
Neurological pain	9	2,161	01:50:08
Total pain	23	5,118	03:09:14
Not pain	38	7,584	03:36:50

V. EXPERIMENTAL RESULTS

A. Pose estimation

We measure the performance of pre-trained HRNet on our dog pain dataset, as well as StanfordExtra and Tigdog. We use the Percentage of Correct Keypoints (PCK) measure, which is the percentage of keypoints detected falling within a normalized distance of the manually annotated ground truth. We use the square root of the bounding box area as the normalization index. The results are shown in Table II.

We can see that the pre-trained HRNet has a good performance in general, and has better performance on the StanfordExtra dataset compared to Tigdog and AUDBD, since StanfordExtra is visually more similar to the Animalpose dataset, on which the HRNet was trained.

TABLE II

10 PERCENT PCK (PERCENTAGE OF CORRECT KEYPOINTS) ON THREE DATASETS. HEAD INCLUDES TWO EYES, AND THE NOSE TIP. SPINE INCLUDES WITHERS AND TAIL END. LEGS INCLUDES FOUR ELBOWS, FOUR KNEES, AND FOUR PAWS.

Dataset	Head	Spine	Legs	Total
StanfordExtra	93.4	88.9	84.0	86.5
Tigdog	85.1	83.3	78.5	81.5
AUDBD	88.4	82.1	77.3	80.9

B. Pain estimation

Our classification models are evaluated using five fold cross-validation (i.e. 80% training, 20% validation), after separating a 13% portion of the entire dataset as an independent test set that is not seen at all during model selection. For each fold, we randomly sample the pain and non-pain classes proportionally in both training and validation folds, which is mainly used for stopping the training. There is no dog used in multiple folds. The test set is not used to train or validate the model.

We contrast several models under this setting. We report the performance of LSTM and C-LSTM branches separately, where the former uses the keypoints, and the latter uses the RGB images. LSTM models are reported under three different settings: LSTM32, LSTM64, and LSTM128, respectively, where the number denoted the number of hidden units. For one-stream networks, we also report baseline results with an I3D network [51] which uses RGB images, and with XGBoost [52] which uses keypoints. The hyperparameters for the baseline XGBoost are optimized via grid search. We use the I3D network that was pre-trained on the Kinetics-300 [51] dataset and fine-tuned the last classification head on our dataset. The I3D model achieves close performance to our C-LSTM model, but uses three times as many parameters as our approach. Finally, we also report results with a frequently used dual-stream approach, where one branch processes RGB image, and the second branch processes the optical flow image of the same video. Table III reports the average accuracy and F1 scores for these models obtained after five fold cross-validation, reported on the test set. The F1-score (i.e. the harmonic mean of precision and recall) is reported in addition to accuracy, since it is informative when the dataset is imbalanced. The model with the highest average F1 score and accuracy is C-LSTM+LSTM128, which has an average F1 score of $76.3\% \pm 4.4$ and an average accuracy of $77.0\% \pm 4.6$.

C. Qualitative Analysis

We use the Grad-CAM method [53] to compute activation (saliency) maps for our best model (i.e. ConvLSTM+LSTM) to visualize the model’s attention. Following [23], we apply Grad-CAM on the last convolutional layer of the RGB stream. In order to obtain the class discrimination saliency map, we first compute the gradient score of the class c , y^c , with respect to the feature map activations A^k (layer output) in the last convolutional layer, i.e. $\frac{\partial y^c}{\partial A^k}$. The gradient flows are then processed by the global average pooling in the width and

TABLE III

RESULTS (%F1-SCORE AND ACCURACY) FOR DIFFERENT MODELS. RGB: IMAGES, KP: KEYPOINTS, OF: OPTICAL FLOW. 32, 64, 128 DENOTE THE NUMBER OF HIDDEN UNITS IN THE LSTM MODELS, WHICH ARE USED TO PROCESS THE KEYPOINTS SEQUENCE.

One stream models	Data type	Avg. F1	Avg. Acc.	#Param.
C-LSTM	RGB	62.9±3.3	64.5±3.6	8,105,986
LSTM64	KP	58.8±5.5	59.1±6.8	137,922
I3D	RGB	74.1±2.5	74.6±1.1	24,327,632
XGBoost	KP	55.1±3.2	61.3±2.7	54,400
Two stream models				
C-LSTM+LSTM32	RGB+KP	70.3±3.3	71.2±3.1	8,143,204
C-LSTM+LSTM64	RGB+KP	72.9±3.5	72.1±4.1	8,243,908
C-LSTM+LSTM128	RGB+KP	76.3±4.4	77.0±4.6	8,635,780
Double C-LSTM	RGB+OF	67.7±5.1	69.2±3.9	16,210,403

height dimensions (denote as i and j respectively) to get the neuron importance weights α_k^c .

$$\alpha_k^c = \frac{\overbrace{\frac{1}{Z} \sum_i \sum_j}^{\text{global average pooling}} \underbrace{\frac{\partial y^c}{\partial A_{ij}^k}}_{\text{gradients via backprop}}}{Z} \quad (9)$$

where Z is the product of width and height. The importance of feature maps (layer output channels) k for a class c are then weighted according to their corresponding gradient magnitude. The feature maps are processed by a RELU function to focus on the features which have a positive impact on the class of interest.

$$L_{\text{Grad-CAM}}^c = \text{ReLU} \left(\underbrace{\sum_k \alpha_k^c A^k}_{\text{linear combination}} \right) \quad (10)$$

The region with the highest magnitude is considered to be the most important for the classification decision. The output feature map is up-sampled with an interpolation to the size of the original image, and superposed on it as a heatmap.

Figure 7 shows two dogs in orthopedic pain condition. The top sequence in Figure 7 is a case that is correctly classed as pain by our ConvLSTM+LSTM model, with a confidence score of 0.95. According to the heatmap, the lower body parts and the head parts are the regions that the model most focuses on. This is reasonable, since a dog with an orthopedic pain would be unstable in their walking motion, and this would show up in their leg movements and head position. We note that the algorithm seems not to be affected too much by the image background, and could follow the temporal patterns of the dog. In the bottom sequence of Figure 7, we see an example sequence where the dog is in pain, but the algorithm classifies it incorrectly as “no pain”, with a low confidence score. From the heatmap, we can see that algorithm is focusing on the back of the dog, and the legs are not clearly visible from this angle. Therefore, the algorithm cannot reliably detect the pain expression.




Fig. 7. Saliency maps for a true positive (top) and a false negative (bottom) case belonging to the pain class.

VI. CONCLUSIONS

Recognizing pain-related behaviors in dogs automatically is a challenging problem. Clinical pain assessment methods, such as the Glasgow composite measure pain scale, use multiple behavioral indicators such as “vocalization, attention to wound, mobility, response to touch, demeanor and posture/activity” for assessment of behaviours [54]. In this paper, we have barely scratched the surface of behavior analysis for this task, and integrated computer vision based tools into a solution that uses a small number of cues. Furthermore, we have worked only on samples with orthopedic and neuropathic pain; other categories, such as organic pain, were not included.

Our results show that both spatial-temporal features in images and keypoint based descriptions of the motion patterns are useful for the visual analysis. While data restrictions and the risk of overfitting prevented us from fine-tuning pose estimators, we show that generic tools have good performance for this task. Contrary to most work on animal pain estimation, we have not focused on the facial region. This remains a future work, but would require further clinical tools, such as validated grimace scales for labeling.

ETHICAL IMPACT STATEMENT

The proposed approach cannot be used as a clinical assessment tool for dog pain estimation without further tests and validation. No dogs were harmed in this research, there was no induced pain for the subjects, and ethics committee approval is obtained for the preparation of the dog video dataset.

REFERENCES

- [1] S. Broome, M. Feighelstein, A. Zamansky, G. Carreira Lencioni, P. Haubro Andersen, F. Pessanha, M. Mahmoud, H. Kjellström, and A. A. Salah, “Going deeper than tracking: A survey of computer-vision based recognition of animal pain and emotions,” *International Journal of Computer Vision*, vol. 131, no. 2, pp. 572–590, 2023.
- [2] F. Hernández-Luquin, B. Gutierrez-Serafin, H. J. Escalante, H. Pérez-Espinosa, L. Villaseñor-Pineda, and V. Reyes-Meza, “Dog emotion recognition from images in the wild: Debiw dataset and first results,” in *Proceedings of the Ninth International Conference on Animal-Computer Interaction*, 2022, pp. 1–13.
- [3] G. Papandreou, T. Zhu, N. Kanazawa, A. Toshev, J. Tompson, C. Bregler, and K. Murphy, “Towards accurate multi-person pose estimation in the wild,” in *Proc. CVPR*, 2017, pp. 4903–4911.
- [4] H.-S. Fang, S. Xie, Y.-W. Tai, and C. Lu, “RMPE: Regional multi-person pose estimation,” in *Proc. CVPR*, 2017, pp. 2334–2343.
- [5] Y. Chen, Z. Wang, Y. Peng, Z. Zhang, G. Yu, and J. Sun, “Cascaded pyramid network for multi-person pose estimation,” in *Proc. CVPR*, 2018, pp. 7103–7112.
- [6] A. Mathis, P. Mamidanna, K. M. Cury, T. Abe, V. N. Murthy, M. W. Mathis, and M. Bethge, “DeepLabcut: markerless pose estimation of user-defined body parts with deep learning,” *Nature neuroscience*, vol. 21, no. 9, pp. 1281–1289, 2018.

- [7] T. D. Pereira, D. E. Aldarondo, L. Willmore, M. Kislin, S. S.-H. Wang, M. Murthy, and J. W. Shaevitz, "Fast animal pose estimation using deep neural networks," *Nature methods*, vol. 16, no. 1, pp. 117–125, 2019.
- [8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, 2016, pp. 770–778.
- [9] J. M. Graving, D. Chae, H. Naik, L. Li, B. Koger, B. R. Costelloe, and I. D. Couzin, "DeepPoseKit, a software toolkit for fast and robust animal pose estimation using deep learning," *eLife*, vol. 8, p. e47994, 10 2019. [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/31570119>
- [10] A. Wu, E. K. Buchanan, M. Whiteway, M. Schartner, G. Meijer, J.-P. Noel, E. Rodriguez, C. Everett, A. Norovich, E. Schaffer *et al.*, "Deep graph pose: a semi-supervised deep graphical model for improved animal pose tracking," *bioRxiv*, 2020.
- [11] K. Sun, B. Xiao, D. Liu, and J. Wang, "Deep high-resolution representation learning for human pose estimation," in *Proc. CVPR*, 2019, pp. 5693–5703.
- [12] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, "Realtime multi-person 2d pose estimation using part affinity fields," in *Proc. CVPR*, 2017, pp. 7291–7299.
- [13] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *Proc. ECCV*. Springer, 2014, pp. 740–755.
- [14] M.-F. Tsai and J.-Y. Huang, "Predicting canine posture with smart camera networks powered by the artificial intelligence of things," *IEEE Access*, vol. 8, pp. 220 848–220 857, 2020.
- [15] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *Proc. ICCV*, 2017, pp. 2980–2988.
- [16] R. Girshick, "Fast R-CNN," in *Proc. CVPR*, 2015, pp. 1440–1448.
- [17] S. Kearney, W. Li, M. Parsons, K. I. Kim, and D. Cosker, "Rgb-dog: Predicting canine pose from rgb-d sensors," in *Proc. CVPR*, 6 2020.
- [18] A. Newell, K. Yang, and J. Deng, "Stacked hourglass networks for human pose estimation," in *Proc. ECCV*. Springer, 2016, pp. 483–499.
- [19] N. D. Lawrence, "Gaussian process latent variable models for visualisation of high dimensional data," in *Proc. NeurIPS*, vol. 2, 2003, p. 5.
- [20] J. Cao, H. Tang, H.-S. Fang, X. Shen, C. Lu, and Y.-W. Tai, "Cross-domain adaptation for animal pose estimation," in *Proc. ICCV*, 2019.
- [21] S. N. Raja, D. B. Carr, M. Cohen, N. B. Finnerup, H. Flor, S. Gibson, F. Keeffe, J. S. Mogil, M. Ringkamp, K. A. Sluka *et al.*, "The revised IASP definition of pain: Concepts, challenges, and compromises," *Pain*, vol. 161, no. 9, p. 1976, 2020.
- [22] T. Boneh-Shitrit, M. Feigelstein, A. Bremhorst, S. Amir, T. Distelfeld, Y. Dassa, S. Yaroshetsky, S. Riemer, I. Shimshoni, D. S. Mills *et al.*, "Explainable automated recognition of emotional states from canine facial expressions: the case of positive anticipation and frustration," *Scientific reports*, vol. 12, no. 1, p. 22611, 2022.
- [23] S. Broomé, K. B. Gleerup, P. H. Andersen, and H. Kjellstrom, "Dynamics are important for the recognition of equine pain in video," in *Proc. CVPR*, 2019, pp. 12 667–12 676.
- [24] Y. Lu, M. Mahmoud, and P. Robinson, "Estimating sheep pain level using facial action unit detection," in *2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)*. IEEE, 2017, pp. 394–399.
- [25] F. Pessanha, M. Mahmoud, and K. McLennan, "Towards automatic monitoring of disease progression in sheep: A hierarchical model for sheep facial expressions analysis from video," in *Proc. IEEE FG*, 2020.
- [26] N. Andresen, M. Wöllhaf, K. Hohlbaum, L. Lewejohann, O. Hellwich, C. Thöne-Reineke, and V. Belik, "Towards a fully automated surveillance of well-being status in laboratory mice using deep learning: Starting with facial expression analysis," *Plos one*, vol. 15, no. 4, p. e0228059, 2020.
- [27] M. Feigelstein, I. Shimshoni, L. R. Finka, S. P. Luna, D. S. Mills, and A. Zamansky, "Automated recognition of pain in cats," *Scientific Reports*, vol. 12, no. 1, p. 9575, 2022.
- [28] S. C. Keating, A. A. Thomas, P. A. Flecknell, and M. C. Leach, "Evaluation of emla cream for preventing pain during tattooing of rabbits: changes in physiological, behavioural and facial expression responses," *Plos one*, vol. 7, no. 9, p. e44437, 2012.
- [29] H. I. Hummel, F. Pessanha, A. A. Salah, T. J. van Loon, and R. C. Veltkamp, "Automatic pain detection on horse and donkey faces," in *Proc. IEEE FG*. IEEE, 2020, pp. 793–800.
- [30] V. Franzoni, A. Milani, G. Biondi, and F. Micheli, "A preliminary work on dog emotion recognition," in *IEEE/WIC/ACM International Conference on Web Intelligence-Companion Volume*, 2019, pp. 91–96.
- [31] K. Ferres, T. Schloesser, and P. A. Gloor, "Predicting dog emotions based on posture analysis using deeplabcut," *Future Internet*, vol. 14, no. 4, p. 97, 2022.
- [32] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.
- [33] —, "Yolo9000: better, faster, stronger," in *Proc. CVPR*, 2017, pp. 7263–7271.
- [34] B. Xiao, H. Wu, and Y. Wei, "Simple baselines for human pose estimation and tracking," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 466–481.
- [35] T. Nath, A. Mathis, A. C. Chen, A. Patel, M. Bethge, and M. W. Mathis, "Using deeplabcut for 3d markerless pose estimation across species and behaviors," *Nature Protocols*, vol. 14, no. 7, pp. 2152–2176, 7 2019. [Online]. Available: <https://doi.org/10.1038/s41596-019-0176-0>
- [36] M. Contributors, "Openmmlab pose estimation toolbox and benchmark," <https://github.com/open-mmlab/mmpose>, 2020.
- [37] F. Karim, S. Majumdar, H. Darabi, and S. Harford, "Multivariate LSTM-FCNs for time series classification," *Neural Networks*, vol. 116, pp. 237–245, 2019.
- [38] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *Proc. NeurIPS*, 2014, vol. 27.
- [39] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [40] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [41] A. Graves, "Generating sequences with recurrent neural networks," *arXiv preprint arXiv:1308.0850*, 2013.
- [42] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell, "Long-term recurrent convolutional networks for visual recognition and description," in *Proc. CVPR*, 2015, pp. 2625–2634.
- [43] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo, "Convolutional lstm network: A machine learning approach for precipitation nowcasting," in *Proc. NeurIPS*, 2015, pp. 802–810.
- [44] C. Feichtenhofer, A. Pinz, and A. Zisserman, "Convolutional two-stream network fusion for video action recognition," in *Proc. CVPR*, 2016, pp. 1933–1941.
- [45] B. Biggs, O. Boyne, J. Charles, A. Fitzgibbon, and R. Cipolla, "Who left the dogs out? 3d animal reconstruction with expectation maximization in the loop," in *Proc. ECCV*. Springer, 2020, pp. 195–211.
- [46] L. Del Pero, S. Ricco, R. Sukthankar, and V. Ferrari, "Articulated motion discovery using pairs of trajectories," in *Proc. CVPR*, 2015.
- [47] A. Khosla, N. Jayadevaprakash, B. Yao, and F.-F. Li, "Novel dataset for fine-grained image categorization: Stanford dogs," in *Proc. CVPRW*, 2011.
- [48] A. Prest, C. Leistner, J. Civera, C. Schmid, and V. Ferrari, "Learning object class detectors from weakly annotated video," in *Proc. CVPR*, 2012, pp. 3282–3289.
- [49] Y. D. Salgırlı, "Evaluation of body postures of belgian malinois dogs during a police dog training in germany," *Ankara Üniv Vet Fak Derg*, vol. 59, pp. 241–246, 2012.
- [50] Y. Salgırlı, "How to approach behavioural problems in dogs," *Journal of Istanbul Veterinary Sciences*, pp. 82 – 82, 2020.
- [51] J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," in *Proc. CVPR*, 2017, pp. 6299–6308.
- [52] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 2016, pp. 785–794.
- [53] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proc. CVPR*, 2017, pp. 618–626.
- [54] J. Reid, A. Nolan, J. Hughes, D. Lascelles, P. Pawson, E. Scott *et al.*, "Development of the short-form glasgow composite measure pain scale (cmps-sf) and derivation of an analgesic intervention score," *Animal Welfare*, vol. 16, pp. 97–104, 2007.