

Incremental Learning for Animal Pose Estimation using RBF k-DPP

Gaurav Kumar Nayak*¹

gauravnayak@iisc.ac.in

Het Shah*¹²

divhet163@gmail.com

Anirban Chakraborty¹

anirban@iisc.ac.in

¹ Indian Institute of Science

Bangalore, India

² BITS Pilani

Goa, India

Abstract

Pose estimation is the task of locating keypoints for an object of interest in an image. Animal Pose estimation is more challenging than estimating human pose due to high inter and intra class variability in animals. Existing works solve this problem for a fixed set of predefined animal categories. Models trained on such sets usually do not work well with new animal categories. Retraining the model on new categories makes the model overfit and leads to catastrophic forgetting. Thus, in this work, we propose a novel problem of “*Incremental Learning for Animal Pose Estimation*”. Our method uses an exemplar memory, sampled using Determinantal Point Processes (DPP) to continually adapt to new animal categories without forgetting the old ones. We further propose a new variant of k-DPP that uses RBF kernel (termed as “*RBF k-DPP*”) which gives more gain in performance over traditional k-DPP. Due to memory constraints, the limited number of exemplars along with new class data can lead to class imbalance. We mitigate it by performing image warping as an augmentation technique. This helps in crafting diverse poses, which reduces overfitting and yields further improvement in performance. The efficacy of our proposed approach is demonstrated via extensive experiments and ablations where we obtain significant improvements over state-of-the-art baseline methods.

1 Introduction

Pose estimation is a multi-regression problem that predicts multiple keypoint locations such as Eyes, Nose, Knees, Paws from a given input image (or video). Often the joints are strongly articulated and can be occluded. Thus, pose estimation is quite challenging as it requires searching in the large space of all possible articulated poses. It is an important problem in computer vision due to its utility in several applications. Pose estimation works have inspired human-aware robotics [51, 53], where the robot adapts to human behavior by predicting the human poses. Moreover, there are a lot of behavioral studies on insects and lab animals, such as *Drosophila* [20, 23, 46] and rats [42]. Pose estimation has also played a vital role in 3D reconstruction of objects from images/videos [6, 22, 59, 50] and action recognition [11, 40].

Human pose estimation, where poses are predicted only for humans, is a fairly popular problem. This is mainly due to large-scale labeled datasets that are publicly available [9, 58]. A lot of deep learning based pose estimation models have been proposed for this task including OpenPose [42], AlphaPose [46] and DeepPose [53]. However, Animal pose estimation [53, 43, 50], predicting poses for the animal classes, is comparatively a tougher problem

to solve. In general, animals show higher inter and intra class variability compared to humans. Due to this variability, training a pose estimation model requires a larger amount of labeled instances. Even labeling such instances is a labor-intensive task that requires basic knowledge about the skeletal composition of the various animal species. Often, due to privacy concerns sharing these datasets publicly may not be possible. Sanakoyeu *et al.* [60] publicly released only a subset of the original dataset. To learn animal pose estimation from a limited amount of samples is a challenging problem.

In a traditional supervised setting, the set of categories that the pose estimation models can learn to predict poses usually remains fixed after training. However, in a real-world scenario, new categories of animals can keep coming up. Hence, our pose estimation models should also learn to predict poses for these new categories. To achieve this, these traditional systems would require retraining on the newly enhanced training set consisting of new class data and old class data. Due to memory constraints, we can only store a subset of old class data. Preserving knowledge for the old classes using only a subset of the dataset is challenging. Hence, selecting samples efficiently in such cases becomes a vital problem.

Additionally, in the traditional supervised systems, one has to retrain the model from scratch for the new task, which is computationally quite expensive. One may think of using a naïve finetuning approach to learn the new categories continually. However, finetuning suffers from overfitting on the new categories causing the pose estimation model to forget the old categories catastrophically. In contrast, humans exhibit the remarkable ability to learn new novel categories and retain previously acquired knowledge. Mimicking this human behavior in artificially intelligent models becomes an interesting and challenging problem at the same time with many exciting applications in the real world.

Research in the domain of Continual Learning, often referred to as Incremental learning, aims to develop such artificially intelligent systems that can continuously learn new categories while preserving previously acquired knowledge for the old classes. Incremental Learning works are often focused on the problems of Image classification [6, 13, 36, 48, 59], Text classification [15, 49, 61] and Object detection [4, 52]. However, this problem is still unexplored for Animal Pose estimation. Additionally, the task of pose estimation involving multi-regression is inherently more challenging than classification or object detection.

In this work, we introduce a novel task of “*Incremental Learning for Animal Pose Estimation*”. In this task, the system is provided with new classes in multiple incremental steps. To mitigate the catastrophic forgetting for old classes, at each incremental step, the system has access to an exemplar memory of fixed size in addition to the new class data. The system has to learn to estimate poses for the new classes without forgetting the old classes. The performance of the incremental learning system should not be affected by the order in which the new classes are added. Thus, we also explore various such scenarios of incrementally learning new classes via several ablation experiments.

Prior incremental learning works [13, 36] often add additional heads to the model to learn the new classes. This is memory inefficient, as the number of the parameters grows with the addition of new classes. Our approach does not add any additional layers for incremental steps, keeping the number of parameters for the model constant with the incremental steps. We propose to use Determinantal Point Process (DPP) [28] as a sampling approach to provide a diverse exemplar memory. DPP has been theoretically proven to sample a diverse set of points. To sample a fixed set of k points, we specifically use k -DPP [29]. We observed that the rank of the kernel is a limitation to the number of points one can sample from k -DPP. The rank of the linear kernel, the kernel used in most of the DPP applications [27, 29], is not sufficient enough to select the required amount of samples. Thus, to overcome this,

we propose a new variant of k-DPP that uses RBF Kernel (named as “*RBF k-DPP*”). Our proposed k-DPP variant overcomes the limitation of the number of points that can be sampled and performs even better than the k-DPP with a linear kernel. Additionally, to augment the exemplars, we also use an image warping technique to generate varied poses for a given image. This helps to overcome the class-imbalance and reduces the overfitting on the new class data, leading to further improvement in performance.

We summarize our contribution as follows:

- To the best of our knowledge, we are the first to introduce and solve the novel problem of “*Incremental Learning for Animal Pose Estimation*”.
- Our proposed method leverages Determinantal Point Processes (DPPs) to select samples for the exemplar memory.
- We propose a new variant of k-DPP that uses RBF Kernel (named as “*RBF k-DPP*”) which overcomes the sampling limitation of k-DPP.
- We also propose to use an image warping-based augmentation technique to generate varied poses for the exemplar memory, which further improves the performance.
- The effectiveness of our proposed approach is shown through extensive experiments and improvements over state-of-the-art baselines.

2 Related Works

Our work is closely related to Pose estimation, Incremental Learning, DPP, and Image Warping. Thus, we provide a brief description of the relevant works in these domains.

Human Pose Estimation: Human pose estimation, both in 2D and 3D, has been extensively studied in the literature. Given that our work focuses mainly on 2D pose estimation, we discuss the works belonging to them. Toshev and Szegedy [55] proposed DeepPose, one of the first major works in human pose estimation using Deep Neural Networks. In recent times, a lot of deep learning based models are proposed such as Stacked Hourglass [45], OpenPose [12], MaskRCNN [21], PoseResnet [58] and AlphaPose [16]. In our current work, we demonstrate our proposed method on the AlphaPose and PoseResnet. However, our framework is independent of the choice of the pose estimation model.

Animal Pose Estimation: Few works exist in this domain, and it is relatively less studied than human pose estimation. Some of the recent works have focused on designing pose estimation systems and studying biological functions for specific animals such as Drosophila flies [20, 35, 46], rats [42] and in the wild animals [34, 44]. There are various toolkits proposed for annotating animal poses, such as DeepPose Kit [19], DeepLabcut [47] and AniPose [23]. However, these works are limited to specific animal classes.

Recently, Cao *et al.* [11] and Sanakoyeu *et al.* [50] showed an effective domain adaptation strategy to predict pose for animal classes using human priors. Li *et al.* [32] and Mu *et al.* [44] proposed unsupervised domain adaptation strategies to predict animal poses. However, our proposed approach does not need any such prior knowledge. A major limitation of their approach is that it cannot adapt to new animal classes incrementally without retraining the model. If exemplars are sampled incorrectly, then the model performance on old classes may deteriorate, eventually leading to low performance on the combined animal classes. Our proposed sampling strategy tries to retain the old classes’ acquired knowledge.

Incremental Learning: The works in this area are limited to Text Classification [15, 49, 51], Object Detection [9, 52], and Image Classification [5, 13, 36, 48, 59]. One way to mitigate catastrophic forgetting in incremental learning is weight regularization, where a regularization term is added to the loss function [2, 25]. Another way is to use exemplar memory built with samples from the old classes. Sampling a set of exemplars is a challenging task

because the small representative set should help the model preserve the knowledge of the old classes. A naïve approach is selecting samples randomly from the entire set. This minimizes forgetting to minor extent and is computationally inexpensive [24, 48]. GDumb [47] is another approach that greedily balances the exemplar memory using random sampling. To improve performance further, Herding [13, 48] is proposed where samples around the class means are chosen as representative samples. Dark Experience Replay (DER) [10] uses reservoir sampling to select the samples for the exemplar memory.

Many approaches that commonly use exemplar memory further use Knowledge Distillation to regularize the model. Rebuffi *et al.* [48] propose iCaRL that uses an exemplar memory, sampled using herding strategy, and knowledge distillation loss to mitigate catastrophic forgetting. iCaRL performs better than LwF [56], which does not have an exemplar memory and performs distillation on new class data. Castro *et al.* [13] improved their performance over iCaRL by adding a balanced finetuning stage to reduce the model’s bias towards the new classes. However, Belouadah and Popescu [6] show that even finetuning on exemplars performs better than distillation on them. Our results show a similar trend.

Image Warping: This technique is widely used in applications such as surveillance [57], Image enhancement [24]. Mesh warping using triangular meshes, like Delaunay Triangulation [50] is a popular approach to transform images. However, getting mesh labels is a challenging task. A simpler warping technique is Thin-Plate Splines proposed by Bookstein [7], which takes an image and control points for before and after transformation to transform the grid space. Unlike most works, we use Image warping, Thin-Plate Splines in particular, as an augmentation technique to generate images with diverse poses for the exemplar memory.

DPP: Selecting samples for exemplar memory is an essential task in incremental learning setup. Determinantal Point Processes (DPPs) [28] are elegant probabilistic models that are theoretically proven to sample a diverse set of points for various tasks, including summarization [18], text summarization, image search, and poses [27]. A conditional DPP, k-DPP [29] is often used in real-world applications as it can sample exactly k samples from the DPP.

3 Proposed Approach

Pose estimation aims to predict the spatial coordinates for a given input image. To train a model in a supervised setting, we have access to the training dataset $\mathcal{D} = (I, G)$, where I are the input images, and G are the ground-truth pose labels. For an i^{th} image, its ground-truth pose label, G^i , can be expressed as,

$$G^i = \{g_1^i, g_2^i, \dots, g_J^i\} \in \mathbb{R}^{J \times 3} \quad (1)$$

where J denotes the number of keypoints and g_j^i is the j^{th} keypoint label for the i^{th} image. Each keypoint label, g_j^i , consists of the spatial coordinates and a value for whether the keypoint is visible in the image or not. Generally, the keypoint labels for an i^{th} image are converted to heatmaps, h_j^i for each $j \in \{1, 2, \dots, J\}$ using the Gaussian Kernel [12, 16, 63]. Specifically, we center the Gaussian kernel on the spatial coordinates of the keypoint $z_j^i = (x_j^i, y_j^i)$. Mathematically, h_j^i can be formulated as follows,

$$h_j^i(x_j, y_j) = \frac{1}{\sqrt{2\pi}\sigma^2} \exp\left(-\frac{[(x-x_j)^2 + (y-y_j)^2]}{2\sigma^2}\right) \quad (2)$$

where (x, y) denote the spatial location in the input image and σ is the standard deviation. The MSE loss between h_j and \hat{h}_j is used for training a supervised pose estimation for all keypoints $j \in \{1, 2, \dots, J\}$, where \hat{h}_j are the predicted heatmaps.

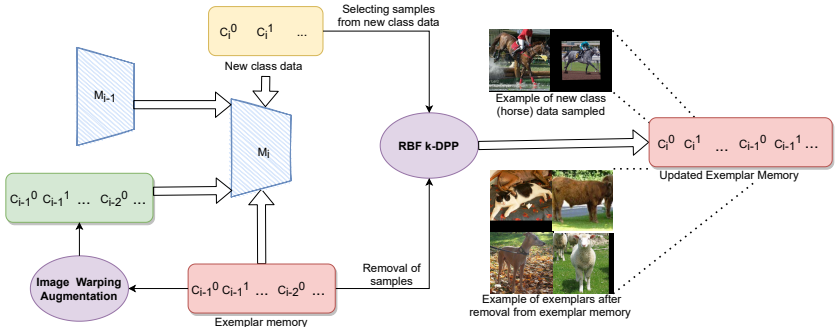


Figure 1: Our proposed strategy for an i^{th} incremental step is shown. After training the model M_i we update the exemplar memory by selecting new class samples and removing old class samples using RBF k-DPP. Image warping augmentation is further used to overcome class-imbalance.

Initially, the pose estimation model, $M_0(\cdot)$, is trained on the dataset \mathcal{D}_0 containing “ a ” base classes denoted by $C_0 = \{C_0^1, C_0^2, \dots, C_0^a\}$ base classes. At each i^{th} incremental step, we add a dataset \mathcal{D}_i containing “ b ” new classes, namely C_i , where $i \in \{1, 2, \dots, k\}$ denotes the incremental step. The goal of the pose estimation model, $M_i(\cdot)$, is to learn to predict keypoints on the new classes, C_i , added at the incremental step i , without forgetting the old classes, i.e. $C_0 \cup C_1 \cup C_2 \dots \cup C_{i-1}$.

To mitigate catastrophic forgetting, we use an exemplar memory [5, 13, 48] consisting of selected samples from old classes data. At each incremental step, the model, $M_i(\cdot)$, has access to the new class data, \mathcal{D}_{new} belonging to the set of classes, C_i , and an exemplar memory, E , containing data, $\mathcal{D}_{exemplar}$, which consists of the classes that the model has already seen, namely $C_0 \cup C_1 \cup C_2 \dots \cup C_{i-1}$. An overview of our proposed approach for Incremental Learning on Animal Pose Estimation is shown in Figure 1. In the next section, we discuss our method of the selection of samples for the exemplar memory in detail.

3.1 Exemplar memory

The exemplar memory, E , plays an important role in incremental learning to mitigate forgetting. This memory helps the model to preserve the knowledge for the old classes. Thus, we focus on developing the best strategy to select exemplars for our pose estimation task that can best represent the classes on which the model is trained. The model uses E as a means to revise the old classes. We primarily investigate the fixed memory setup, where the memory size is fixed to Γ . The size of the memory is independent of the number of classes. The number of samples per class, n stored in the exemplar memory, is $n = \lfloor \frac{\Gamma}{c} \rfloor$, where c is the number of classes that the model M has previously seen. After training the model at an incremental step, the exemplar memory is updated by removing samples from the current memory (Sec. 3.1.1) and adding selected samples from the new class data (Sec. 3.1.2).

3.1.1 Selection of new samples for the exemplar memory using DPP

The exemplar memory plays a pivotal role in the continual learning system, as the model uses only the exemplar memory to retain the knowledge for old classes. A diverse set of exemplars can help the model preserve the acquired knowledge. Selecting samples for the exemplar memory from the new class data is, therefore, a vital task for incremental learning.

Determinantal point processes (DPP) [18, 27, 28, 29] are elegant probabilistic models that can capture negative correlation and are theoretically proven to sample a diverse set of

samples. Let the set of pose labels for N samples be denoted as $\mathcal{V} = \{G^0, G^1, \dots, G^N\}$. The probability of selecting a random subset, $S \subseteq \mathcal{V}$ from the power set, $2^{\mathcal{V}}$, is given as:

$$P_L(S; L) = \frac{\det(L_S)}{\det(I + L)} \quad (3)$$

where L is real positive semi-definite kernel matrix, referred as L-ensemble [8], $L_S = [L_{ij}]_{y_i, y_j \in S}$ is the restriction of K to the rows and columns indexed by S , and I is the Identity matrix.

However, DPPs have a limitation, they are not constrained on the size of the optimal subset selection. This is not desirable for building exemplar memory in our case, as we need a fixed set of samples from the DPP. Thus, a conditional DPP has to be used. A variation of DPP, k-DPP [29], models only sets with cardinality k . The k-DPP is formulated as,

$$P_L^k(S; L) = \frac{\det(L_S)}{\sum_{|S'|=k} \det(L_{S'})} \quad (4)$$

where the cardinality of the set S is k , i.e. $|S| = k$, and L is a positive semi-definite kernel. It is important to note that $k \leq \text{rank}(L)$. To find the most likely subset from a k-DPP, we can simply use a MAP estimate,

$$S_{k\text{-DPP}}^{\text{MAP}} = \text{argmax}_S P_L^k(S; L) \quad (5)$$

Finding the MAP estimate is a NP-Hard problem [26]. Thus, it's greedy approximation is used [27, 28]. We discuss two different kernels that are used to construct the k-DPPs.

Linear kernel: The linear kernel [27, 28] function, F_L , is used to calculate the entries, x_{ij} (i^{th} row and j^{th} column), in the L-ensemble kernel matrix. The value of x_{ij} is denoted by $F_L(\mathbf{x}_i, \mathbf{x}_j)$ which is calculated as,

$$F_L(\mathbf{x}_i, \mathbf{x}_j) = (\mathbf{x}_i \cdot \mathbf{x}_j) \quad \forall \mathbf{x}_i, \mathbf{x}_j \in \mathcal{V} \quad (6)$$

To compute the L-Ensemble matrix (L) using F_L , we can simply do, $L = AA^T$, where A is the matrix of the pose labels for a particular class. The shape of this A matrix is $(N, J, 3)$; where N is the number of data points for a particular class and J are the number of keypoints, each keypoint consists of (x, y) co-ordinate location of the keypoint and whether the keypoint is visible or not. We reshape the A matrix to $(N, J \times 3)$ for simplicity. Thus, the rank of L is given by

$$\text{rank}(L) = \min(\text{rank}(A^T), \text{rank}(A)) \quad (7)$$

The maximum rank of L can thus be $J \times 3$. Thus, the maximum value of k , the number of points that can be sampled in k-DPP is $J \times 3$, which can be very less compared to the required amount of samples.

To overcome the above limitation, we cluster the entire data into h clusters using K-Means [9, 17, 41] based on the pose labels (named as '*k-DPP w/ clustering*'), where $h = \lfloor N / (J \times 3) \rfloor$. From each cluster, we sample $J \times 3$ exemplars giving us the desired number of points to be sampled for each class. This is computationally expensive because of the computational costs of K-Means and k-DPP. Due to varying cluster sizes and densities, K-Means can often give sub-optimal clustering of the data. Also, the initialization of cluster centroids can further heavily influence the performance in selecting a diverse set of points.

Proposed k-DPP variant (Radial Basis Function (RBF) Kernel based k-DPP): To overcome the bottleneck of K-Means clustering for the linear kernel and the computational costs,

we use the Radial Basis Function (RBF) to generate the L-ensemble for k-DPP (named as ‘RBF k-DPP’). In order to use RBF for constructing an L-Ensemble, we need to show that the RBF is positive semi-definite or a kernel. The proof for this can be referred from Supplementary. The RBF kernel function used to calculate entries, x_{ij} (i^{th} row and j^{th} column), in the L-ensemble kernel matrix is formulated as,

$$F_{RBF}(\mathbf{x}_1, \mathbf{x}_2) = \exp(-\gamma \|\mathbf{x}_1 - \mathbf{x}_2\|^2) \quad (8)$$

where $\gamma \geq 0$ is an hyper-parameter.

The rank of the RBF Kernel matrix, in the worst case, grows polynomially with the data dimension [54]. For our case, the rank of the RBF Kernel is sufficiently large enough to sample any amount of samples. This removes the dependency on the use of K-Means clustering and thus is only required to be applied once per class, thereby reducing the computational costs significantly. Moreover, we observed RBF k-DPP also performs better than the k-DPP w/ clustering.

3.1.2 Removal of samples from the Exemplar Memory

After each Incremental step, the Exemplar Memory, E , is updated to accommodate new classes in the memory. To achieve this, we first select a subset of samples for the classes already present in Exemplar memory to make room for new class samples. Removal of samples for the old classes is done by applying the k-DPP again on exemplary memory. After this operation, the removed samples are never used again.

3.2 Image Warping Augmentation

To further augment the exemplar memory, we use Thin Plate Splines Image warping (TPS) [7] to generate samples with different poses. TPS achieves image warping by taking two sets of control points, the input, and the output control points. For our case, we take the limbs visible in the image for an animal, specifically the knee and the paw keypoints. We rotate these keypoints by a small random angle about the elbow keypoint for the respective limb to get the output control points. The rotation of these keypoints is done under skeletal constraints for each animal. TPS then aligns the input and the output control point sets by warping the grid. We further use the image In-painting technique based on the work by Telea [54] to fill in the empty portions of the image. This in-painting starts from the region’s boundary, then gradually fills points towards the center of the region. Each pixel is assigned a normalized value from a small neighborhood of pixels around it. Refer supplementary for visualization of samples, before and after image warping and in-painting operations.

4 Experiments

We use Finetuning as our first baseline, where the model is finetuned on the new class data and has no access to the exemplar memory. Further, we compare the performance of our k-DPP constructed using RBF and Linear Kernel against Herding strategy and Random Sampling strategy [13, 48]. Additionally, we compare our proposed method against some of the heuristic baselines. We modify popular class-incremental learning works on image classification and adapt them to our pose estimation task, which serve as strong baselines for comparison. Oracle denotes the performance of the model when it has access to the entire old class samples, an upper bound for all the approaches.

Random Sampling Strategy: This is one of the easiest ways where n samples are randomly selected from the entire set of samples to build the memory.

Herding Strategy: This has been widely used as a sampling technique for exemplar memory [13, 48, 57]. We sort the samples based on the distance from the mean pose for each class. We then select the top- n samples for each class from these sorted lists of samples.

Adapted-iCaRL: We modify iCaRL [48], originally proposed for classification, to make it suitable for pose estimation task. We replace the cross-entropy loss term with MSE loss for the new class data. We perform distillation on the exemplar memory, built using herding strategy, with MSE loss between the model from the previous incremental step and model at the current step. We give equal weightage to both the MSE loss terms.

Adapted-EEIL: Following the work Castro *et al.* [13], we add a balanced finetuning step, where equal number of samples for all the classes are considered and distillation loss on the new class data is also applied.

Adapted-DER++: We use the Dark Experience Replay (DER) sampling strategy proposed by Buzzega *et al.* [10] to select the samples for the exemplar memory at each incremental step. We modify the cross-entropy loss with the MSE loss and use the regularization losses proposed in their work.

Adapted-GDumb: We use the Greedy sampling strategy proposed by Bin *et al.* [58] to select the samples for the exemplar memory. We use the MSE loss for training the model at each incremental step.

Refer supplementary for more details on adapted baselines.

4.1 Implementation and Setting

We use AlphaPose [16] and PoseResnet [58] as the pose estimation models in our experiments. We select 17 keypoints similar to Cao *et al.* [10]. We set the value of $\sigma = 2$ for generating Gaussian heatmaps. We perform experiments on the Animal-Pose Dataset [10] which contains a total of 5117 instances from five animal classes, namely {Cat, Cow, Dog, Horse, Sheep}. We consider two different sizes for the fixed exemplar memory, $\Gamma = 500$ and $\Gamma = 1000$ samples which are only 9.7% and 19.5% of the total dataset size. We repeat the experiments three times and report the mean and standard deviation of the PCK@0.05 values for the predictions made by the pose estimation model on all the classes.

4.2 Results and Discussion

We first begin with one base class, ‘cat’, and then incrementally add one animal class at each step in the order of ‘horse’, ‘cow’, ‘dog’, ‘sheep’.

Approach	Memory Sizes (Γ)							
	1000				500			
	Incremental steps				Incremental steps			
	1	2	3	4	1	2	3	4
Oracle	87.49 ± 0.20	87.51 ± 0.16	84.01 ± 0.22	83.47 ± 0.23	87.49 ± 0.20	87.51 ± 0.16	84.01 ± 0.22	83.47 ± 0.23
Finetuning	64.51 ± 0.50	56.76 ± 0.16	56.56 ± 0.37	52.47 ± 0.47	64.50 ± 0.49	56.76 ± 0.16	56.56 ± 0.37	52.17 ± 0.47
Random	83.86 ± 0.16	78.14 ± 0.25	66.02 ± 0.61	64.22 ± 0.36	80.64 ± 0.17	74.12 ± 0.28	64.00 ± 0.22	60.88 ± 0.26
Herding	82.93 ± 0.29	75.90 ± 0.34	65.96 ± 0.34	61.69 ± 0.18	80.56 ± 0.14	71.82 ± 0.31	62.64 ± 0.18	59.16 ± 0.27
k-DPP w/ clustering (Ours)	84.14 ± 0.12	79.07 ± 0.29	67.75 ± 0.31	65.48 ± 0.43	81.63 ± 0.65	74.32 ± 0.43	64.65 ± 0.14	61.71 ± 0.10
RBF k-DPP ($\gamma = 0.5$) (Ours)	84.20 ± 0.48	78.99 ± 0.20	71.00 ± 0.28	68.77 ± 0.31	81.56 ± 0.06	74.71 ± 0.09	64.73 ± 0.12	61.88 ± 0.34
RBF k-DPP ($\gamma = 50$) (Ours)	83.80 ± 0.25	79.21 ± 0.35	71.38 ± 0.28	68.72 ± 0.08	81.56 ± 0.30	74.59 ± 0.25	65.26 ± 0.44	62.42 ± 0.27
RBF k-DPP ($\gamma = 100$) (Ours)	83.55 ± 0.11	79.21 ± 0.33	71.07 ± 0.42	68.86 ± 0.52	81.54 ± 0.39	75.01 ± 0.13	65.04 ± 0.18	62.46 ± 0.14

Table 1: Mean and Standard deviation of the overall PCK@0.05 of the model across classes are reported in the table. We show results for two fixed memory sizes, $\Gamma = 500$ and $\Gamma = 1000$ samples. We perform three ablations on γ for RBF kernel, with the values 0.5, 50 and 100.

Efficacy of k-DPP: From Table 1, we observe that our proposed approach of k-DPP w/ clustering performs better than the popularly used Herding strategy by 2.5% for the memory $\Gamma = 500$ samples, and 3.8% for $\Gamma = 1000$ samples, at the final incremental step. For the case

of $\Gamma = 500$ samples, we see a small improvement in performance for RBF k-DPP. Note that k-DPP with K-Means has its own shortcomings. It is computationally expensive since we have to perform the sampling using k-DPP for each cluster, whereas in RBF k-DPP, sampling is performed only once. For $\Gamma = 1000$ samples, we get an improvement of 3.2% in the performance for RBF k-DPP ($\gamma = 50$) against k-DPP w/ clustering. The effectiveness of RBF k-DPP can be seen with the increase in the number of incremental steps. Selection of a set of diverse exemplars becomes crucial with the increase in incremental steps as the number of samples per class decreases. RBF k-DPP is more effective in handling this compared to traditional k-DPP with K-Means.

Approach	Memory sizes (Γ)							
	1000				500			
	Incremental steps				Incremental Steps			
	1	2	3	4	1	2	3	4
k-DPP w/ clustering (Ours)	84.14 \pm 0.12	79.07 \pm 0.29	67.75 \pm 0.31	65.48 \pm 0.43	81.63 \pm 0.65	74.32 \pm 0.43	64.65 \pm 0.14	61.71 \pm 0.10
RBF k-DPP ($\gamma = 50$) (Ours)	83.80 \pm 0.25	79.21 \pm 0.35	71.38 \pm 0.28	68.72 \pm 0.08	81.56 \pm 0.30	74.59 \pm 0.25	65.26 \pm 0.44	62.42 \pm 0.27
RBF k-DPP ($\gamma = 50$) + Augmentation (Ours)	84.57 \pm 0.22	81.88 \pm 0.35	74.52 \pm 0.44	71.53 \pm 0.34	81.79 \pm 0.17	74.97 \pm 0.06	66.49 \pm 0.23	63.01 \pm 0.11

Table 2: Effect of Augmentation; Performance of our proposed RBF k-DPP ($\gamma=50$) with augmentation against RBF k-DPP($\gamma=50$) and k-DPP w/ clustering.

Effectiveness of Proposed Augmentation: Table 2 shows the effectiveness of our proposed image warping augmentation strategy on top of the proposed DPP variant (RBF k-DPP). We obtain further improvement of 0.6% and 2.8% at the final incremental step for $\Gamma = 500$ and $\Gamma = 1000$ samples, respectively. This improvement can be attributed mainly to the enhanced intra-class variance and diversity in poses after augmentation. It thus helps the model to overfit less on the new classes by overcoming class imbalance.

Approach	Memory sizes (Γ)							
	1000				500			
	Incremental steps				Incremental Steps			
	1	2	3	4	1	2	3	4
Adapted - iCaRL	61.90 \pm 0.09	48.49 \pm 0.43	38.51 \pm 0.39	34.88 \pm 0.47	61.64 \pm 0.49	47.15 \pm 0.50	39.24 \pm 2.62	37.53 \pm 1.16
Adapted - EEIL	72.17 \pm 0.63	64.86 \pm 0.66	60.43 \pm 0.69	57.41 \pm 0.01	69.15 \pm 0.15	59.99 \pm 0.32	58.64 \pm 0.31	56.05 \pm 0.04
Adapted - DER++	76.29 \pm 0.32	68.93 \pm 0.71	59.60 \pm 0.12	56.28 \pm 0.25	72.56 \pm 0.33	64.80 \pm 0.43	55.57 \pm 0.13	51.67 \pm 0.52
Adapted - GDumb	83.11 \pm 0.16	77.48 \pm 0.35	66.89 \pm 0.23	63.22 \pm 0.42	81.12 \pm 0.28	73.20 \pm 0.17	63.93 \pm 0.22	60.40 \pm 0.46
RBF k-DPP ($\gamma = 50$) + Augmentation (Ours)	84.57 \pm 0.22	81.88 \pm 0.35	74.52 \pm 0.44	71.53 \pm 0.34	81.79 \pm 0.17	74.97 \pm 0.06	66.49 \pm 0.23	63.01 \pm 0.11

Table 3: Performance of our proposed RBF k-DPP ($\gamma=50$) with augmentation against adapted incremental learning baselines using AlphaPose model [14]

Comparison against Adapted Incremental Learning Baselines: We compare our pose incremental learning model’s performance against the state-of-the-art iCaRL, EEIL, DER++ and GDumb adapted for the pose estimation task in Table 3. We observe a significant improvement of 2.6%, 11.3%, 7% and 25.5% in performance against GDumb, DER++, EEIL and iCaRL, respectively at $\Gamma = 500$ samples. Further, at $\Gamma = 1000$ samples, our proposed method again obtains significant improvement of 8.31%, 15.25%, 14.1% and 36.6% over GDumb, DER++, EEIL and iCaRL, respectively.

Other Incremental Learning Setups: We further perform experiments with different setups for incremental steps as shown in Figure 2. Our proposed approach consistently performs better than the adapted versions of the state-of-the-art works.

Comparison with other Augmentation Strategy: In Table 4, we compare the performance of our proposed Image Warping augmentation method with other augmentation strategies, namely Rotation, Flipping and Gaussian Noise. We obtain 1.5%, 0.3% and 0.1% improvement in performance over Noise, Rotation and Flipping respectively at $\Gamma = 1000$. Also, we

further observe an improvement of 1.7%, 0.9% and 0.8% compared to Noise, Rotation and Flipping respectively for $\Gamma = 500$.




Figure 2: Experimental Results for different setups of incremental steps. Base classes for (a) and (b) are {‘cat’, ‘cow’, ‘dog’} while {‘horse’} and {‘sheep’} are added at incremental steps. Similarly, the base class for (c) and (d) is ‘cat’, where {‘horse’, ‘cow’} and {‘dog’, ‘sheep’} are added at incremental steps. We set $\Gamma = 1000$ samples for (a) and (c) and $\Gamma = 500$ samples for (b) and (d).

Approach	Memory sizes (Γ)			
	1000		500	
	Incremental steps		Incremental steps	
	1	2	1	2
RBF k-DPP ($\gamma = 50$) + Flipping	81.53 \pm 0.23	75.46 \pm 0.35	81.53 \pm 0.28	73.41 \pm 0.36
RBF k-DPP ($\gamma = 50$) + Rotation	81.84 \pm 0.42	75.31 \pm 0.24	81.13 \pm 0.24	73.29 \pm 0.51
RBF k-DPP ($\gamma = 50$) + Noise	81.81 \pm 0.18	74.09 \pm 0.19	81.39 \pm 0.42	72.53 \pm 0.25
RBF k-DPP ($\gamma = 50$) + Augmentation (Ours)	82.05 \pm 0.15	75.59 \pm 0.13	82.03 \pm 0.34	74.19 \pm 0.33

Table 4: Experimental results of our proposed augmentation strategy against other augmentation strategies, Flipping, Rotation and Gaussian Noise. We consider {‘cat’, ‘cow’, ‘dog’} as base classes while {‘horse’} and {‘sheep’} are added at incremental steps.

Experiments on PoseResnet model [58]: We compare the performance of our proposed approach against the adapted state-of-the-art Incremental Learning baselines iCaRL, EEIL, DER++ and GDumb on PoseResnet [58] architecture. As observed in Table 5, our proposed approach achieves better performance than the adapted baselines for both the exemplar memory settings ($\approx 4.2\%$ for $\Gamma = 1000$ and $\approx 1.5\%$ for $\Gamma = 500$).

Approach	Memory sizes (Γ)			
	1000		500	
	Incremental steps		Incremental Steps	
	1	2	1	2
Adapted - iCaRL	50.30 \pm 0.31	37.80 \pm 0.40	46.60 \pm 0.26	35.93 \pm 0.26
Adapted - EEIL	45.30 \pm 0.19	39.97 \pm 0.46	41.33 \pm 0.36	38.72 \pm 0.55
Adapted - DER++	54.50 \pm 0.24	47.92 \pm 0.31	52.00 \pm 0.30	47.27 \pm 0.35
Adapted - GDumb	55.35 \pm 0.28	51.45 \pm 0.49	51.21 \pm 0.43	47.64 \pm 0.39
RBF-DPP ($\gamma = 50$) + Augmentation (Ours)	57.93 \pm 0.32	55.72 \pm 0.22	52.45 \pm 0.31	49.1 \pm 0.36

Table 5: Performance of our proposed RBF k-DPP ($\gamma = 50$) against adapted incremental learning baselines using PoseResnet model [58]. We consider {‘cat’, ‘cow’, ‘dog’} as base classes while {‘horse’} and {‘sheep’} are added at incremental steps.

5 Conclusion

This paper proposes a novel task of incrementally learning to predict pose for animals. Our experimental results show: i) selection of samples for the exemplar memory plays a pivotal role in incremental learning, ii) the effectiveness of proposed k-DPP over the popular heuristic baselines, iii) our proposed RBF k-DPP performs significantly better than state-of-the-art baselines and yields more gain over k-DPP besides being less computationally expensive, iv) augmentation using image warping handles class imbalance and improves performance.

6 Acknowledgements

This work is supported by a project grant from Wipro IISc Research and Innovation Network (WIRIN). We would like to extend our gratitude to all the reviewers for their valuable suggestions. We would also like to thank Siddharth Seth for helpful discussions on this project.

References

- [1] Saeid Agahian, Farhood Negin, and Cemal Köse. An efficient human action recognition framework with pose-based spatiotemporal features. *Engineering Science and Technology, an International Journal*, 23(1):196–203, 2020. ISSN 2215-0986. doi: <https://doi.org/10.1016/j.jestch.2019.04.014>. URL <https://www.sciencedirect.com/science/article/pii/S2215098618312345>.
- [2] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. 11 2017.
- [3] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [4] David Arthur and Sergei Vassilvitskii. K-means++: The advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '07, page 1027–1035, USA, 2007. Society for Industrial and Applied Mathematics. ISBN 9780898716245.
- [5] Eden Belouadah and Adrian Popescu. Il2m: Class incremental learning with dual memory. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 583–592, 2019. doi: 10.1109/ICCV.2019.00067.
- [6] Benjamin Biggs, Thomas Roddick, Andrew Fitzgibbon, and Roberto Cipolla. Creatures great and smal: Recovering the shape and motion of animals from video. In C.V. Jawahar, Hongdong Li, Greg Mori, and Konrad Schindler, editors, *Computer Vision – ACCV 2018*, pages 3–19, Cham, 2019. Springer International Publishing. ISBN 978-3-030-20873-8.
- [7] F.L. Bookstein. Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, 1989. doi: 10.1109/34.24792.
- [8] A. Borodin and A. Soshnikov. Janossy densities. i. determinantal ensembles. *Journal of Statistical Physics*, 113:595–610, 2002.
- [9] C. Brust, Christoph Käding, and J. Denzler. Active and incremental learning with weak supervision. *KI - Künstliche Intelligenz*, 34:165–180, 2020.
- [10] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. 2020.

- [11] J. Cao, H. Tang, H. Fang, X. Shen, Y. Tai, and C. Lu. Cross-domain adaptation for animal pose estimation. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9497–9506, Los Alamitos, CA, USA, 2019. IEEE Computer Society. doi: 10.1109/ICCV.2019.00959. URL <https://doi.ieeecomputersociety.org/10.1109/ICCV.2019.00959>.
- [12] Z. Cao, G. Hidalgo, T. Simon, S. E. Wei, and Y. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1):172–186, 2018. doi: 10.1109/TPAMI.2019.2929257.
- [13] Francisco M. Castro, Manuel J. Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 241–257, Cham, 2018. Springer International Publishing. ISBN 978-3-030-01258-8.
- [14] Arslan Chaudhry, Puneet Dokania, Thalaiyasingam Ajanthan, and Philip Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. 01 2018.
- [15] ZhiHang Chen, Liping Huang, and Yi L. Murphey. Incremental learning for text document classification. In *2007 International Joint Conference on Neural Networks*, pages 2592–2597, 2007. doi: 10.1109/IJCNN.2007.4371367.
- [16] H. Fang, S. Xie, Y. Tai, and C. Lu. Rmpe: Regional multi-person pose estimation. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2353–2362, 2017. doi: 10.1109/ICCV.2017.256.
- [17] Pasi Fränti and Sami Sieranoja. How much can k-means be improved by using better initialization and repeats? *Pattern Recognition*, 93:95–112, 2019. ISSN 0031-3203. doi: <https://doi.org/10.1016/j.patcog.2019.04.014>. URL <https://www.sciencedirect.com/science/article/pii/S0031320319301608>.
- [18] Boqing Gong, W.-L Chao, K. Grauman, and Fei Sha. Diverse sequential subset selection for supervised video summarization. *Advances in Neural Information Processing Systems*, 3:2069–2077, 01 2014.
- [19] Jacob Graving, Daniel Chae, Hemal Naik, Liang Li, Benjamin Koger, Blair Costelloe, and Iain Couzin. Deepposekit, a software toolkit for fast and robust animal pose estimation using deep learning. 04 2019. doi: 10.1101/620245.
- [20] Semih Günel, Helge Rhodin, Daniel Morales, João Campagnolo, Pavan Ramdya, and Pascal Fua. Deepfly3d: A deep learning-based approach for 3d limb and appendage tracking in tethered, adult drosophila. 05 2019. doi: 10.1101/640375.
- [21] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017. doi: 10.1109/ICCV.2017.322.
- [22] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7122–7131, 2018. doi: 10.1109/CVPR.2018.00744.

- [23] Pierre Karashchuk, Katie L. Rupp, Eryn S. Dickinson, Elischa Sanders, Eiman Azim, Bingni W. Brunton, and John C. Tuthill. Anipose: a toolkit for robust markerless 3d pose estimation. *bioRxiv*, 2020. doi: 10.1101/2020.05.26.117325. URL <https://www.biorxiv.org/content/early/2020/05/29/2020.05.26.117325>.
- [24] D. King. Applications for real time image warping. In *Southcon/96 Conference Record*, pages 298–302, 1996. doi: 10.1109/SOUTHCON.1996.535083.
- [25] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114, 12 2016. doi: 10.1073/pnas.1611835114.
- [26] Chun-Wa Ko, Jon Lee, and Maurice Queyranne. An exact algorithm for maximum entropy sampling. *Operations Research*, 43(4):684–691, 1995. ISSN 0030364X, 15265463. URL <http://www.jstor.org/stable/171694>.
- [27] Alex Kulesza. Determinantal point processes for machine learning. *Foundations and Trends® in Machine Learning*, 5(2-3):123–286, 2012. ISSN 1935-8245. doi: 10.1561/22000000044. URL <http://dx.doi.org/10.1561/22000000044>.
- [28] Alex Kulesza and Ben Taskar. Structured determinantal point processes. In J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc., 2010. URL <https://proceedings.neurips.cc/paper/2010/file/1f50893f80d6830d62765ffad7721742-Paper.pdf>.
- [29] Alex Kulesza and Ben Taskar. k -dpps: Fixed-size determinantal point processes. In *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML’11*, page 1193–1200, Madison, WI, USA, 2011. Omnipress. ISBN 9781450306195.
- [30] D. Lee and B. Schachter. Two algorithms for constructing a delaunay triangulation. *International Journal of Parallel Programming*, 9:219–242, 06 1980. doi: 10.1007/BF00977785.
- [31] Timothy E. Lee, Jonathan Tremblay, Thang To, Jia Cheng, Terry Mosier, Oliver Kroemer, Dieter Fox, and Stan Birchfield. Camera-to-robot pose estimation from a single image, 2020.
- [32] Chen Li and Gim Hee Lee. From synthetic to real: Unsupervised domain adaptation for animal pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1482–1491, June 2021.
- [33] Jia Li, Wen Su, and Zengfu Wang. Simple pose: Rethinking and improving a bottom-up approach for multi-person pose estimation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07):11354–11361, 2020. doi: 10.1609/aaai.v34i07.6797. URL <https://ojs.aaai.org/index.php/AAAI/article/view/6797>.

- [34] Shuyuan Li, Jianguo Li, Hanlin Tang, Rui Qian, and Weiyao Lin. Atrw: A benchmark for amur tiger re-identification in the wild. *Proceedings of the 28th ACM International Conference on Multimedia*, 2020. doi: 10.1145/3394171.3413569. URL <http://dx.doi.org/10.1145/3394171.3413569>.
- [35] Siyuan Li, Semih Gunel, Mirela Ostrek, Pavan Ramdya, Pascal Fua, and Helge Rhodin. Deformation-aware unpaired image translation for pose estimation on laboratory animals. pages 13155–13165, 06 2020. doi: 10.1109/CVPR42600.2020.01317.
- [36] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, 2018. doi: 10.1109/TPAMI.2017.2773081.
- [37] O. Lie and A. Toet. Applications of digital image warping in surveillance and navigation. *Displays*, 19:133–139, 1998.
- [38] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 740–755, Cham, 2014. Springer International Publishing. ISBN 978-3-319-10602-1.
- [39] M. Loper, Naureen Mahmood, J. Romero, Gerard Pons-Moll, and Michael J. Black. Smpl: a skinned multi-person linear model. *ACM Trans. Graph.*, 34:248:1–248:16, 2015.
- [40] Diogo C. Luvizon, David Picard, and Hedi Tabia. 2d/3d pose estimation and action recognition using multitask deep learning. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5137–5146, 2018. doi: 10.1109/CVPR.2018.00539.
- [41] J. MacQueen. Some methods for classification and analysis of multivariate observations. 1967.
- [42] Alexander Mathis, Pranav Mamidanna, Kevin M. Cury, Taiga Abe, V. Murthy, M. Mathis, and M. Bethge. Deeplabcut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, 21:1281–1289, 2018.
- [43] Mackenzie Weygandt Mathis and Alexander Mathis. Deep learning tools for the measurement of animal behavior in neuroscience. *Current Opinion in Neurobiology*, 60:1–11, 2020. ISSN 0959-4388. doi: <https://doi.org/10.1016/j.conb.2019.10.008>. URL <https://www.sciencedirect.com/science/article/pii/S0959438819301151>. *Neurobiology of Behavior*.
- [44] Jiteng Mu, Weichao Qiu, Gregory D. Hager, and Alan L. Yuille. Learning from synthetic animals. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12383–12392, 2020. doi: 10.1109/CVPR42600.2020.01240.
- [45] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 483–499, Cham, 2016. Springer International Publishing. ISBN 978-3-319-46484-8.

- [46] Talmo D Pereira, Diego E Aldarondo, Lindsay Willmore, Mikhail Kislin, Samuel S-H Wang, Mala Murthy, and Joshua W Shaevitz. Fast animal pose estimation using deep neural networks. *Nat Methods*, 16:117–125, 2019. ISSN 1548-7105. doi: 10.1038/s41592-018-0234-5.
- [47] Ameya Prabhu, Philip Torr, and Puneet Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *The European Conference on Computer Vision (ECCV)*, August 2020.
- [48] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. icarl: Incremental classifier and representation learning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5533–5542, 2017. doi: 10.1109/CVPR.2017.587.
- [49] Nekkanti Sailaja, L. Padmasree, and Nimmala Mangathayaru. Incremental learning for text categorization using rough set boundary based optimized support vector neural network. *Data Technologies and Applications*, ahead-of-print, 07 2020. doi: 10.1108/DTA-03-2020-0071.
- [50] A. Sanakoyeu, Vasil Khalidov, Maureen S. McCarthy, A. Vedaldi, and N. Neverova. Transferring dense pose to proximal animal classes. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5232–5241, 2020.
- [51] Guangxu Shan, Shiyao Xu, Li Yang, Shengbin Jia, and Yang Xiang. Learn#: A novel incremental learning method for text classification. *Expert Systems with Applications*, 147:113198, 2020. ISSN 0957-4174. doi: <https://doi.org/10.1016/j.eswa.2020.113198>. URL <https://www.sciencedirect.com/science/article/pii/S0957417420300245>.
- [52] K. Shmelkov, C. Schmid, and Alahari Karteek. Incremental learning of object detectors without catastrophic forgetting. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 3420–3429, 2017.
- [53] M. Svenstrup, S. Tranberg, H. J. Andersen, and T. Bak. Pose estimation and adaptive robot behaviour for human-robot interaction. In *2009 IEEE International Conference on Robotics and Automation*, pages 3571–3576, 2009. doi: 10.1109/ROBOT.2009.5152690.
- [54] A. Telea. An image inpainting technique based on the fast marching method. *Journal of Graphics Tools*, 9:23 – 34, 2004.
- [55] Alexander Toshev and Christian Szegedy. Deeppose: Human pose estimation via deep neural networks. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1653–1660, 2014. doi: 10.1109/CVPR.2014.214.
- [56] Ruoxi Wang, Yingzhou Li, and Eric Darve. On the numerical rank of radial basis function kernels in high dimensions. *SIAM Journal on Matrix Analysis and Applications*, 39(4):1810–1835, 2018. ISSN 1095-7162. doi: 10.1137/17m1135803. URL <http://dx.doi.org/10.1137/17M1135803>.

- [57] Max Welling. Herding dynamical weights to learn. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09*, page 1121–1128. Association for Computing Machinery, 2009. ISBN 9781605585161. doi: 10.1145/1553374.1553517. URL <https://doi.org/10.1145/1553374.1553517>.
- [58] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. 04 2018. ISBN 978-3-030-01230-4. doi: 10.1007/978-3-030-01231-1_29.
- [59] Junting Zhang, Jie Zhang, Shalini Ghosh, Dawei Li, Serafettin Tasci, Larry Heck, Heming Zhang, and C.-C. Jay Kuo. Class-incremental learning via deep model consolidation. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1120–1129, 2020. doi: 10.1109/WACV45572.2020.9093365.
- [60] Teng Zhang, Liangchen Liu, Kun Zhao, Arnold Wiliem, Graham Hemson, and Brian Lovell. Omni-supervised joint detection and pose estimation for wild animals. *Pattern Recognition Letters*, 132:84–90, 2020. ISSN 0167-8655. doi: <https://doi.org/10.1016/j.patrec.2018.11.002>. URL <https://www.sciencedirect.com/science/article/pii/S0167865518308742>. Multiple-Task Learning for Big Data (MTL4BD).
- [61] Silvia Zuffi, Angjoo Kanazawa, and Michael J. Black. Lions and tigers and bears: Capturing non-rigid, 3d, articulated shape from images. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3955–3963, 2018. doi: 10.1109/CVPR.2018.00416.