

# Few-shot Keypoint Detection with Uncertainty Learning for Unseen Species

Changsheng Lu<sup>†</sup>, Piotr Koniusz<sup>\*,§,†</sup>

<sup>†</sup>The Australian National University    <sup>§</sup>Data61/CSIRO

ChangshengLuu@gmail.com, firstname.lastname@anu.edu.au

## Abstract

Current non-rigid object keypoint detectors perform well on a chosen kind of species and body parts, and require a large amount of labelled keypoints for training. Moreover, their heatmaps, tailored to specific body parts, cannot recognize novel keypoints (keypoints not labelled for training) on unseen species. We raise an interesting yet challenging question: how to detect both base (annotated for training) and novel keypoints for unseen species given a few annotated samples? Thus, we propose a versatile Few-shot Keypoint Detection (FSKD) pipeline, which can detect a varying number of keypoints of different kinds. Our FSKD provides the uncertainty estimation of predicted keypoints. Specifically, FSKD involves main and auxiliary keypoint representation learning, similarity learning, and keypoint localization with uncertainty modeling to tackle the localization noise. Moreover, we model the uncertainty across groups of keypoints by multivariate Gaussian distribution to exploit implicit correlations between neighboring keypoints. We show the effectiveness of our FSKD on (i) novel keypoint detection for unseen species, (ii) few-shot Fine-Grained Visual Recognition (FGVR) and (iii) Semantic Alignment (SA) downstream tasks. For FGVR, detected keypoints improve the classification accuracy. For SA, we showcase a novel thin-plate-spline warping that uses estimated keypoint uncertainty under imperfect keypoint correspondences.

## 1. Introduction

Deep fully-supervised pose estimation has proven its excellence in detecting keypoints on humans [6, 14, 18, 35, 47]. However, these keypoint detectors are dedicated to specific species and body parts on which they are trained. They are not reusable for unseen species, and consume large amounts of annotated data. In contrast, given a few samples, a child can adequately recognize and generalize a keypoint on a paw of rabbit, cat, dog, kangaroo under varying poses *e.g.*, jumping, crouching, or walking. By comparison, machine

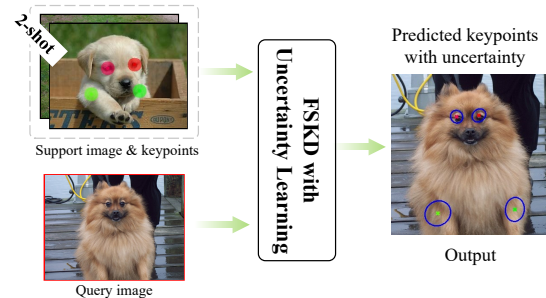


Figure 1. Illustration of few-shot keypoint detection (4-way, 2-shot protocol) with uncertainty learning. Our FSKD model successfully detects four novel keypoints and estimates the localization uncertainty given two annotated instances in unseen species.

perception is inferior to biological perception [51]. Inspired by the progress in few-shot learning [25, 26, 41, 45, 49, 58], we propose Few-shot Keypoint Detection (FSKD) which can learn from few keypoints and recognise previously unseen keypoint types even for species unseen during training (Fig. 1).

As keypoints provide crucial structural and semantic information, FSKD has numerous promising applications such as semi-automatic labelling, face alignment [27, 44], Fine-Grained Visual Recognition (FGVR) [46], animal behavior analysis [37], *etc.* The use of keypoints can also simplify the labor-intensive surveillance of wildlife [34, 56].

In this paper, we propose a versatile FSKD approach which copes with various levels of domain shift. Fig. 2 shows that *the categories of species used for training and testing may be the same or different (top branches), and also the keypoint types of specific body parts may be the same or disjoint (bottom branches)*. The easiest problem setting assumes the same kind of species and the same types of keypoints throughout training and testing. However, if the species as well as keypoint types used for training and testing are both disjoint, the problem is challenging due to high levels of domain shift.

We note that learning and generalizing based on a few of samples is hard due to the limited number of annotations and a large variability of samples. Moreover, there exist

\*The corresponding author.

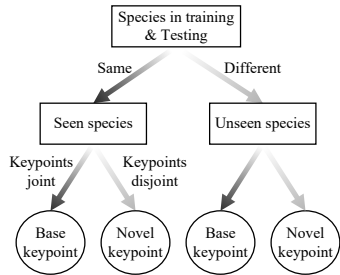


Figure 2. Different settings of FSKD lead to different levels of domain shift.

large amounts of interfering noise and similar local patterns in an image, which can challenge FSKD and deem it arguably a much harder problem than Few-shot Object Detection (FSOD) [13, 22, 55]. In contrast to objects which usually have well-defined boundaries, keypoints lack the clear boundary and correspond to some local regions rather than exact coordinates. Consequently, keypoints inherently exhibit ambiguity and location uncertainty, which are reflected in both groundtruth keypoints and predictions. Thus, we develop an FSKD approach which can deal with domain shifts and model the uncertainty of localization.

Our FSKD firstly extracts the deep representations of support keypoints to build the keypoint prototypes (one per keypoint type), which are correlated against the query feature map to yield keypoint-specific attentive features. After descriptor extraction, attentive features are transformed into descriptors which will be used for subsequent keypoint localization. To address the limited types of base training keypoints, we introduce the generated auxiliary keypoints into learning. Though these keypoints show poor matching between support and query images, they boost the keypoint diversity and significantly help infer novel keypoints. The difficulty of FSKD results in imperfect keypoint predictions over novel keypoints. To compensate for this effect and deal with inherent ambiguity and noise of keypoints, we propose to model the localization uncertainty by learning the covariance for individual or multiple keypoints, which allows larger tolerance of those noisy keypoints by the loss function. Our contributions are summarized as follows:

- i. We propose a flexible few-shot keypoint detection (FSKD) model that can detect varying types and numbers of keypoints given one or more annotated samples.
- ii. Both localization and semantic uncertainty are modeled within our localization networks, where an uncertainty-aided grid-based locator (UC-GBL) is proposed. Moreover, we propose the multi-scale UC-GBL in order to reduce the risk of mislocalization.
- iii. We employ the low-quality auxiliary keypoints during learning, and model the covariance for coupled main and auxiliary keypoints to improve generalization.

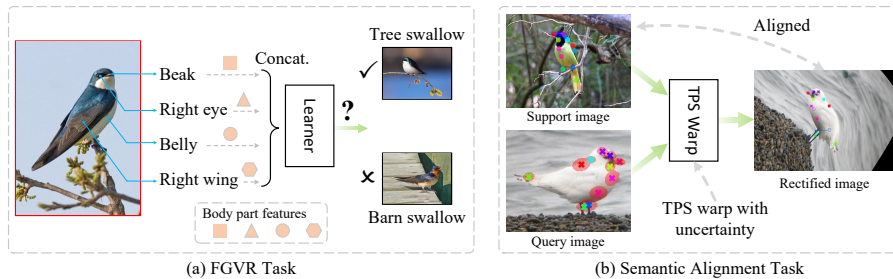


Figure 3. Downstream tasks on which we evaluate the proposed FSKD. (a) Few-shot Fine-Grained Visual Recognition (FGVR); (b) Semantic Alignment (SA).

- iv. Convincing experiments show that our FSKD can detect novel keypoints on unseen species. With a simple modification, we extend our FSKD to act as a simple keypoint detector applied to few-shot fine-grained visual recognition and semantic alignment (Fig. 3).

To our best knowledge, our work is the first attempt to model keypoint detection as few-shot learning.

## 2. Related Work

**Few-shot Learning (FSL):** Initially, FSL was dedicated to image classification based on few samples [15, 51]. The current FSL pipelines are based on deep backbones and mainly focus on i) metric learning [25, 42, 45, 49]; ii) optimization, *e.g.*, MAML [16] rapidly adapts to new tasks; and iii) data hallucination [54, 57]. FSL has expanded into other computer vision tasks including few-shot segmentation [28, 30] and object detection [13, 22, 55, 57].

**Keypoint Detection:** Compared to traditional keypoint detectors [11, 32], the deep learning methods are more general and can be categorized into two kinds, where the first kind uses the heatmap regression followed by post-processing to search the keypoint with maximal heat value [6, 14, 18, 35, 44], and the second kind is directly performing regression on the keypoint position [7, 47], which is also adopted in our work. As for the heatmap regression based approaches, they can be further divided into top-down [14, 18, 43] and bottom-up pose estimators [6, 8, 35]. Recently, some works [5, 29] perform cross-domain adaptation or shape deformation by leveraging the large-scale source datasets. Novotny *et al.* [36] use self-supervision to learn matching features between a pair of transformed images. Though these approaches can limit annotation burden when learning, they cannot directly detect novel keypoints on unseen species based on a few samples.

**Uncertainty in Computer Vision:** The uncertainty generally consists of *aleatoric uncertainty* and *epistemic uncertainty* [2, 23] modeled by a Gaussian over the predictions and placing a distribution over the model weights (*i.e.*, Bayesian Neural Network [2]), respectively. In this paper, we mainly focus on the heteroscedastic aleatoric un-

certainly, popular in a number of applications. Kendall *et al.* [24] use uncertainty to relatively weigh multi-task loss functions between depth regression and segmentation. He *et al.* [20] and Choi *et al.* [9] incorporate uncertainty into bounding box regression of Faster R-CNN [39] and YOLOv3 [38]. However, these models treat multiple variables independently, while we model uncertainty by covariance to capture underlying relations between variables.

### 3. Few-shot Keypoint Detection

#### 3.1. Architecture Overview

Given the sampled support image and keypoints, FSKD aims to detect the corresponding keypoints in query image, where the hardest setting includes species and keypoint types that are both disjoint between training and testing. For  $N$  support keypoints and  $K$  support images, the problem is dubbed as  $N$ -way- $K$ -shot detection.

An overview of our FSKD pipeline is shown in Fig. 4, which includes feature encoder  $\mathcal{F}$ , feature modulator  $\mathcal{M}$ , descriptor extractor  $\mathcal{P}$ , and uncertainty-aided grid-based locator  $\mathcal{G}$  (UC-GBL). Moreover, we also learn the semantic distinctiveness (SD) by adding a side branch  $\mathcal{D}$  (SD head). In the following, we will describe each module in detail.

#### 3.2. Keypoint Prototypes and Descriptors Building

**Keypoint Embedding:** Let support and query images  $\mathbf{I}_s$  and  $\mathbf{I}_q$  be mapped to  $\mathcal{F}(\mathbf{I}_s)$  and  $\mathcal{F}(\mathbf{I}_q)$  in feature space  $\mathbb{R}^{l \times l \times C}$  using a weight-shared convolutional encoder  $\mathcal{F}$ . Given a support keypoint at location  $\mathbf{u} = [x, y]^T$ , we extract the keypoint representation  $\Phi = \mathcal{A}(\mathcal{F}(\mathbf{I}_s), \mathbf{u})$  via  $\mathcal{A}$ . Operator  $\mathcal{A}$  could be integer-based indexing [36], bilinear interpolation [18, 21], or Gaussian pooling. Feature representations  $\Phi$  must contain some local context of keypoint to be sufficiently discriminative during the matching step. Therefore, we employ Gaussian pooling. Let  $\mathbf{u}_{k,n}$  be the  $n$ -th keypoint in the  $k$ -th support image  $\mathbf{I}_s^k$  and  $\mathcal{F}(\mathbf{I}_s^k)(\mathbf{x})$  be the feature vector at position  $\mathbf{x} \in \mathbb{R}^2$ . The keypoint representation  $\Phi_{k,n}$  can be obtained as

$$\Phi_{k,n} = \sum_{\mathbf{x}} \exp(-\|\mathbf{x} - \mathbf{u}_{k,n}\|_2^2 / 2\xi^2) \cdot \mathcal{F}(\mathbf{I}_s^k)(\mathbf{x}), \quad (1)$$

where  $\xi$  is the standard deviation. Following prototypical networks [42], each support keypoint prototype (SKP)  $\Phi_n$  can be obtained by averaging the same type of keypoint representations across support images as

$$\Phi_n = \frac{1}{K} \sum_{k=1}^K \Phi_{k,n}, \quad (2)$$

where  $n = 1, \dots, N$ . Thus we have  $N$  SKPs generated in the  $N$ -way  $K$ -shot learning process.

**Support and Query Correlation:** In order to guide FSKD to discover the corresponding keypoints in the query image, SKP  $\Phi_n$  needs to be correlated against query feature

map  $\mathcal{F}(\mathbf{I}_q)$ . To this end, we adopt a feature modulator  $\mathcal{M}$ , which takes both  $\Phi_n$  and  $\mathcal{F}(\mathbf{I}_q)$  as input to produce the attentive features  $\mathbf{A}_n = \mathcal{M}(\mathcal{F}(\mathbf{I}_q), \Phi_n), n = 1, 2, \dots, N$ . For  $\mathcal{M}$  we could choose simple correlation, spatial attention, channel attention, *etc.* In this paper, we use simple correlation for its efficiency, which is given as  $\mathbf{A}_n(\mathbf{p}) = \mathcal{F}(\mathbf{I}_q)(\mathbf{p}) \odot \Phi_n$ , where  $\mathbf{p} \in \mathbb{R}^2$  runs over the feature map of size  $l \times l$ , and  $\odot$  is the channel-wise multiplication. Feature correlation performs the similarity learning between support and query, which ‘activates’ the local regions in the query feature map that significantly correlate with SKP.

**Keypoint Descriptor Extraction:** After performing feature correlation, we then project each attentive feature  $\mathbf{A}_n$  into keypoint descriptor  $\Psi_n$  via a descriptor extractor  $\mathcal{P}$  to decrease the dimensionality, namely,  $\Psi_n = \mathcal{P}(\mathbf{A}_n)$ .

#### 3.3. Keypoint Localization & Uncertainty Learning

##### 3.3.1 Vanilla Grid-based Locator (GBL)

Instead of regressing global location for each keypoint descriptor  $\Psi_n$ , we approach the keypoint localization as grid classification and local offset regression, where a grid classifier  $\mathcal{G}_c$  and an offset regressor  $\mathcal{G}_o$  are used. Let  $\mathbf{v} \in \mathbb{R}^2$  be the offset to the center of a grid where the keypoint falls. The predicted grid  $\mathbf{g} \in \{0, \dots, S-1\} \times \{0, \dots, S-1\}$  is obtained by the 2D index of the maximum in the grid probability map  $\mathbf{P} \in \mathbb{R}^{S \times S}$ , where  $\mathbf{P} = \text{Softmax}(\mathcal{G}_c(\Psi_n))$ , and the predicted offset  $\mathbf{v}$  can be retrieved from the vector field  $\mathcal{G}_o(\Psi_n) \in \mathbb{R}^{S \times S \times 2}$  at grid  $\mathbf{g}$  (see Fig. 4). In addition, we construct the groundtruth (GT) offset  $\mathbf{v}^*$  via

$$\mathbf{t} = \mathbf{u}^* S / l_0 \quad \mathbf{z} = \lfloor \mathbf{t} \rfloor + 0.5 \quad \mathbf{v}^* = 2(\mathbf{t} - \mathbf{z}), \quad (3)$$

where  $\mathbf{u}^* \in \mathbb{R}^2$  is a GT keypoint in square-padded query image with edge length  $l_0$ ,  $\mathbf{t} \in \mathbb{R}^2$  is the transformed coordinate in the grid frame, and  $\mathbf{z}$  is the grid center. Furthermore, the grid label can be formed as  $g^* = \lfloor \mathbf{t}_y \rfloor S + \lfloor \mathbf{t}_x \rfloor$ , and  $\mathbf{v}^* \in [-1, 1)^2$ . With  $\mathbf{v}^*$  and  $g^*$ , we design a cross-entropy grid classification loss  $\mathcal{L}_{cls}$ , and an offset regression loss  $\mathcal{L}_{os}$  using MSE to minimize the vanilla GBL ( $\mathcal{G}_c, \mathcal{G}_o$ ).

##### 3.3.2 Localization and Semantic Uncertainty

Unlike approaches [9, 20, 24, 33], we use covariance  $\Sigma$  to model the localization uncertainty of individual or multiple keypoints. Let  $\mathcal{N}(\mathbf{x}; \mathbf{v}^*, \Sigma)$  be the multivariate Gaussian distribution with  $\mathbf{x}, \mathbf{v}^* \in \mathbb{R}^k, k \geq 2$ , and  $\Sigma \in \mathbb{S}_{++}^k$ . Let  $\mathbf{x}$  be the predicted offset for a keypoint (or stacked multiple keypoints) and  $\mathbf{v}^*$  be the GT, then one may write the negative log-likelihood (NLL) loss as

$$\begin{aligned} \mathcal{L}_{os-nll} &= -\mathbb{E} \log \mathcal{N}(\mathbf{x}; \mathbf{v}^*, \Sigma) \\ &\equiv \frac{1}{2} \mathbb{E} [(\mathbf{x} - \mathbf{v}^*)^T \Sigma^{-1} (\mathbf{x} - \mathbf{v}^*) + \log \det(\Sigma)]. \end{aligned} \quad (4)$$

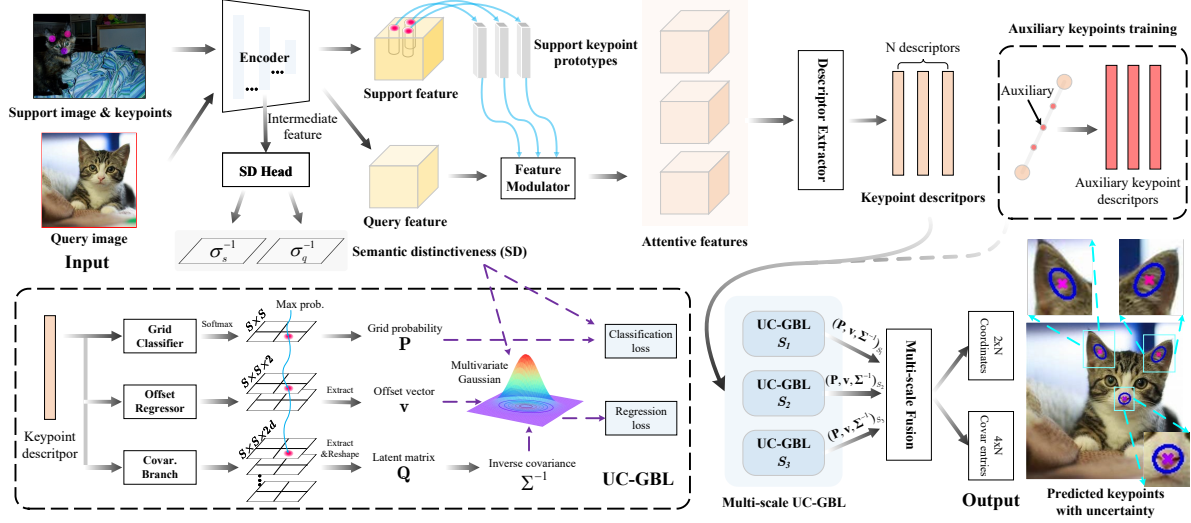


Figure 4. Few-shot keypoint detection pipeline. The whole model aims to predict the keypoints with uncertainty in a query image given the support keypoints. The prototypes guide the query feature to generate keypoint-specific attentive features, which will be condensed into descriptors for keypoint localization through multi-scale UC-GBL. Via interpolation, the auxiliary keypoints with poor matching quality are also incorporated into the learning process to boost the generalization ability of FSKD.

However,  $\mathcal{L}_{\text{os-nll}}$  in Eq. 4 relies on the computation of the inverse of covariance matrix  $\Sigma$ , which is costly and unstable in back-propagation especially when  $k \geq 4$ . Thus, we replace  $\Sigma$  with the precision matrix  $\Omega = \Sigma^{-1}$ , and obtain

$$\mathcal{L}_{\text{os-nll}} = \frac{1}{2} \mathbb{E}[(\mathbf{x} - \mathbf{v}^*)^\top \Omega (\mathbf{x} - \mathbf{v}^*) - \log \det(\Omega)]. \quad (5)$$

As a result, the  $\mathcal{L}_{\text{os-nll}}$  can be easily computed as long as  $\Omega \succ \mathbf{0}$ . To guarantee this, we let  $\Omega = \frac{1}{d} \mathbf{Q} \mathbf{Q}^\top$ , where  $\mathbf{Q} \in \mathbb{R}^{k \times d}$  ( $d \geq k$ ) is the latent matrix learned from our covariance branch network. In extreme case, a small  $\epsilon \rightarrow 0$  can be added to ensure  $\det(\Omega) > 0$ .

Firstly, we investigate learning the covariance per keypoint by adding a covariance branch  $\mathcal{G}_v$  to GBL (Fig. 4), whose output is the latent covariance field  $\mathcal{G}_v(\Psi_n) \in \mathbb{R}^{S \times S \times 2d}$ . Then  $\mathbf{Q} \in \mathbb{R}^{2 \times d}$  which encodes the covariance information for a given keypoint can be extracted from grid  $\mathbf{g}$ . Secondly, we investigate learning relations for multiple keypoints within group (*i.e.*,  $m$  keypoints per group) by a multi-keypoint covariance branch  $\mathcal{G}_{\text{mkv}}$  whose output is  $\mathbf{Q}_{\text{mkv}} \in \mathbb{R}^{2m \times 2m}$ .

In addition to  $\Sigma$  which reflects the localization uncertainty over keypoints, following [23, 36], we also model the semantic uncertainty  $\sigma$  by learning a single-channel SD map  $\sigma_s^{-1} \in \mathbb{R}^{H \times W \times 1}$  via SD head  $\mathcal{D}$  (Fig. 4) whose values are in range (0, 1). The higher the value, the more distinctive perceptually a keypoint is. Let  $\sigma_s^{-1}$  and  $\sigma_q^{-1}$  be the support and query SD maps. Thus, for each keypoint descriptor  $\Psi_n$ , we extract the corresponding SD scalar as  $w_n = \frac{1}{2}(\sigma_{s, \mathbf{u}_n}^{-1} + \sigma_{q, \mathbf{u}'_n}^{-1})$ , where  $\sigma_{s, \mathbf{u}_n}^{-1}$  and  $\sigma_{q, \mathbf{u}'_n}^{-1}$  are support and query SD map values at keypoint locations  $\mathbf{u}_n$  and  $\mathbf{u}'_n$ .

Let  $\mathbf{W} = \text{diag}([w_1, w_1, \dots, w_m, w_m]) \in \mathbb{S}_{++}^{2m}$  be a diagonal matrix with entries  $w_n$ . We introduce  $\mathbf{W}$  into Eq. 5:

$$\mathcal{L}_{\text{uc}} = \frac{1}{2} \mathbb{E}[(\mathbf{x} - \mathbf{v}^*)^\top (\Omega + \beta \mathbf{W}) (\mathbf{x} - \mathbf{v}^*) - \log(\det(\Omega \mathbf{W}^\beta))], \quad (6)$$

where  $\beta$  is a trade-off (we set  $\beta = 1$ ). We also use  $w_n$  to reweight the cross-entropy grid classification loss  $\mathcal{L}_{\text{cls-uc}} = -\mathbb{E}[\sqrt{w_n} \langle \mathbb{I}(\mathbf{g}^*), \log(\text{vec}(\mathbf{P})) \rangle]$  where  $\text{vec}(\cdot)$  vectorizes a matrix,  $\mathbb{I}(\cdot)$  is a one-hot encoding of scalar. Both  $\mathcal{L}_{\text{uc}}$  and  $\mathcal{L}_{\text{cls-uc}}$  are used by the uncertainty-aided GBL (UC-GBL).

Compared to vanilla GBL, our proposed UC-GBL has a couple of advantages: 1) The model enjoys a larger tolerance over prediction and label noise, and reduces the impact of noise which degrades the learning performance as the learned  $\Sigma$  and  $\mathbf{W}$  can serve as attenuation; 2) eigenvalues and eigenvectors of  $\Sigma$  provide the localization uncertainty.

### 3.3.3 Multi-scale UC-GBL and Uncertainty Fusion

Increasing the scale  $S$  will increase the precision of grids but also result in more grids. In order to reduce the risk of mislocalization, we employ multi-scale UC-GBL in our pipeline. Let the loss function at scale  $S$  be  $\mathcal{L}^{(S)} = \alpha_1 \mathcal{L}_{\text{uc}} + \alpha_2 \mathcal{L}_{\text{cls-uc}}$  (we set  $\alpha_1 = \alpha_2 = 1$ ). Then the multi-scale localization loss  $\mathcal{L}_{\text{ms}}$  is formulated as

$$\mathcal{L}_{\text{ms}} = \frac{1}{N_S} \sum_{i=1}^{N_S} \mathcal{L}^{(S_i)}, \quad (7)$$

where  $N_S$  is the number of scales used in FSKD. The unified keypoint prediction  $\mathbf{u}$  is computed as

$$\mathbf{u} = \frac{1}{N_S} \sum_{i=1}^{N_S} \frac{l_0}{S_i} (\mathbf{g}^{(S_i)} + 0.5 + 0.5 \mathbf{v}^{(S_i)}), \quad (8)$$

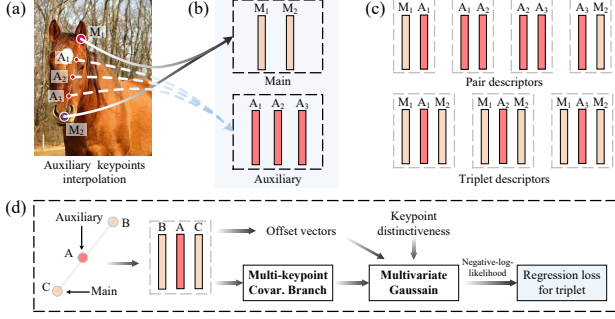


Figure 5. Keypoint grouping strategies and covariance modeling for multiple keypoints. (a) Interpolation with three nodes between body parts; (b) main and auxiliary descriptors; (c) grouping with pairs or triplets; (d) multi-keypoint covariance modeling.

where  $\mathbf{g}^{(S_i)} \in \{0, \dots, S_i - 1\} \times \{0, \dots, S_i - 1\}$  and  $\mathbf{v}^{(S_i)}$  are the predicted 2D grid index and offset at scale  $S_i$ . The localization uncertainty  $\Sigma$  is obtained by

$$\Sigma = \frac{1}{4N_S} \sum_{i=1}^{N_S} \left( \frac{l_0}{S_i} \right)^2 \Sigma^{(S_i)}, \quad (9)$$

where  $\Sigma^{(S_i)}$  is the covariance at scale  $S_i$  obtained by inverting the precision matrix  $\Omega^{(S_i)}$ . A glance of unified estimated keypoints and uncertainty is shown in Fig. 4.

### 3.4. Learning with Auxiliary Keypoints

In addition to the main training keypoints provided by annotations, we adopt the auxiliary keypoints into learning, which are generated via interpolation  $\mathcal{T}(t; [\mathbf{u}_1, \mathbf{u}_2])$  on a path whose end points are keypoints  $[\mathbf{u}_1, \mathbf{u}_2]$ , and  $t \in (0, 1)$  is the so-called interpolation node, as shown in Fig. 5(a). We use an off-the-shelf saliency detector [52] to prune auxiliary keypoints that do not lie on the foreground. Similarly, we build auxiliary descriptor  $\Psi_n$  (Fig. 5(b)) and localization loss  $\tilde{\mathcal{L}}_{ms}$  (Eq. 7) for auxiliary keypoints of query image. Even though auxiliary keypoints between support and query do not match well, they add visual diversity beyond appearances of training keypoints. We also group main and auxiliary keypoints into pairs or triplets (Fig. 5(c)) and model the multi-keypoint covariance  $\Sigma_{mkv}$  via a branch  $\mathcal{G}_{mkv}$  (Fig. 5(d)). Thus, we propose a multi-keypoint offset regression loss  $\mathcal{L}_{ms-mk}$ .

### 3.5. Objective Functions

Our pipeline has three loss terms, which are the main keypoint loss  $\mathcal{L}_{ms}$ , auxiliary keypoint loss  $\tilde{\mathcal{L}}_{ms}$ , and multi-keypoint offset regression loss  $\mathcal{L}_{ms-mk}$ .  $\mathcal{L}_{ms-mk}$  is used by default with the triplet grouping strategy. The overall loss ( $\gamma_1 = \gamma_2 = \gamma_3 = 1$ ) is

$$\mathcal{L} = \gamma_1 \mathcal{L}_{ms} + \gamma_2 \tilde{\mathcal{L}}_{ms} + \gamma_3 \mathcal{L}_{ms-mk}. \quad (10)$$

## 4. Experiments

### 4.1. Experiment Setup

**Datasets:** 1) *Animal pose* [5] dataset has over 6,000 annotated instances with 5 species *cat, dog, cow, horse, and sheep*. 17 body parts are used in our experiments; 2) *CUB* [50] consists of 200 species with 15 keypoint annotations; 3) *NABird* [48] has 555 categories, 11 keypoint annotations, and 48,562 images. For each dataset, we select the following keypoint types into the *novel keypoint set*: i) *two eyes and four knees* for animal pose dataset; ii) *forehead, two eyes, and two wings* for CUB; iii) *two eyes and two wings* for NABird. The remaining keypoints are used by the *base keypoint set*. See §B of **Suppl. Material** (table with splits).

**Metric:** We use the percentage of correct keypoints (PCK) as evaluation metric [36, 53]. The distance of keypoint to GT should be less than  $\tau \cdot \max(w_{bbx}, h_{bbx})$ , where  $w_{bbx}$  and  $h_{bbx}$  are the edges of object bounding box. We set  $\tau = 0.1$ .

**Compared Methods:** We adapt for comparison the work of Novotny *et al.* [36] who developed the probabilistic introspection matching loss (*ProbIntr*). Moreover, we build a *baseline* using vanilla GBL at scale of  $S = 8$  (no auxiliary keypoints used, no uncertainty, *etc.*). For our FSKD approach, we introduce two variants for comparisons, which are *FSKD (rand)* and *FSKD (default)*. They share the architecture but differ by the types of interpolation path. The former, in each episode, randomly selects a number of paths formed by arbitrary two body parts for interpolation, while the latter interpolates auxiliary keypoints on limbs.

**Implementation Details:** Feature encoder  $\mathcal{F}$  uses modified ResNet50 [19]. Three UC-GBLs with scale  $S = \{8, 12, 16\}$  are stacked to form the multi-scale UC-GBL. Random flips and rotations are used for data augmentation, and the image is resized and padded to  $384 \times 384$ . For auxiliary keypoints generation, the number of interpolation paths is set to 6 and the so-called interpolation nodes  $t = \{0.25, 0.5, 0.75\}$ . All the models are trained with 80k episodes in the animal pose dataset and 40k episodes in CUB and NABird. The results are reported over 1000 episodes in testing.

### 4.2. Experiments on Few-shot Keypoint Detection

Firstly, we set experiments on detecting novel keypoints for unseen species (Fig. 2). For animal pose dataset, we alternately select one animal kind as unseen species for testing while the remaining four as seen species for training (five leave-one-out subproblems). For CUB, 100 species are for training, 50 for validation, and 50 for testing. In terms of NABird, the split is 333, 111, and 111 for training, validation and testing respectively. We use all-way 1-shot protocol on all tasks. In each episode, there is one support image and all base (or novel) keypoints are used as support keypoints. The same-species episode is used. Secondly, we explore detecting novel keypoints for seen species, where 70%

Table 1. Results on 1-shot keypoint detection for unseen or seen species across three datasets. The PCK scores are reported.

Species	Setting	Method	Animal Pose Dataset					CUB	NABird	
			Cat	Dog	Cow	Horse	Sheep			Avg
Unseen	Novel	Baseline	27.30	24.40	19.40	18.25	21.22	22.11	66.12	39.14
		ProbIntr	28.54	23.20	19.55	17.94	17.03	21.25	68.07	48.70
		FSKD (rand)	46.05	40.66	37.55	38.09	31.50	38.77	<b>77.90</b>	54.01
		FSKD (default)	<b>52.36</b>	<b>47.94</b>	<b>44.07</b>	<b>42.77</b>	<b>36.60</b>	<b>44.75</b>	77.89	<b>56.04</b>
	Base	Baseline	51.08	40.44	45.27	35.72	43.03	43.11	81.16	75.74
		ProbIntr	45.96	42.49	37.87	40.53	37.04	40.78	73.46	70.56
		FSKD (rand)	<b>57.12</b>	51.12	47.83	49.71	43.71	49.90	<b>87.94</b>	<b>87.84</b>
		FSKD (default)	56.38	<b>51.29</b>	<b>48.24</b>	<b>49.77</b>	<b>43.95</b>	<b>49.93</b>	87.71	86.99
Seen	Novel	Baseline	29.41	24.43	19.95	19.59	21.95	23.07	67.56	43.52
		ProbIntr	26.09	21.44	19.71	16.95	17.83	20.40	64.13	46.71
		FSKD (rand)	55.31	44.08	39.80	41.52	32.32	42.61	78.11	56.33
		FSKD (default)	<b>60.84</b>	<b>53.44</b>	<b>47.78</b>	<b>49.21</b>	<b>38.47</b>	<b>49.95</b>	<b>78.17</b>	<b>58.35</b>
	Base	Baseline	62.30	49.33	51.33	42.98	44.18	50.02	84.02	75.92
		ProbIntr	57.23	48.58	42.65	48.70	36.15	46.66	76.57	70.47
		FSKD (rand)	67.55	57.54	<b>53.47</b>	57.40	44.80	56.15	87.75	87.88
		FSKD (default)	<b>68.66</b>	<b>59.24</b>	52.70	<b>58.53</b>	<b>45.04</b>	<b>56.83</b>	<b>90.80</b>	<b>88.16</b>

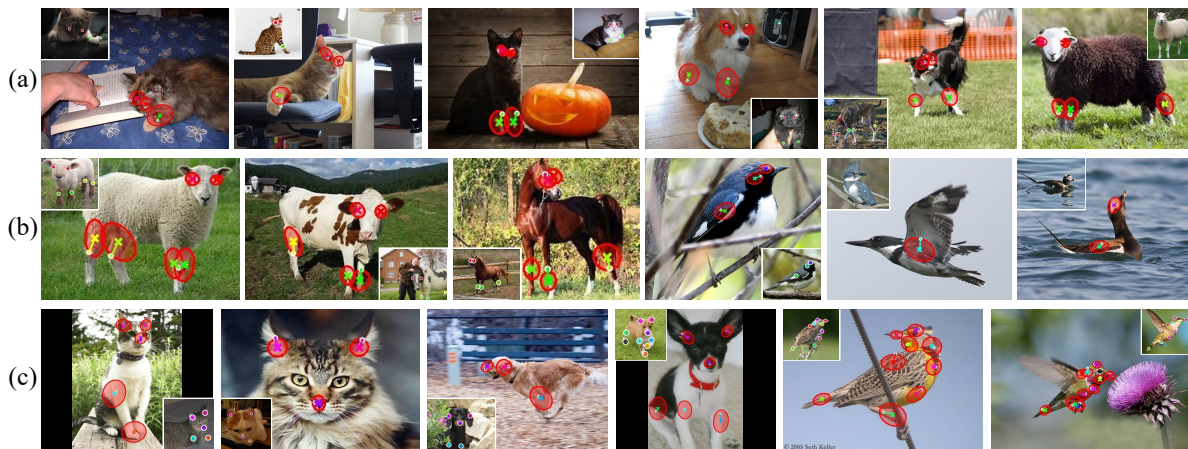


Figure 6. Examples of 1-shot keypoint detection for unseen species. The small image on the corner is the support while the big image is the query. The support keypoints and GT query keypoints are shown by circle dots. Our FSKD keypoint prediction is shown by tilted cross centered by an ellipse which represents the localization uncertainty with 99.7% confidence. Not all support keypoints have the corresponding keypoints in query image as the GT may not exist. The line segment that connects from prediction to GT reveals the localization error. The rows (a)~(b) show the results on detecting novel keypoints while row (c) shows results for base keypoints.

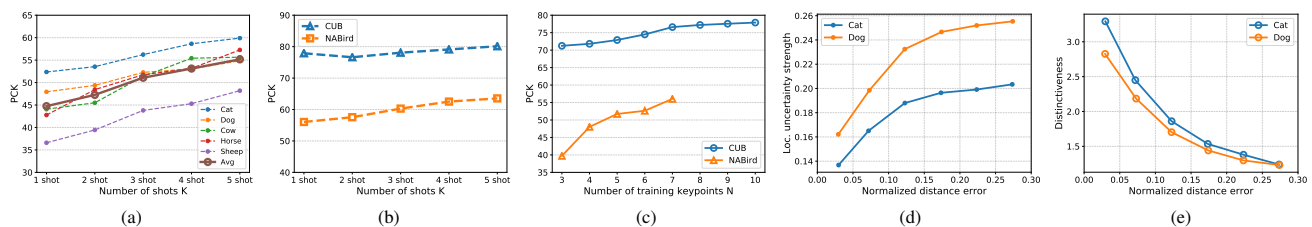


Figure 7. Study on shots, number of training keypoints, and uncertainty. (a)~(b) impact of shots  $K$ ; (c) impact of number of training keypoints  $N$ ; (d) and (e) show the localization uncertainty strength  $J'$  and keypoint distinctiveness  $w$  vs. the normalized distance error  $d'$ .

of images are for training and 30% for testing (all datasets). Thirdly, for the above two settings, we also report testing results for base keypoints (using identical trained model).

Table 1 shows the results (95% confidence intervals are

below 1.2%). We observe that: 1) The baseline is competitive compared to ProbIntr, which may be due to the success of feature modulator and vanilla GBL; 2) our FSKD variants significantly improve the scores in detecting novel

Table 2. FSKD (novel keypoints, unseen classes) on several backbones and 5 datasets. The score (\*27.75) is achieved by *Baseline*.

FSKD with		Animal	CUB	NABird	DeepF2	AwA
<i>ResNet50</i>	[19]	44.75	77.89	56.04	33.04	<b>64.76</b> (*27.75)
<i>HRNet-W32</i>	[43]	47.61	78.24	56.89	33.67	<b>70.99</b>
<i>HRNet-W48</i>	[43]	<b>48.81</b>	<b>79.45</b>	<b>57.11</b>	<b>34.29</b>	<b>72.20</b>

Table 3. Ablation study on each component of FSKD. *UC-GBL* is by default at scale  $S = 8$  and \* means no use of  $\mathcal{L}_{ms-mk}$ . *Aux* stands for adding auxiliary keypoints for training; *MS UC-GBL* means multi-scale UC-GBL involved  $S = \{8, 12, 16\}$ . The results on the Animal dataset are the average over five subproblems.

One shot, PCK@ $\tau = 0.1$	Animal	CUB	NABird
1: <i>Baseline</i>	22.11	66.12	39.14
2: <i>Baseline+UC-GBL*</i>	24.17	68.29	41.16
3: <i>Baseline+UC-GBL*+Aux.</i>	41.70	74.50	51.62
4: <i>Baseline+UC-GBL+Aux.</i>	42.60	76.25	54.27
5: <i>Baseline+UC-GBL<sup>(12)</sup>+Aux.</i>	42.65	76.90	54.17
6: <i>Baseline+UC-GBL<sup>(16)</sup>+Aux.</i>	42.61	75.86	54.15
7: <i>Baseline+MS UC-GBL+Aux.</i>	<b>44.75</b>	<b>77.89</b>	<b>56.04</b>

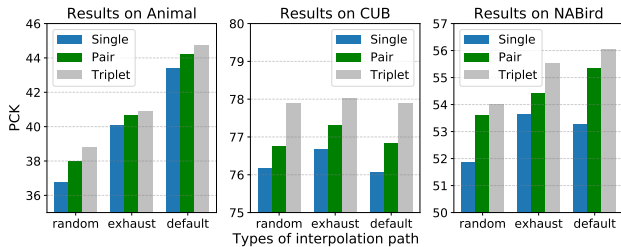


Figure 8. The impact of keypoint grouping strategies for various interpolation path strategies, where *single* means no grouping is used, and *exhaust* means using the exhaustive path in interpolation.

keypoints for unseen species across three datasets, reaching 38.77%, 77.90%, 54.01% for FSKD (rand) and 44.75%, 77.89%, 56.04% for FSKD (default). Though FSKD (rand) uses noisy auxiliary keypoints, it performs very well in CUB and NABird and even surpasses FSKD (default) in CUB; 3) Results on base keypoints are higher than on novel keypoints due to smaller domain shift. Similar trend is observed for the seen species. The qualitative results by the FSKD (default) are shown in Fig. 6. The localization uncertainty is represented by an ellipse whose major/minor axes are  $(3\sqrt{\lambda_1}, 3\sqrt{\lambda_2})$ , where  $\lambda_i$  are eigenvalues of  $\Sigma$ , and eigenvectors determine the orientation. Our FSKD localizes keypoints well and estimates uncertainty which matches GT and correlates with the shape of body parts.

**FSKD on DeepFashion2 & AwA:** Table 2 shows results of FSKD (default variant, 1-shot novel keypoint detection, unseen classes) w.r.t. 3 backbones/5 datasets, including *DeepFashion2* [17] (training on up-clothing categories/testing on lower-clothing categories) and the diverse *AwA Pose* [1] (novel keypoints types as in Animal test set, rest for train-

ing). Table 2 shows FSKD+HRNet-W48 yields 72.20% on AwA Pose ( $\sim 23\%$  over 48.81% of Animal dataset).

### 4.3. Ablation Study

Below we validate the effectiveness of each component using FSKD (default) under the novel keypoints detection.

**Number of Shots:** Typically, few-shot learning scores increase as the number of shots increases. Fig. 7(a) and (b) show that PCK scores at 5-shot yield improvements of 10.42% (average over five subproblems on Animal), 7.50% in NABird, and 2.23% in CUB, compared to 1-shot.

**Number of Training Keypoints:** We vary the number of training keypoints from base keypoint set and test the novel keypoints in CUB and NABird. Fig. 7(c) suggests that including more keypoints into training increases the keypoint diversity and helps FSKD generalize in novel keypoints.

**Localization and Semantic Uncertainty:** The statistical trend between the localization uncertainty and the distance error is shown in Fig. 7(d). We use  $J = 3(\sqrt{\lambda_1} + \sqrt{\lambda_2})$  to depict the ‘uncertainty strength’ for a keypoint prediction, then normalize  $J$  and distance error  $d$  by bounding box as  $(J', d') = (J, d) / \max(w_{\text{bbx}}, h_{\text{bbx}})$ . For keypoints, we divide  $d'$  into intervals of size 0.05 and calculate the average of  $d'$  and average of corresponding  $J'$  for each range. The plot shows  $J'$  becomes larger as  $d'$  increases, which validates that the learned uncertainty indicates the quality of predictions and can be used to suppress the noise in the loss function. Similarly, Fig. 7(e) shows the relation between the distinctiveness  $w$  of keypoints and  $d'$ , that is,  $w$  is lower when  $d'$  is larger because a keypoint is harder to localize when it is less semantically distinctive. Including both localization and semantic uncertainty into our UC-GBL, the 2nd row of Table 3 shows up to 3% gain across three datasets.

**Analysis of Auxiliary Keypoints:** Generating auxiliary keypoints (akin to self-supervision) boosts the visual diversity of training keypoints. The resulting noise is handled by our FSKD due to uncertainty modeling. After adding auxiliary keypoints, the scores improve dramatically (3rd row, Table 3). Moreover, when grouping auxiliary and main keypoints as triplets and modeling uncertainty for triplets, the performance improves further (4th row, Table 3). In addition, we also study the effects of different keypoint grouping strategies under various types of interpolation path. In Fig. 8, we can see up to 3% gain when using pairs or triplets.

**Improvements on Multi-scale UC-GBL:** Table 3 shows that multi-scale UC-GBL outperforms single-scale models. Therefore, multi-scale learning limits mislocalization.

## 5. Downstream Tasks

### 5.1. Few-shot Fine-grained Visual Recognition

Following [46], we adopt pose normalization (PN) that uses the concatenation of body part features to capture the distinctive features across fine-grained classes (Fig. 3(a)).

Table 4. Few-shot FGVR results. *AuxKps* means adding auxiliary keypoints to form augmented prototypes for classification.

Datasets	Models		1-shot	5-shot	all-shot
CUB	Proto	[42]	23.03	38.05	41.79
	Proto+BP	[31]	18.43	33.63	38.34
	Proto+bbN	[46]	23.97	40.22	44.61
	Proto+PN	[46]	35.92	58.66	63.51
	Ours		<b>37.45</b>	<b>61.22</b>	<b>66.25</b>
	Ours+ <i>AuxKps</i>		<b>38.04</b>	<b>61.74</b>	<b>66.37</b>
NABird	Proto+PN	[46]	26.17	50.55	60.03
	Ours		<b>27.68</b>	<b>51.81</b>	<b>61.56</b>

First, we modify our FSKD into a simple keypoint extractor by leveraging the universal keypoint prototype (UKP) computed by averaging the SKPs on additional 1000 episodes after training. In testing, UKPs guide FSKD to detect the keypoints, and thus FSKD *no longer* needs the support input as a reference. We use FGVR model from [46], based on ProtoNet (Proto) [42], whereas our FSKD supplies keypoint predictions. Bilinear pooling (BP) [31] and bounding box normalization (bbN) [46] based methods are also compared. All models are evaluated in all-way setting, 1-, 5- and all-shot results are reported. Table 4 shows our model achieves the best scores which validates the quality of FSKD. Including the features from auxiliary keypoints further improves results by making the prototypes more discriminative.

## 5.2. Semantic Alignment

Semantic alignment (SA) is used in tasks such as recognition and graphics [4, 40]. We demonstrate that, under the imperfect keypoint predictions, the query image  $\mathbf{I}_q$  can be warped into the rectified image  $\mathbf{I}'_q$  which aligns well with the support image  $\mathbf{I}_s$  (Fig. 3(b)).

**Uncertainty-weighted Thin-plate-spline Warp:** Unlike classic TPS warp [3, 12], the key idea of our uncertainty-weighted TPS warp is letting the unequal warping contributions of keypoints based on their uncertainty, as the well-matching correspondences should be encouraged to warp while uncertain ones not. Let  $\mathbf{P} = [\mathbf{p}_1, \dots, \mathbf{p}_N] \in \mathbb{R}^{2 \times N}$  be the support keypoints,  $\mathbf{P}' = [\mathbf{p}'_1, \dots, \mathbf{p}'_N] \in \mathbb{R}^{2 \times N}$  be the predicted query keypoints, and  $\hat{\mathbf{P}} = [\mathbf{1}, \mathbf{P}'^\top]^\top \in \mathbb{R}^{3 \times N}$ . Let the estimated uncertainty strength for each query keypoint be  $J_i$ , and  $\mathbf{D} = \text{diag}([J_1, \dots, J_N]) \in \mathbb{S}_{++}^N$  be a diagonal matrix. Then we obtain the transformation parameters<sup>1</sup>  $\mathbf{T} \in \mathbb{R}^{2 \times (N+3)}$  of uncertainty-weighted TPS warp as

$$\mathbf{T} = \left( \begin{bmatrix} \mathbf{R} + \lambda \mathbf{D}^2 & \hat{\mathbf{P}}^\top \\ \hat{\mathbf{P}} & \mathbf{0}_{3 \times 3} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{P}'^\top \\ \mathbf{0}_{3 \times 2} \end{bmatrix} \right)^\top, \quad (11)$$

Matrix  $\mathbf{R} \in \mathbb{R}^{N \times N}$  has entries  $r_{i,j} = d_{i,j}^2 \log d_{i,j}^2$ ,  $d_{i,j}$  is the Euclidean distance between  $\mathbf{p}_i$  and  $\mathbf{p}_j$ , and  $\lambda \geq 0$  is the

<sup>1</sup>Uncertainty-weighted TPS warp is derived in §A of Suppl. Material.

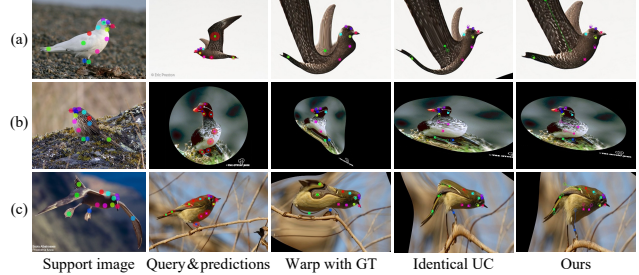


Figure 9. Comparison of semantic alignment approaches. The first column shows the support keypoints & image; the second column shows the query image with the predicted keypoints (marked by tilted crosses) and uncertainty (red shadow ellipses); the last three columns are the results achieved by *Warp with GT* [3], *Identical UC*, and our uncertainty-weighted TPS warp.

warping penalty. By using  $\mathbf{T}$ , every pixel grid  $\mathbf{q}'_i = [x_i, y_i]^\top$  in the rectified image  $\mathbf{I}'_q$  has the mapped pixel in the query image  $\mathbf{I}_q$  following the transformation  $\mathbf{q}_i = \mathbf{T} \tilde{\mathbf{q}}'_i$ , where  $\tilde{\mathbf{q}}'_i = [r_{1,i}, \dots, r_{N,i}, 1, x_i, y_i]^\top$  and  $r_{n,i} = d_{n,i}^2 \log d_{n,i}^2$ , the  $d_{n,i}$  is the Euclidean distance between the  $n$ -th support keypoint  $\mathbf{p}_n$  and the pixel grid  $\mathbf{q}'_i$ . After image remapping  $\mathbf{I}'_q[\mathbf{q}'_i] = \mathbf{I}_q[\mathbf{q}_i]$ , we obtain the rectified image  $\mathbf{I}'_q$ .

**Results:** We perform SA for unseen species using 1-shot FSKD model trained on mix-species episodes. We set  $\lambda = 1$  and compare our approach with 1) *Warp with GT* [3], which uses GT query keypoints; and 2) *Identical UC* using predicted keypoints with identical uncertainty  $\mathbf{D} = \text{diag}([s, \dots, s])$  where  $s = 20^2 \log 20^2$  was chosen experimentally. Fig. 9 shows that our approach penalizes the warp of uncertain keypoints, reduces the risks of unacceptable deformations, and produces a good alignment with support image. Fig. 9(a) shows the detected wing (GT is right-wing) with large uncertainty due to the occlusion, which results in a small warp and a large distance difference w.r.t. the corresponding support keypoint (Fig. 9(a), 5th column).

## 6. Conclusion

We have extended few-shot learning into the challenging task of keypoint detection by introducing a novel FSKD approach which learns the localization uncertainty of keypoints. FSKD is very flexible as it can detect keypoints of various types (seen vs. unseen) on various species (seen vs. unseen). Our simple uncertainty model deals with the keypoint noise and elegantly produces the uncertainty distribution of the likely position of GT keypoints. With the help of auxiliary keypoints, multi-keypoint covariance, and multi-scale localization, FSKD significantly boosts the detection performance. Moreover, FSKD can be successfully applied to a variety of downstream tasks such as FGVR and semantic alignment, where our novel uncertainty-weighted TPS warp leverages uncertainty. We hope our FSKD model will provide the starting point for the vision community and inspire the future research on few-shot keypoint detection.



## References

- [1] Prianka Banik, Lin Li, and Xishuang Dong. A novel dataset for keypoint detection of quadruped animals from images. *arXiv preprint arXiv:2108.13958*, 2021. 7
- [2] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pages 1613–1622. PMLR, 2015. 2
- [3] Fred 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. 8, 12, 16
- [4] Steve Branson, Grant Van Horn, Serge Belongie, and Pietro Perona. Bird species categorization using pose normalized deep convolutional nets. *arXiv preprint arXiv:1406.2952*, 2014. 8
- [5] Jinkun Cao, Hongyang Tang, Hao-Shu Fang, Xiaoyong Shen, Cewu Lu, and Yu-Wing Tai. Cross-domain adaptation for animal pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9498–9507, 2019. 2, 5
- [6] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Openpose: realtime multi-person 2d pose estimation using part affinity fields. *IEEE transactions on pattern analysis and machine intelligence*, 43(1):172–186, 2019. 1, 2
- [7] Joao Carreira, Pulkit Agrawal, Katerina Fragkiadaki, and Jitendra Malik. Human pose estimation with iterative error feedback. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4733–4742, 2016. 2
- [8] Bowen Cheng, Bin Xiao, Jingdong Wang, Honghui Shi, Thomas S Huang, and Lei Zhang. Higherhrnet: Scale-aware representation learning for bottom-up human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5386–5395, 2020. 2
- [9] Jiwoong Choi, Dayoung Chun, Hyun Kim, and Hyuk-Jae Lee. Gaussian yolov3: An accurate and fast object detector using localization uncertainty for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 502–511, 2019. 3
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 13
- [11] Konstantinos G Derpanis. The harris corner detector. *York University*, pages 2–3, 2004. 2
- [12] Gianluca Donato and Serge J Belongie. *Approximation methods for thin plate spline mappings and principal warps*. Cite-seer, 2003. 8
- [13] Qi Fan, Wei Zhuo, Chi-Keung Tang, and Yu-Wing Tai. Few-shot object detection with attention-rpn and multi-relation detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4013–4022, 2020. 2
- [14] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. Rmpe: Regional multi-person pose estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 2334–2343, 2017. 1, 2
- [15] Li Fei-Fei, Rob Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE transactions on pattern analysis and machine intelligence*, 28(4):594–611, 2006. 2
- [16] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017. 2, 15
- [17] Yuying Ge, Ruimao Zhang, Xiaogang Wang, Xiaoou Tang, and Ping Luo. Deepfashion2: A versatile benchmark for detection, pose estimation, segmentation and re-identification of clothing images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5337–5345, 2019. 7
- [18] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 1, 2, 3
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5, 7
- [20] Yihui He, Chenchen Zhu, Jianren Wang, Marios Savvides, and Xiangyu Zhang. Bounding box regression with uncertainty for accurate object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2888–2897, 2019. 3
- [21] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. *Advances in neural information processing systems*, 28:2017–2025, 2015. 3
- [22] Bingyi Kang, Zhuang Liu, Xin Wang, Fisher Yu, Jiashi Feng, and Trevor Darrell. Few-shot object detection via feature reweighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8420–8429, 2019. 2
- [23] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *NIPS*, 2017. 2, 4
- [24] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018. 3
- [25] Gregory Koch, Richard Zemel, Ruslan Salakhutdinov, et al. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2. Lille, 2015. 1, 2
- [26] Piotr Koniusz and Hongguang Zhang. Power normalizations in fine-grained image, few-shot image and graph classification. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 1
- [27] Marek Kowalski, Jacek Naruniec, and Tomasz Trzcinski. Deep alignment network: A convolutional neural network for robust face alignment. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 88–97, 2017. 1

- [28] Gen Li, Varun Jampani, Laura Sevilla-Lara, Deqing Sun, Jonghyun Kim, and Joongkyu Kim. Adaptive prototype learning and allocation for few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8334–8343, 2021. [2](#)
- [29] Siyuan Li, Semih Gunel, Mirela Ostrek, Pavan Ramdya, Pascal Fua, and Helge Rhodin. Deformation-aware unpaired image translation for pose estimation on laboratory animals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13158–13168, 2020. [2](#)
- [30] Xiang Li, Tianhan Wei, Yau Pun Chen, Yu-Wing Tai, and Chi-Keung Tang. Fss-1000: A 1000-class dataset for few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2869–2878, 2020. [2](#)
- [31] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear cnn models for fine-grained visual recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 1449–1457, 2015. [8](#)
- [32] G Lowe. Sift-the scale invariant feature transform. *Int. J.* 2:91–110, 2004. [2](#)
- [33] Zhengxiong Luo, Zhicheng Wang, Yan Huang, Liang Wang, Tieniu Tan, and Erjin Zhou. Rethinking the heatmap regression for bottom-up human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13264–13273, 2021. [3](#)
- [34] Julian Marstaller, Frederic Tausch, and Simon Stock. Deepbees-building and scaling convolutional neuronal nets for fast and large-scale visual monitoring of bee hives. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. [1](#)
- [35] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016. [1, 2](#)
- [36] David Novotny, Samuel Albanie, Diane Larlus, and Andrea Vedaldi. Self-supervised learning of geometrically stable features through probabilistic introspection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3637–3645, 2018. [2, 3, 4, 5, 13](#)
- [37] 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. *Nature methods*, 16(1):117–125, 2019. [1](#)
- [38] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. [3](#)
- [39] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99, 2015. [3](#)
- [40] Baoguang Shi, Xinggang Wang, Pengyuan Lyu, Cong Yao, and Xiang Bai. Robust scene text recognition with automatic rectification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4168–4176, 2016. [8](#)
- [41] Christian Simon, Piotr Koniusz, and Mehrtash Harandi. Meta-learning for multi-label few-shot classification. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3951–3960, 2022. [1](#)
- [42] Jake Snell, Kevin Swersky, and Richard S Zemel. Prototypical networks for few-shot learning. *arXiv preprint arXiv:1703.05175*, 2017. [2, 3, 8, 13](#)
- [43] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5693–5703, 2019. [2, 7](#)
- [44] Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deep convolutional network cascade for facial point detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3476–3483, 2013. [1, 2](#)
- [45] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1199–1208, 2018. [1, 2, 13](#)
- [46] Luming Tang, Davis Wertheimer, and Bharath Hariharan. Revisiting pose-normalization for fine-grained few-shot recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14352–14361, 2020. [1, 7, 8](#)
- [47] Alexander Toshev and Christian Szegedy. Deeppose: Human pose estimation via deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1653–1660, 2014. [1, 2](#)
- [48] Grant Van Horn, Steve Branson, Ryan Farrell, Scott Haber, Jessie Barry, Panos Ipeirotis, Pietro Perona, and Serge Belongie. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 595–604, 2015. [5](#)
- [49] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29:3630–3638, 2016. [1, 2](#)
- [50] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. [5](#)
- [51] Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM Computing Surveys (CSUR)*, 53(3):1–34, 2020. [1, 2](#)
- [52] Zhe Wu, Li Su, and Qingming Huang. Stacked cross refinement network for edge-aware salient object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7264–7273, 2019. [5](#)
- [53] Yi Yang and Deva Ramanan. Articulated human detection with flexible mixtures of parts. *IEEE transactions on pattern analysis and machine intelligence*, 35(12):2878–2890, 2012. [5](#)
- [54] Hongguang Zhang, Jing Zhang, and Piotr Koniusz. Few-shot learning via saliency-guided hallucination of samples.

- In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2770–2779, 2019. [2](#)
- [55] Shan Zhang, Lei Wang, Naila Murray, and Piotr Koniusz. Kernelized few-shot object detection with efficient integral aggregation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022. [2](#)
- [56] 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. [1](#)
- [57] Weilin Zhang and Yu-Xiong Wang. Hallucination improves few-shot object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13008–13017, 2021. [2](#)
- [58] Hao Zhu and Piotr Koniusz. Ease: Unsupervised discriminant subspace learning for transductive few-shot learnin. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022. [1](#)