

A Mathematical Analysis of Learning Loss for Active Learning in Regression

Megh Shukla ✉

Shuaib Ahmed

Mercedes-Benz Research and Development India

megh.shukla@daimler.com

Abstract

Active learning continues to remain significant in the industry since it is data efficient. Not only is it cost effective on a constrained budget, continuous refinement of the model allows for early detection and resolution of failure scenarios during the model development stage. Identifying and fixing failures with the model is crucial as industrial applications demand that the underlying model performs accurately in all foreseeable use cases. One popular state-of-the-art technique that specializes in continuously refining the model via failure identification is Learning Loss[24]. Although simple and elegant, this approach is empirically motivated. Our paper develops a foundation for Learning Loss which enables us to propose a novel modification we call LearningLoss++. We show that gradients are crucial in interpreting how Learning Loss works, with rigorous analysis and comparison of the gradients between Learning Loss and LearningLoss++. We also propose a convolutional architecture that combines features at different scales to predict the loss. We validate LearningLoss++ for regression on the task of human pose estimation (using MPII and LSP datasets), as done in Learning Loss. We show that LearningLoss++ outperforms in identifying scenarios where the model is likely to perform poorly, which on model refinement translates into reliable performance in the open world.

1. Introduction

The deep learning era has heralded a paradigm shift in the approaches we take to model our data. Cheaper hardware, clean data and robust modelling has resulted in a large number of industrial applications such as driver monitoring systems and activity recognition. However, like all data-driven paradigms, deep learning has an insatiable appetite for data. Industry applications of human pose estimation often require millions of labelled images for acceptable performance to eliminate risks associated with faulty predictions when the model is deployed on-ground. This introduces a two-fold challenge: finding algorithms to recognize

failure cases for the model, as well as reducing the costs associated with the large scale collection of annotated data.

One potential solution lies in active learning, a class of algorithms aimed at reducing annotation costs involved in training the model. The essence behind all active learning algorithms is to allow the model to choose a set of images which would impart maximum information about the dataset if labelled. Although active learning is well explored, active learning in human pose estimation remains a challenge. This is because human pose models [15, 20, 21] regress two dimensional heatmaps to estimate the location of joints, making this approach different from those tasks having a probabilistic interpretation of outputs. The task of regressing heatmaps is sufficiently different from regressing the joint coordinates [3]; with DeepPose (2014) [22] being the last significant technique to do so in human pose estimation.

This paper builds upon Learning Loss For Active Learning [24], a task agnostic active learning algorithm having performance comparable with state-of-the-art in classification, object detection and human pose estimation. This is remarkable, as we do not specify domain constraints such as geometry of the human skeletal structure explicitly. However, this approach has two drawbacks: 1) The learning loss objective is intuitively defined and 2) The architecture of the learning loss module discards spatial information. With LearningLoss++, our contributions include:

1. We establish equivalency between Learning Loss' empirically driven objective and the KL divergence objective proposed by LearningLoss++ which is hyperparameter free.
2. We analyze the gradient to provide intuition into the training process of Learning Loss. We then compare the expected gradients for Learning Loss and LearningLoss++, which allows us to prove the advantages associated with smoothed gradients.
3. A convolutional architecture to combine features at different scales, replacing the global average pooling - fully connected architecture defined in Learning Loss.




	(a)	(b)	(c)
			
True Loss: l	0.45	0.47	2.87
Pred Loss: \hat{l}	1.48	1.51	1.63
Image Pairs(i, j)	$(i = (a), j = (b))$		$(i = (a), j = (c))$
$\ \nabla_w(l_i, l_j, \hat{l}_i, \hat{l}_j)\ $: (LearningLoss)	$\ \theta_i - \theta_j\ $		$\ \theta_i - \theta_j\ $
$\ \nabla_w(l_i, l_j, \hat{l}_i, \hat{l}_j)\ $: (LearningLoss++)	≈ 0		$0.32\ \theta_i - \theta_j\ $

Figure 1. Both Learning Loss and LearningLoss++ (ours) use ranking loss to train the loss predictor network. However, the gradient response for both the approaches significantly differ. For the image pair (a, b), the true loss as well as predicted loss are similar, agreeing with intuition. While LearningLoss++ concurs with this notion, Learning Loss wrongly penalizes the network. For the image pair (a,c), the true loss for (c) is significantly larger than (a) whereas the predicted loss is similar. In this scenario, LearningLoss++ enforces a larger gradient whereas the Learning Loss gradient has the same response as in the previous scenario, in spite of the wrong predictions.

We perform experimental validation using the MPII [1] and LSP-LSPET [7] datasets and compare our approaches with other methods, with an emphasis on the ability of LearningLoss++ to identify failures. Successful identification of failure cases translates into continuous refinement of the model and thus improved reliability in the open world use cases.

2. Related Work

The early foundation for active learning is summarized in [18], with approaches involving uncertainty studied in [8, 9, 11, 13]. Query Based Committee [19] and other ensemble approaches [2, 14] explore the performance of a group of models with different parameter spaces for an image and compute the level of agreement between the models for the image. Expected Gradient Length based approaches [25], [26] use gradient as a measure for determining the *informativeness* of an image. Diversity promoting approaches such as [23], [4] and [17] use a small sample set to represent the larger pool of unlabelled data for annotation. Bayesian active learning approaches [5], [10] provide beautiful theoretical results using aleatoric and epistemic uncertainties, but to the best of our knowledge, there are no available results for human pose estimation.

Bayesian uncertainty has been used in hand pose estimation (DeepPrior) [3] but a direct application is not possible for human pose estimation. DeepPrior directly regresses the joint coordinates, unlike human pose estimation which regresses entire heatmaps. Estimating the epistemic uncertainty for entire heatmaps is not studied, limiting its application to human pose estimation. Since DeepPrior consists of

fully connected layers, applying dropout is a standard technique. However, the use of dropout is absent in fully convolutional architectures such as those used in human pose estimation. Finally, the extension of DeepPrior when joints are occluded is not clear. Aleatoric uncertainty has been used in human pose estimation [6] to estimate the location of occluded joints, however this approach does not extend to active learning.

Extending these active learning approaches to human pose estimation is not trivial. Traditional uncertainty based approaches find applications in classification due to the availability of a posterior distribution, which is unavailable for human pose estimation. Ensemble approaches are memory intensive for deployment on edge devices. Methods relying on gradient length are computationally expensive which limit their use in real-time applications. Diversity techniques lack the ability to detect use cases where the model is likely to fail. Bayesian uncertainty approaches not only restrict the network to a Bayesian Convolutional Network (BCN), they rely on the use of Dropout and multiple forward propagation runs reducing their use in real-time applications. The first algorithms dedicated to active learning for human pose estimation was [12] which involved the use of multi-peak entropy. It can be shown that multi-peak entropy value is essentially a lower bound for the standard entropy approach. Learning loss for active learning [24] proposed the use of a general purpose auxiliary model trained to predict an indicative loss for tasks involving classification, detection and pose estimation. The underlying philosophy for learning loss is to detect images where the model performs poorly. Our work builds upon this paper to im-

prove the correlation between predicted and true loss which enables better detection of failures therefore improving the model early on during the development phase.

3. Learning Loss (Yoo and Kweon) [24]

The learning loss module is an auxiliary network $h(\cdot)$ attached to the intermediate layers of the main model (classifier/pose estimator) $f(\cdot)$. For a given image x , we have: $\hat{y} = f(x)$ where \hat{y} is the prediction of the model f with the ground truth y . The value of true loss l is $l = L_{model}(y, \hat{y})$ where $L_{model}(\cdot)$ is the loss function such as cross-entropy / mean square error. Intermediate representations of the model $f^h(x)$ are inputs for the learning loss network to predict a 'loss' \hat{l} , where $\hat{l} = h(f^h(x))$.

3.1. Method

The notion of predicted loss \hat{l} is similar to that of the true loss l , with a high value of \hat{l} implying that the model has likely produced a wrong prediction. In the absence of true loss as is the case with unlabelled images, the predicted loss steps in to compute the performance of the model on the input image. Therefore, images with a high \hat{l} are selected for annotation to refine the use cases where the model is predicted to perform poorly. To train the learning loss network, the authors use a ranking loss to compare a pair of images. Let (x_i, l_i, \hat{l}_i) , (x_j, l_j, \hat{l}_j) represent a pair of (image, true loss, predicted loss). The objective to train the learning loss network is:

$$\mathbb{L}_{loss} = \max \left(0, -\text{sign}(l_i - l_j)(\hat{l}_i - \hat{l}_j) + \xi \right) \quad (1)$$

The idea behind Eq: 1 is simple, if the true loss l_i corresponding to x_i is greater than l_j , then the predicted loss \hat{l}_i has to be greater than \hat{l}_j by a margin ξ so as to not incur any loss. Similarly, if $l_j > l_i$ then $\hat{l}_j > \hat{l}_i$. The learning loss network incurs a loss if the predicted loss is not greater than ξ or if the learning loss network predicts the opposite ($l_i > l_j$ but $\hat{l}_i < \hat{l}_j$). The decision to use a pairwise comparison of images might seem strange, especially when mean square error $(l - \hat{l})^2$ can be used to train the learning loss network to learn the mapping between an input image x and the true loss l . The authors on the contrary argue that the network trained on MSE fails to learn anything meaningful in their experimental study.

Active learning using learning loss is straightforward. We train the task specific model $f(\cdot)$ and the learning loss network $h(\cdot)$ using all labelled images ($x \in \mathcal{L}$). We then select a subset of unlabelled images ($x \in \mathcal{U}$) for annotation which have a high predicted loss $\hat{l} = h(f^h(x))$. This process continues cyclically for continuous refinement of the model.




Figure 2. We use an example to demonstrate that Learning Loss forces the weight vector w to align along the most discriminative component $\theta_i - \theta_j$ to explain away the predicted losses

4. LearningLoss++

We first establish equivalency between the intuition driven objective (Eq: 1) and the KL divergence based objective of LearningLoss++ by comparing the gradients. We then analyse the gradient formulation to provide insights into the training of the learning loss network. This is followed by highlighting some shortcomings associated with the Learning Loss gradient, and finally we analyse the LearningLoss++ gradient to show that it implicitly absorbs the margin ξ hyperparameter.

4.1. Gradient Analysis

Let θ represent the output of the *penultimate* layer in the fully connected learning loss network. By definition, the predicted loss in terms of θ is: $\hat{l} = \theta^T w$ where w represents the weights of the final learning loss layer. While we consider the gradients associated with the penultimate layer, we note that these gradients are backpropagated and hence our analysis extends to any general fully connected network.

Gradient for Learning Loss: The learning loss objective Eq: 1 can now be written as: $\mathbb{L}_{loss} = \max(0, -\text{sign}(l_i - l_j)(\theta_i^T w - \theta_j^T w) + \xi)$. The gradient for learning loss is:

$$\nabla_w \mathbb{L}_{loss} \in \{0, \pm(\theta_i - \theta_j)\} \quad (2a)$$

$$\nabla_{\theta} \mathbb{L}_{loss} \in \{0, \pm w\} \quad (2b)$$

Eq: 2 provides an interesting insight into the behaviour of the module. *The weights Eq: 2(a) are aligned to emphasize on the most discriminative component between the intermediate features.* If we treat θ as a vector in $n - \text{dimensional}$ space and θ_i and θ_j differ along one component only, the weight vector amplifies the response along that particular component to predict the indicative loss, as we discuss below.

So How Does Learning Loss work? We highlight the discriminative property of the gradients with an example. For this example, let us assume $\theta_i = [1, 2, 1]^T$,

$\theta_j = [1, 1, 1]^T$ and $w = [0, 3, 0]$ in a three dimensional vector space as shown in Fig: 2. We use the same value of margin ξ as recommended in the paper. Depending upon the true loss (l_i and l_j), the objective leads us to two cases: 1) $l_i > l_j$ and 2) $l_i < l_j$. Substituting the values of θ_i, θ_j and w in \mathbb{L}_{loss} (Eq: 1) for case 1, we get: $\mathbb{L}_{loss}(w, \theta_i, \theta_j) = \max(0, -1 * (6 - 3) + 1) = 0$.

Intuitively, this corresponds to taking a projection of both the penultimate layer outputs θ along the weight vector w . Since the current weight vector is already aligned along the discriminative component, the learning loss model does not incur a penalty. Things get interesting if we consider the second case, $l_i < l_j$. In this scenario, the objective incurs a penalty ($\mathbb{L}_{loss} = \max(0, 1 * (6 - 3) + 1) = 4$). Using Eq: 2(a), we get $\nabla_w \mathbb{L}_{loss} = [0, 1, 0]^T$, with gradient descent acting against this direction ($-\nabla_w \mathbb{L}_{loss}$). This step enables the weight vector to reverse direction and ultimately point along the direction that minimizes the loss.

Gradient for LearningLoss++: As we saw in the previous example, the discriminative property ($\theta_i - \theta_j$) of the gradient allows the learning loss network to align weights in a manner to explain the predicted losses. We show that the proposed KL divergence based training objective has a gradient with a similar form.

As before, let $\hat{l}_i = \theta_i^T w$ and $\hat{l}_j = \theta_j^T w$ represent the predicted loss for images i and j . We now consider two images in the minibatch and compute softmax over \hat{l} : $q_i = e^{\hat{l}_i} / (e^{\hat{l}_i} + e^{\hat{l}_j})$ with q_j defined similarly. This interpretation can be viewed as the probability of sampling x_i over x_j and vice versa for annotation. While we have defined a probabilistic interpretation for the predicted losses, a similar one needs to be defined for the true losses. Since true losses ($l_i, l_j > 0$), we use simple scaling $p_i = \frac{l_i}{l_i + l_j}$ to denote the probability of x_i having a higher true loss than x_j . Intuitively, an image x_i having a true loss n -times greater than x_2 is n -times more likely to get sampled for annotation. The objective to minimize is:

$$\mathbb{L}_{loss}(w, \theta_i, \theta_j) = \text{KL}(p||q) = p_i \log \frac{p_i}{q_i} + p_j \log \frac{p_j}{q_j} \quad (3)$$

For brevity, we reproduce the final solution, referring the reader to the supplementary material for the full derivation. The solution is simple and delightfully familiar:

$$\nabla_w \mathbb{L}_{loss}(w, \theta_i, \theta_j) = (q_i - p_i)(\theta_i - \theta_j) \quad (4a)$$

$$\nabla_{\theta} \mathbb{L}_{loss}(w, \theta_i, \theta_j) = (q_i - p_i)w \quad (4b)$$

Active Learning with LearningLoss++: Although we have introduced softmax and KL divergence with LearningLoss++, the process of active learning sampling remains the same as in Learning Loss. The images corresponding to the top- k predicted losses \hat{l} are chosen for annotation. We




Figure 3. Statistical results show that the training losses (true loss) follow a gamma distribution for a sufficiently converged regression model. This figure shows the true loss histogram for a human pose estimation model. We can approximate this particular histogram with $\gamma(k = 4, \Theta = 0.066)$ obtained using maximum likelihood estimation.

use softmax and KL divergence to train the loss prediction network only. Therefore, while the gradient computation is probabilistic, the core active learning sampling process is deterministic.

4.1.1 LearningLoss++ Advantages

We have successfully proven that Learning Loss and LearningLoss++ share the same underlying principles in training. The most prominent feature of the LearningLoss++ gradient is the presence of $(q_i - p_i)$ that smoothens the gradient in comparison to Learning Loss. This smoothing, as we will show, allows for better detection of *lossy* images. We additionally note that LearningLoss++ implicitly absorbs the margin hyperparameter ξ . **Fig: 1 depicts the two scenarios which highlight the advantages associated with the LearningLoss++ gradient.** We discuss the two scenarios in detail:

Scenario 1 (Fig: 1, image pair-(a, b)): In the event we sample a pair of images with similar true losses, Learning Loss incurs a loss even when the predicted losses are similar, since the margin ξ forces the predicted losses to have a prediction loss margin of at least ξ . With LearningLoss++, the network does not incur a penalty. Since $l_i \approx l_j \implies p_i \approx p_j \approx 0.5$ and $\hat{l}_i \approx \hat{l}_j \implies q_i \approx q_j \approx 0.5$, we have $(q_i - p_i) \approx 0$. Although one question remains; how likely are we to sample a pair of images having similar true losses? Or formally, if $P(X = l_i, Y = l_j)$ represents the probability of sampling images with losses (l_i, l_j) , then given a true loss margin δ (not to be confused with ξ), what is $P(|X - Y| \leq \delta)$?

Fortunately, a closed form solution exists to compute the

δ	0.02	0.04	0.06	0.08	0.1	0.125	0.15
$P_{X,Y,\gamma}$	0.094	0.185	0.274	0.358	0.437	0.527	0.607

Table 1. We compute $P(|(X = l_i) - (Y = l_j)| \leq \delta)$ using the closed form solution (Eq:5-supplementary) for true loss $-l_i, l_j$ distributed according to $\gamma(k = 4, \theta = 0.066)$. We have verified the correctness of our solution with a computer simulation.

probability of sampling a pair of images with similar losses. Statistical regression under the condition of homoscedasticity assumes that the residuals (losses) are distributed as per the normal distribution $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. The distribution of squared residual ε^2 follows $\mathcal{N}^2(0, \sigma^2) = \gamma(\frac{1}{2}, 2\sigma^2)$. This result can be generalized to a summation of n gaussian distributions: $p(\varepsilon^2) = \sum_{i=1}^n \mathcal{N}_i(0, \sigma^2) = \gamma(\frac{n}{2}, 2\sigma^2)$. Therefore, the probability of sampling a loss follows a gamma distribution. : $P(X = l_i) = \gamma(k, \Theta)$ as shown in Fig: 3.

The probability of sampling two images with true losses within δ is:

$$P(|X - Y| \leq \delta) = \int_0^\delta \gamma(x, k, \Theta) \int_0^{x+\delta} \gamma(y, k, \Theta) dy dx + \int_\delta^\infty \gamma(x, k, \Theta) \int_{x-\delta}^{x+\delta} \gamma(y, k, \Theta) dy dx \quad (5)$$

To compute a closed form solution for Eq: 5, we restrict $k \in \mathbb{Z}^+$ allowing us to compute the integral of the gamma distribution analytically. Since the rest of the derivation involves extensive simplification, we request the reader to refer to the final result Eq: (5) from the supplementary material. To provide intuition into Eq:5-supplementary, we use the loss distribution from Fig: 3 as an example to compute the $P(|X - Y| \leq \delta)$ for various true loss margin δ in Table: 1. We observe that there is a high probability (10% - 43%) of sampling a pair of images having similar true losses. For a well trained loss prediction module that correctly predicts similar \hat{l} for similar l , this results in upto 10% - 43% of the resultant gradient updates being noisy. Fortunately, LearningLoss++ does not suffer from this issue.

Scenario 2 (Fig: 1, image pair-(a, c)): We also consider the case where the true losses (l_i, l_j) for a pair of images are significantly different and their predicted losses (\hat{l}_i, \hat{l}_j) are similar. In this scenario, Learning Loss correctly incurs a gradient, however the gradient formulation is the same as when the true and predicted losses were similar. *LearningLoss++ gives greater weightage to the gradient when the predicted losses for a pair of images do not reflect the fact that one of the images sampled in the pair has a high true loss than the other.* To prove this statement, we compute the expected gradient (Eq: 4) for LearningLoss++ for a fixed true loss margin δ . For a true loss pair $(X = l_i, Y = l_j)$ sampled from $\gamma(k, \Theta)$ where $Y = X + \delta$

$\delta \rightarrow$	0.0	0.1	0.2	0.3	0.4	0.5
LL++	$q_i-0.5$	$q_i-0.39$	$q_i-0.3$	$q_i-0.25$	$q_i-0.21$	$q_i-0.18$
LL	\leftarrow constant c_1 \rightarrow					

Table 2. We show the expected gradient response $K(\theta_i - \theta_j)$ where K is tabulated above for different values of the true loss margin δ using Eq:12-supplementary. These values are computed assuming that the true loss values l are distributed according to $\gamma(k = 4, \Theta = 0.066)$ (Fig: 3). LL++: LearningLoss++, LL: Learning Loss

(results hold true when $Y = X - \delta$ by symmetry) and $p_i = l_i/(l_i + l_j)$, the expected gradient with respect to $\gamma_{k,\theta}$ is:

$$\mathbb{E}_{x,y|\delta_2} \left[\nabla_w \mathbb{L}(w, \theta_i, \theta_j) \right] = \lim_{\delta_1 \rightarrow \delta_2} \int_{x=0}^{x=\infty} \int_{y=x+\delta_1}^{y=x+\delta_2} (q_i - \frac{x}{2x + \delta_2})(\theta_i - \theta_j) \frac{\gamma(x, k, \Theta) \gamma(y, k, \Theta)}{p(y - x = \delta_2)} dy dx \quad (6)$$

The continuous nature of the gamma distribution leads us to use $\delta_1 \rightarrow \delta_2$ to faithfully infer area under the curve as probability. The simplification of Eq: 6 is along similar lines as deriving $P(|X - Y| \leq \delta)$, hence we refer the curious reader to the supplementary material for a complete derivation. We note that the final solution (Eq: 12-supplementary) can be written as a function of true loss margin δ : $\mathbb{E}_{x,y|\delta} [\nabla_w \mathcal{L}] = (q_i - \phi(\delta))(\theta_i - \theta_j)$. Since Eq: 12-supplementary is verbose and not very intuitive, we turn to Table: 2 for a more intuitive outlook. When $\delta = 0$, the loss prediction network incurs no loss if it predicts $\hat{l}_i \approx \hat{l}_j \implies q_i \approx q_j \approx 0.5$. However, a response of $q_i \approx q_j \approx 0.5$ when the true loss margin $\delta = 0.5$ incurs a large larger gradient response: $(0.5 - 0.18)(\theta_i - \theta_j)$. Only when the loss predictor network predicts $\hat{l}_i < \hat{l}_j \implies q_i < q_j$, the network incurs a lower gradient response. Therefore, for a pair of images where one image has a higher true loss, the loss predictor network is forced to predict a high value of predicted loss \hat{l} for that particular image, in the absence of which the gradient penalty is steep. This translates into the network better identifying images with a high true loss (or faulty inferences). We also note that the softmax ($q = \text{softmax}(\hat{l}_i, \hat{l}_j)$) associated with LearningLoss++ removes the need for a predicted loss margin ξ hyperparameter.

So far, we have rigorously analyzed learning loss for regression. We have successfully shown that the objectives of Learning Loss and LearningLoss++ are equivalent. We have highlighted the role played by the gradient in aligning the weights in a manner that explains the predicted loss. We have shown that a non-trivial number of Learning Loss gradient updates are noisy. We later derived the expected gradient formulation for LearningLoss++, which ensures that the images with a high true loss are identified correctly. Our

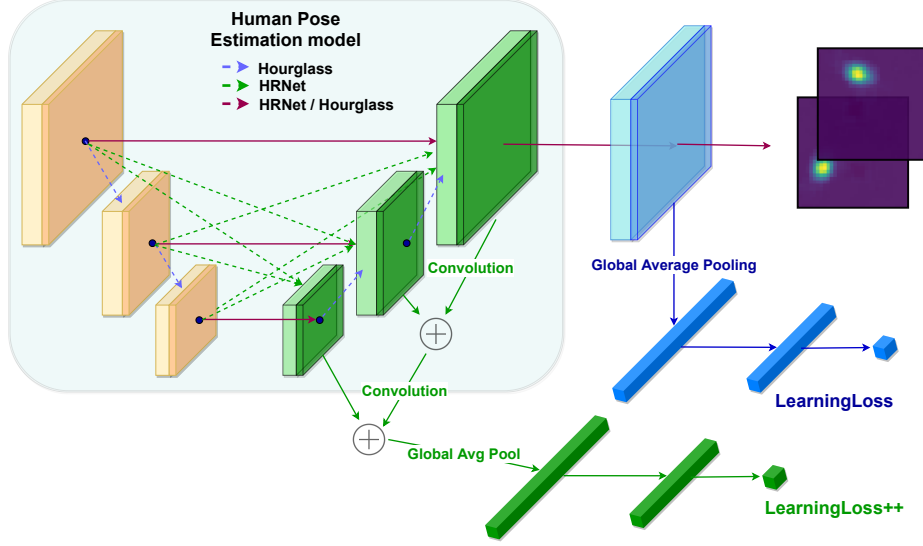


Figure 4. Learning Loss uses global average pooling to collapse the spatial dimensions of the intermediate feature maps into a vector. We argue that such an approach helps only in trivial cases: when the losses can be inferred from cues such as image backgrounds. Human pose estimation relies on spatial interaction between various features, which is lost with pooling. A convolutional feature extractor for learning loss hence captures features at multiple scales without spatial loss, allowing the learning loss network to perform well even when analyzing complex poses.

discussion till now on Learning Loss and LearningLoss++ was valid for all regression tasks. In the next section, we propose a convolutional architecture for the loss prediction module and discuss its associated advantages specific to human pose estimation.

4.2. Convolutional Architecture: Human Pose Estimation

Instead of using global average pooling (GAP) (Fig: 4), we suggest the use of a convolutional feature extractor for LearningLoss++, since GAP removes any spatial dependency between the features which the learning loss model can use to predict the indicative loss. The convolutional network takes as input intermediate features at multiple scales from the human pose estimation model. These multi-resolution intermediate representations are a vital feature of state-of-the-art models such as HRNet and Stacked Hourglass that allow the network to have a microscopic as well as a macroscopic view of the images.

We represent the multiscale representations of spatial dimensions $n \times n$ as $\mathcal{H}_{n \times n}$. The stacked hourglass has features at five such scales in an hourglass: $\mathcal{H}_{4 \times 4}, \mathcal{H}_{8 \times 8} \dots \mathcal{H}_{64 \times 64}$. To combine features at multiple scales, we suggest the use of strided convolutions. Let $\mathcal{H}_{n_1 \times n_1}, \mathcal{H}_{n_2 \times n_2}$ represent features at consecutive scales where $n_1 > n_2$. Then, the combination of the two features is $\mathcal{H}_{n_2 \times n_2} += \text{Conv}(\text{stride} = \frac{n_1}{n_2}, \text{kernel_size} = \frac{n_1}{n_2})(\mathcal{H}_{n_1 \times n_1})$. If n_1 is not perfectly divisible by n_2 , the convolution takes place with stride, kernel = $\lfloor \frac{n_1}{n_2} \rfloor$ followed

by another convolution with stride=1 and an appropriate kernel size chosen so that the output of this convolution matches $\mathcal{H}_{n_2 \times n_2}$. We perform a superimposition $\mathcal{H}_{n_2 \times n_2} + (\mathcal{H}_{n_1 \times n_1} \xrightarrow{\text{Conv}} \mathcal{H}_{n_2 \times n_2})$ as concatenation would drastically increase the size of the network. The reduction process from $\mathcal{H}_{n_1 \times n_1} \rightarrow \mathcal{H}_{n_2 \times n_2} \dots$ is carried out till we reach the final representation (smallest spatial dimension) after which we use a global average pooling to reduce the features to a one dimensional vector. This is followed by a standard fully connected network to infer the predicted losses.

5. Experimentation and Result

Our code (written in PyTorch [16]) will be made available on <https://github.com/meghshukla/math-analysis-learningloss> and we use open sourced code wherever possible. We report results on the LSP-LSPET datasets and MPII datasets (as done in [12, 24]) with the following experiments: 1) Correlation between predicted loss and true loss [24] 2) Simulating 5-7 active learning cycles [17]. We take a moment to remind the reader that active learning is well suited for large datasets ($> 100,000$ images) and we use these experiments to draw intuition to support the transferability of active learning for human pose estimation to practical applications.

Dataset: The MPII [1] dataset contains images capturing every day human activity. In contrast, the LSP-LSPET [7] dataset represents sports poses, such as those encountered in athletics and parkour. We use standard testing splits: Newell validation split [15], [24] for the MPII




Figure 5. LearningLoss++ yields a higher degree of correlation with the true loss. Not only is the mean loss for the sampled images (black) higher, the true loss variance within the sampled points is lower. (Best viewed when zoomed)

dataset, and the latter thousand images of the LSP dataset [7] for the combined LSP-LSPET dataset. While the MPII dataset serves to measure the general performance of our algorithms, the LSP-LSPET dataset truly represents active learning as a process in industrial applications. Real world applications imply non-stationarity of data, hence the training as well as validation set at any time instant do not represent the entire dataset (which is not the case with the Newell validation split). The LSP-LSPET experiment allows us to study this characteristic, where our initial training and testing pool consists of samples drawn only from the LSP dataset, with subsequent stages allowing the model to sample from the similar yet different LSPET data distribution. This represents a challenge to active learning algorithms to generalize to changes in the data distribution based on the data which is already annotated.

Experiment Design: We simulate active learning cycles using the stacked hourglass model by training an initial base model on 1000 images, with each active learning stage selecting a new set of 1000 images from the remaining unlabelled pool of training data. The initial 1000 images for MPII are randomly sampled from the training data, whereas the initial 1000 images for LSP-LSPET are the first 1000 images from the LSP dataset (LSP consists of 2000 images, split 1000 for train and remaining form the testing set). We follow standard evaluation metrics by using PCKh@0.5 to evaluate model performance on MPII and PCK@0.2 for the LSP-LSPET images. We repeat each experiment five times and report the mean and standard deviation for the same. We limit our experiments to single person pose estimation, hence extract multiple persons in an image into separate im-

ages using the ground truth / location information from the datasets.

Failure identification: A key question remains: *How do we quantify model failures?* To eliminate subjective bias when qualitatively defining poor performance, we use the PCK/PCKh accuracy metric to identify poor performance over a set of images. A low PCK/PCKh score corresponds to the model consistently drawing wrong inferences for that particular set of images. Our justification for using PCK/PCKh is that since these are universally accepted metrics to quantify the performance of the model, they can also be used to identify the degree to which the models perform poorly on the learning loss sampled images.

Comparison Algorithms: We compare our work with Learning Loss and Coreset. As noted in the related work section, DeepPrior by Caramau *et al.* [3] uses dropouts as well as directly regress the joint coordinates, both of which prevent its application in human pose estimation. Entropy based approaches [12] have been compared with Learning Loss [24] where the latter performs better than the former in identifying images with high true losses.

Results: We first compare the correlation between the predicted losses and the true losses for Learning Loss and LearningLoss++, which is shown in Fig: 5. The graph is computed for all the images from the LSPET dataset, with the model trained on the initial pool of images from the LSP dataset. Since the model is trained on LSP dataset and not LSPET, *both the approaches are tested on their ability to identify failures and generalize to datasets from a different distribution.* We observe that the true loss mean of the images sampled with LearningLoss++ is higher as

(a) Failure Detection: PCK scores for the images sampled at Stage n . (Lower PCK values indicate better identification of faulty inferences.)

# images	LSP-LSPET (PCK@0.2)					MPII (PCKh@0.5)				
	2000	3000	4000	5000	6000	1000	2000	3000	4000	5000
Random	0.430 ±0.017	0.527 ±0.012	0.593 ±0.007	0.624 ±0.009	0.645 ±0.007	0.663 ±0.012	0.739 ±0.013	0.766 ±0.003	0.792 ±0.007	0.797 ±0.006
Coreset	0.288 ±0.017	0.438 ±0.020	0.447 ±0.017	0.493 ±0.013	0.556 ±0.010	0.384 ±0.014	0.522 ±0.009	0.608 ±0.012	0.697 ±0.009	0.755 ±0.029
LL	0.305 ±0.013	0.253 ±0.021	0.358 ±0.025	0.520 ±0.011	0.617 ±0.017	0.311 ±0.036	0.465 ±0.024	0.621 ±0.017	0.735 ±0.012	0.777 ±0.010
LL++	0.250 ±0.011	0.186 ±0.022	0.385 ±0.011	0.533 ±0.020	0.627 ±0.012	0.291 ±0.022	0.439 ±0.018	0.610 ±0.020	0.705 ±0.023	0.762 ±0.014
LL++conv	0.209 ±0.018	0.214 ±0.028	0.400 ±0.010	0.545 ±0.011	0.635 ±0.012	0.309 ±0.029	0.439 ±0.011	0.603 ±0.016	0.704 ±0.022	0.777 ±0.008

(b) Testing performance: PCK scores for the testing dataset after each sampling iteration. (Higher values are better)

# images	LSP-LSPET (PCK@0.2)					MPII (PCKh@0.5)				
	Stage 1	Stage 2	Stage 3	Stage 4	Stage 5	Stage 1	Stage 2	Stage 3	Stage 4	Stage 5
Random	0.803 ±0.003	0.818 ±0.002	0.827 ±0.003	0.834 ±0.003	0.841 ±0.002	0.76 ±0.006	0.783 ±0.006	0.803 ±0.009	0.813 ±0.004	0.822 ±0.007
Coreset	0.797 ±0.008	0.814 ±0.004	0.822 ±0.004	0.831 ±0.004	0.837 ±0.003	0.766 ±0.006	0.792 ±0.007	0.812 ±0.007	0.822 ±0.011	0.83 ±0.011
LL	0.796 ±0.004	0.814 ±0.003	0.823 ±0.004	0.833 ±0.003	0.842 ±0.005	0.763 ±0.007	0.793 ±0.005	0.814 ±0.004	0.829 ±0.004	0.838 ±0.005
LL++	0.798 ±0.005	0.810 ±0.003	0.824 ±0.006	0.831 ±0.004	0.841 ±0.002	0.763 ±0.007	0.791 ±0.008	0.811 ±0.005	0.822 ±0.005	0.827 ±0.009
LL++conv	0.798 ±0.005	0.813 ±0.002	0.825 ±0.003	0.832 ±0.005	0.840 ±0.005	0.760 ±0.009	0.790 ±0.008	0.808 ±0.005	0.828 ±0.003	0.836 ±0.006

Table 3. LearningLoss++ outperforms others methods in failure identification Tab:(a), with similar testing accuracy as others Tab:(b). Identifying and fixing faulty inferences (high predicted loss) improves reliability of the model in open world use cases. Detecting poor inferences also allows for specific use case based collection of data, removing the need for large scale data collection in an attempt to identify model failure cases.

well as exhibits a lower variance among the true loss values in comparison to Learning Loss, implying that LearningLoss++ consistently identifies images where the model performs poorly. While Learning Loss performs better than entropy based approaches (shown in [24]), our method improves upon Learning Loss to better identify images with a high value of true loss. Similar behaviour, albeit with a smaller performance gap is obtained over the MPII dataset.

The results of the simulations for the active learning cycles are shown in Table: 3. We identify two phases when training a deep learning model: development and saturation phase. The development phase forms the first few cycles where the model has not converged with respect to the validation dataset. In comparison, the saturation stage marks the convergence of the model on the validation dataset, as the training dataset effectively represents the validation dataset. In practice, our models are rarely saturated because of the shift in distribution of the new incoming data as well as the potentially significant number of failure cases that arise during open world usage. Therefore, the development phase dominates model training and is of practical interest to us. Table: 3(a) highlights the performance of various methods in failure detection. We see that both LearningLoss++ variants return lower PCK scores during the first few stages for both the datasets, indicating successful identification of images where the inference is faulty. With the larger MPII dataset, this trend continues into the model saturation phase. However, the smaller LSPET dataset has a limited number of *tough* examples which are identified early on by LearningLoss++ as well as the original Learning Loss module. The lack of other difficult images cause a sharp rise in the PCK scores for the subsequent stages. We quickly note that detecting failures accurately does not necessarily improve general model performance; failure cases

are usually sparsely represented in the testing dataset which is why a quantifiable increase in accuracy is not detected. We also focus on the results in Stage 1, since stage-1 sampling is performed on the base models shared by all approaches. The ability of the convolutional network to generalize to LSPET images is highlighted in the Stage 1 results for LSP-LSPET dataset. With a PCK@0.2 value of 0.209, the convolutional network (LL++conv) comfortably outperforms its non convolutional counterpart (LL++) which has a PCK@0.2 value of 0.250. We attribute the superior performance of the convolutional architecture in LSP-LSPET to the complex poses that are usually encountered in sporting events. These complex poses are efficiently modeled by the convolutional architecture since it maintains spatial dependencies at multiple scales.

6. Conclusion

This work provides a strong mathematical foundation for Learning Loss, a popular active learning technique. We develop insights into the training of the learning loss module, and propose LearningLoss++ that uses an equivalent KL divergence based objective with additional benefits. Not only is the proposed objective free of any hyperparameters, the resultant gradient is smooth, which allows for better detection of faulty inferences. We also propose a convolutional architecture that exploits the spatial dependency in human pose estimation model. Our experiments show that LearningLoss++ delivers a strong performance in continuous model refinement through identification and early detection of faulty model inferences.

Acknowledgment: We thank Brijesh Pillai and Partha Bhattacharya at Mercedes-Benz R&D India for providing the funding and compute hardware required for this work.

References

- [1] 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)*, June 2014.
- [2] William H Beluch, Tim Genewein, Andreas Nürnberger, and Jan M Köhler. The power of ensembles for active learning in image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9368–9377, 2018.
- [3] Razvan Caramalau, Binod Bhattarai, and Tae-Kyun Kim. Active learning for bayesian 3d hand pose estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3419–3428, 2020.
- [4] Ehsan Elhamifar, Guillermo Sapiro, Allen Yang, and S Shankar Sasrty. A convex optimization framework for active learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 209–216, 2013.
- [5] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. *arXiv preprint arXiv:1703.02910*, 2017.
- [6] Nitesh B Gundavarapu, Divyansh Srivastava, Rahul Mitra, Abhishek Sharma, and Arjun Jain. Structured aleatoric uncertainty in human pose estimation. In *CVPR Workshops*, volume 2, 2019.
- [7] Sam Johnson and Mark Everingham. Learning effective human pose estimation from inaccurate annotation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2011.
- [8] Ajay J Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos. Multi-class active learning for image classification. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2372–2379. IEEE, 2009.
- [9] Tejaswi Kasarla, Gattigorla Nagendar, Guruprasad M Hegde, Vineeth Balasubramanian, and CV Jawahar. Region-based active learning for efficient labeling in semantic segmentation. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1109–1117. IEEE, 2019.
- [10] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in neural information processing systems*, pages 5574–5584, 2017.
- [11] David D Lewis and William A Gale. A sequential algorithm for training text classifiers. In *SIGIR'94*, pages 3–12. Springer, 1994.
- [12] Buyu Liu and Vittorio Ferrari. Active learning for human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4363–4372, 2017.
- [13] Christoph Mayer and Radu Timofte. Adversarial sampling for active learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3071–3079, 2020.
- [14] Prem Melville and Raymond J Mooney. Diverse ensembles for active learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 74, 2004.
- [15] 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.
- [16] Adam Paszke and Sam Gross et al. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [17] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. *arXiv preprint arXiv:1708.00489*, 2017.
- [18] Burr Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.
- [19] H Sebastian Seung, Manfred Opper, and Haim Sompolinsky. Query by committee. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 287–294, 1992.
- [20] 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 (CVPR)*, June 2019.
- [21] Jonathan J Tompson, Arjun Jain, Yann LeCun, and Christoph Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. In *Advances in neural information processing systems*, pages 1799–1807, 2014.
- [22] A. Toshev and C. Szegedy. Deeppose: Human pose estimation via deep neural networks. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1653–1660, 2014.
- [23] Yi Yang, Zhigang Ma, Feiping Nie, Xiaojun Chang, and Alexander G Hauptmann. Multi-class active learning by uncertainty sampling with diversity maximization. *International Journal of Computer Vision*, 113(2):113–127, 2015.
- [24] D. Yoo and I. S. Kweon. Learning loss for active learning. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 93–102, 2019.
- [25] X. You, R. Wang, and D. Tao. Diverse expected gradient active learning for relative attributes. *IEEE Transactions on Image Processing*, 23(7):3203–3217, 2014.
- [26] Y. Yuan, S. Chung, and H. Kang. Gradient-based active learning query strategy for end-to-end speech recognition. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2832–2836, 2019.