

Hinge-Wasserstein: Estimating Multimodal Aleatoric Uncertainty in Regression Tasks

Ziliang Xiong¹, Arvi Jonnarth^{1,3}, Abdelrahman Eldesokey²,
Joakim Johnander⁴, Bastian Wandt¹, Per-Erik Forssén¹

¹Computer Vision Laboratory, Department of Electrical Engineering, LiU, Sweden

²Visual Computing Center, KAUST, Saudi Arabia

³Husqvarna Group, Huskvarna, Sweden, ⁴Zenseact, Sweden

{name.surname}@{liu.se, kaust.edu.sa, zenseact.com}

Abstract

Computer vision systems that are deployed in safety-critical applications need to quantify their output uncertainty. We study regression from images to parameter values and here it is common to detect uncertainty by predicting probability distributions. In this context, we investigate the regression-by-classification paradigm which can represent multimodal distributions, without a prior assumption on the number of modes. Through experiments on a specifically designed synthetic dataset, we demonstrate that traditional loss functions lead to poor probability distribution estimates and severe overconfidence, in the absence of full ground truth distributions. In order to alleviate these issues, we propose hinge-Wasserstein – a simple improvement of the Wasserstein loss that reduces the penalty for weak secondary modes during training. This enables prediction of complex distributions with multiple modes, and allows training on datasets where full ground truth distributions are not available. In extensive experiments, we show that the proposed loss leads to substantially better uncertainty estimation on two challenging computer vision tasks: horizon line detection and stereo disparity estimation.

1. Introduction

Deep neural networks have revolutionized computer vision, producing accurate predictions on a large variety of tasks. However, for safety-critical applications, it is crucial to also quantify the uncertainty of predictions. Observations in many tasks are inherently stochastic, e.g., low-resolution measurements or occlusions of the region of interest. These observations are usually referred to as being subject to *aleatoric uncertainty* [17], which cannot be reduced, even given more collected data. In many common regression tasks, e.g. depth estimation and object pose es-

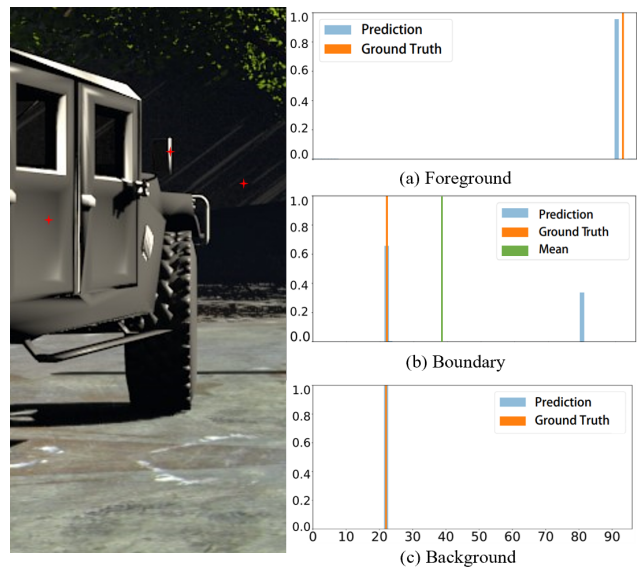


Figure 1. Edge pixels are subject to multimodal aleatoric uncertainty. Left: Three pixels are marked with red stars of the left frame. Right: Predicted disparity distributions of the three pixels, horizontal axis is disparity (same range for three subfigures), vertical axis is probability on (a) foreground (door of the car), (b) boundary pixel (edge of the mirror), and (c) background.

timation, the aleatoric uncertainty is usually described with multimodal distributions. Popular training objectives for regression tasks are the L_1 and L_2 loss. However, these losses assume that the data follows, at least partially, a Gaussian (L_2) or Laplacian (L_1) distribution [24], thereby neglecting multimodal distributions and returning the mean of multiple modes. However, in the multimodal case as shown in Fig. 1, the mean can deviate from the *most likely mode* (which is usually the desired prediction) and is instead located in a region of low likelihood. Specifically, boundary pixels will have a high likelihood of both the foreground and the background, e.g., the predicted disparity map in Fig. 1 (b) shows

that the disparity likelihood is 70% at 20px, and 30% at 80px. The mean (38px) is by no means the real foreground disparity, and for some applications, *e.g.*, autonomous driving, decisions based on this perception might be fatal.

A flexible approach to model multimodal distributions is *regression-by-classification* (RbC) [6, 12, 34] where the regression space is partitioned into a fixed set of bins, and the task is to predict the probability of the regressed value falling into each bin. Figs. 1 and 2 show examples of RbC outputs from the proposed method on stereo disparity and horizon line estimation respectively. Furthermore, in RbC it is often argued that the Wasserstein loss should be used [8, 12, 22], as the regression space partitioning induces inter-class correlations. For instance, a predicted bin close to the ground truth mode is generally better than the one that is far away. Another advantage of the Wasserstein loss is that it allows for multimodal training targets, which improves regression performances [8]. Despite the merits of Wasserstein loss, there is a remaining issue untackled: **in most realistic tasks, such multimodal ground truth labels do not exist**. To the best of our knowledge, none of the previous RbC works consider aleatoric uncertainty estimation in scenarios where multimodal ground truth distributions are unavailable.

To analyze the influence of different loss functions on estimating aleatoric uncertainty, we create a synthetic dataset for which the aleatoric uncertainty can be controlled. This analysis reveals a major downside of the plain Wasserstein loss, namely poor uncertainty estimation when the full target distribution is missing. To mitigate this, we introduce *hinge-Wasserstein*, an improved version of the Wasserstein loss comprising a hinge-like mechanism during loss computation. This allows weak secondary modes to exist in the predicted distribution, and is thereby able to reduce overconfidence, especially when full ground truth distributions are unavailable. We provide an ablation study that shows the effectiveness of the proposed loss on representing multimodal distributions with unimodal targets by comparing with plain Wasserstein and synthetic multimodal targets. We further demonstrate that the proposed loss significantly improve uncertainty estimation while maintaining the main task performances on two common regression tasks, namely horizon line detection and stereo disparity estimation.

To summarize, our main **contributions** are as follows:

- We show that the plain Wasserstein loss trained with unimodal targets leads to poor uncertainty estimation. Moreover, we prove that the Wasserstein loss combined with a terminal softmax layer leads to vanishing gradients.
- We propose *hinge-Wasserstein*, which is a easy fix to plain Wasserstein, as a training loss for regression tasks with prevalent aleatoric uncertainty. The extra computation for hinge operation is negligible. Additionally, we prove the proposed loss to be a proper scoring rule.

- Through extensive experiments, we show that hinge-Wasserstein gives better uncertainty estimation while maintaining the main task performance, especially under multimodal aleatoric uncertainty.

2. Related work

Uncertainty quantification: Uncertainty quantification approaches can be divided into (i) *parametric approaches*, which are usually more task-specific; (ii) *ensemble methods* that are expensive but can be applied to most tasks; and (iii) *regression-by-classification*, which outputs full distributions [12]. Parametric approaches assume an output distribution as a form of inductive bias. The neural network is then designed and trained to predict the parameters of this assumed distribution. The target distribution is highly specific for a certain task. *E.g.*, He *et al.* [14] predict the parameters of a Gaussian distribution for the box width and height in object detection. A more complex distribution, Poisson multi-Bernoulli, is adopted in the work of Hess *et al.* [15] for probabilistic object detection. Contrarily, ensemble methods estimate uncertainty by producing a set of outputs that together characterize the uncertainty. A well-known example is the deep ensemble [20] where multiple neural networks – trained for the same task but using different initial seeds and batch shuffling – make predictions on each data point. Another approach that avoids the need for multiple neural networks is Monte Carlo dropout [7]. Inference is conducted via multiple forward passes through a single neural network with random dropout active. Ensemble methods typically require no modifications to the model architecture and can therefore be combined with the parametric approach to achieve better calibration [16]. One disadvantage with ensemble methods, however, is that they are relatively computationally expensive.

Regression-by-classification: The main principle for RbC is to transform regression tasks into classification tasks. This is achieved by discretizing the continuous target variables into bins, given some prior knowledge about the problem or automatically computed ranges. Several classical regression approaches adopted this strategy such as Support Vector Regression [35], ordinal regression [5], and decision tree regression [23]. In deep learning, these classical approaches were adapted to be differentiable, and thus operate in end-to-end learning frameworks [5, 19, 28]. Other approaches employed this concept to solve specific problems. Niu *et al.* [27] tackled the problem of age estimation as ordinal regression using a series of binary classification sub-problems to utilize inherent order of labels. Liu *et al.* [22] incorporated inter-class correlations in a Wasserstein training framework for pose estimation. Garg *et al.* [8] also utilized a Wasserstein training loss and proposed to have an additional offset prediction for each bin. They revealed the benefit of training with synthetic multimodal ground truth




Figure 2. Horizon line detection should be framed as a probabilistic regression problem due to its inherently stochastic nature. Upper Left: Image where horizon line detection is easy (red line) and direct regression would work. Upper Middle and Right: Images where the horizon line is ambiguous. Bottom row: Plots below the images show the output probability distributions for the horizon line parameters (α, ρ) , from the proposed method. Red: Gaussian-smoothed ground truth; Blue: predicted density. Images are from the HLW dataset [34].

for improving stereo disparity. Despite the success of this approach, one major issue is that multimodal ground truth does not exist in real-world datasets and as we will show, training with unimodal ground truth will result in overconfidence. We therefore propose an improvement to the Wasserstein loss that is able to reduce the overconfidence, and thus allow training on real-world datasets.

Multimodal aleatoric uncertainty: The term *aleatoric uncertainty* refers to inherent uncertainty in the observation, e.g. in the form of measurement noise or limited precision, which cannot be alleviated given more data [17]. A task with prevalent measurement noise is horizon line detection, which is to estimate the horizon from a single image. The images are subject to low exposure, motion blur, and occlusions, resulting in aleatoric uncertainty. Workman *et al.* [34] first proposed to use a CNN to directly estimate the horizon. They also published a benchmark dataset, *Horizon Lines in the Wild* (HLW), containing real-world images with labeled horizon lines. The HLW-Net [34] has a GoogleNet [31] backbone trained by soft-argmax and a cross-entropy loss, predicting line parameters. Brachmann *et al.* [3] achieved better performance by predicting a set of 2D points and then fits a line with RANSAC. Although each point has a sampling score, it fails to transfer to the probabilistic distribution of the line. SLNet [21] is an architecture to find semantic lines and surpassed the other approaches. It proposes line pooling layers that extract line features from feature maps. Though a classification head of SLNet decides if a candidate line is semantic, it does not address the distribution of the line. In the experiment, we build upon the work of Workman *et al.* [34], given the open-source dataset and its end-to-end model architecture.

Disparity estimation is another task with inherent aleatoric uncertainty [17], caused by the depth discontinuities in most natural scenes. Especially boundary pixels between foreground and background objects are likely to be inherently multimodal, having two modes for both objects. Garg *et al.* [8] revealed the benefit of training with synthetic multimodal ground truth for improving stereo dis-

parity. However, it neither assesses the uncertainty evaluation nor addresses the unavailability of multimodal ground truth. Häger *et al.* [12] improve the uncertainty estimation by setting a maximum-entropy distribution as the target for those occluded pixels, and evaluate uncertainty using *spar-sification plots*. We adopt their uncertainty evaluation and show that the proposed hinge- W_1 can further improve uncertainty estimation, in particular on non-synthetic datasets, where it alleviates the lack of multimodal ground truth.

3. Method

We first review the theory for training Regression-by-Classification Networks in Sec. 3.1 and the closed-form solution of the Wasserstein distance in Sec. 3.2. We then introduce hinge-Wasserstein loss in Sec. 3.3 and prove it to be a proper scoring rule. Finally we introduce uncertainty evaluation metrics in Sec. 3.4.

3.1. Training regression-by-classification networks

In *regression-by-classification* (RbC), a regression variable $y \in \mathbb{R}$ is discretized into K bins. A neural network Z then predicts a conditional probability $p(y|\mathbf{x})$ given the evidence \mathbf{x} , which in our case is an image. The output of the network $\hat{p}_y = Z(\mathbf{x})$ in the K -probability-simplex is a vector where each element represents the probability that the regression variable y lies in a specific interval k ,

$$\hat{p}_y[k] \approx P(v_k < y < v_{k+1} | \mathbf{x}), \quad (1)$$

where $\{v_k\}_{k=1}^{K+1}$ are the bin edges. The final regression prediction \hat{y} is obtained by applying a decoding function to the output vector \hat{p}_y , i.e., $\hat{y} = \text{dec}(\hat{p}_y)$. The decoding function dec can be defined in different ways as explained in [6] depending on the exact representation of p_y . A straightforward decoding is to extract the maximum value of the output vector \hat{p}_y as in [34]. Besides decoding, in Sec. 3.4 we discuss computing an uncertainty measure from \hat{p}_y .

Training RbC networks requires defining a loss on the output vector \hat{p}_y with respect to some ground truth annota-

tion p_y^* . Usually only a single continuous value, y^* , is given as ground truth. This can be interpreted as the corresponding distribution being a Dirac impulse, $p_y^* = \delta[y - y^*]$. A natural choice of loss for this setting is the NLL of y^* under the predicted distribution, \hat{p}_y [34]. Kendall *et al.* [18] instead propose to minimize the decoding error, \hat{y} . But Häger *et al.* [12] argue that this can lead to a biased output when aleatoric uncertainty is present.

3.2. Wasserstein loss

The output labels in a RbC network have an inherent order, see (1). The errors caused by predicting a bin close to the correct one are thus less severe than predicting one farther away. This detail is ignored by the standard cross-entropy loss, KL divergence, etc, and it has thus been argued that the Wasserstein loss is a better fit for regression-by-classification [22]. The Wasserstein loss between a predicted distribution $p(y)$ and the ground truth $q(y)$ over some variable y is defined as

$$W_m(p, q) = \inf_{\gamma \in \Gamma(p, q)} \int \int |y' - y|^m \gamma(y, y') dy dy' \quad , \quad (2)$$

where $\Gamma(p, q)$ is the set of all possible transport plans that take q to p . For $m = 1$ the search over transport plans in (2) can be avoided [32, 33], and in (1), the output distribution is discretized, which leads to:

$$W_1(p, q) = \sum_{k=1}^K |P[k] - Q[k]| \quad . \quad (3)$$

Here, P and Q are the CDFs of p and q respectively, *i.e.* $P[k] = \sum_{k=1}^K p[k]$ and $Q[k] = \sum_{k=1}^K q[k]$. Further, for non-synthetic datasets *e.g.* [34], the annotations are Dirac distributions, $q(y) = \delta(y - y^*)$. This also avoids the search over transport plans [9]:

$$W_m(p_y, q) = \left(\sum_k p_y[k] (\hat{y}(p) - y^*(q))^m \right)^{1/m} \quad . \quad (4)$$

Softplus Activation: Ensuring that \hat{p}_y in (1) is on the K -probability-simplex is most commonly achieved via the softmax function. Below, we show that vanishing gradients occur for the combination of a softmax layer and the W_1 distance with Dirac ground truth. Softmax, commonly used as the final layer for multi-class classifiers, is defined in (5).

$$g_i(\mathbf{z}) = \frac{e^{z_i}}{\sum_{j=1}^N e^{z_j}} \quad , \quad \text{for } i = 1, 2, \dots, N, \quad (5)$$

where N is the number of classes, and $\mathbf{z} = [z_i, i = 1, 2, \dots, N]$. Softmax normalizes its output to sum to one. Thus, we can treat the output $\mathbf{g}(\mathbf{z}) = [g_1(\mathbf{z}), g_2(\mathbf{z}), \dots, g_N(\mathbf{z})]$ as the predicted categorical distribution for each class. The ground truth label is $\delta[j - j^*]$,

where j^* is the true class index. In this case, the W_1 distance can be simplified as

$$W_1(\delta[j - j^*], \mathbf{g}(\mathbf{z})) = \sum_i |i - j^*| g_i(\mathbf{z}), \quad (6)$$

The partial derivative of (6) with respect to z_k is

$$\frac{\partial W_1(\delta[j - j^*], \mathbf{g}(\mathbf{z}))}{\partial z_k} = g_k(\mathbf{z})(|k - j^*| - \sum_i |i - j^*| g_i(\mathbf{z})). \quad (7)$$

There are two problematic cases for (7):

Case 1: Low initial value for the correct bin. For the correct bin, where $k = j^*$, $|k - j^*|$ in (7) is zero. Further, if $g_k(\mathbf{z}) \approx 0$, we have a small, but always negative, contribution from the sum, since $g_k(\mathbf{z})$ is a factor of (7). Thus, if $g_k(\mathbf{z})$ starts out as a small value, it will be hard to change the value. In the limit case, where $g_k(\mathbf{z}) = 0$ the gradient is zero, resulting in no learning at all.

Case 2: Dominant, but incorrect mode. If there is a dominant, but incorrect mode $m \neq j^*$, the contribution to the gradient will also be low. In the limit we have $g_m(\mathbf{z}) = 1$, $g_i(\mathbf{z}) = 0$, for $i \neq m$ (as $\sum_i g_i(\mathbf{z}) = 1$). In this case, both terms in (7) also become close to 0, for all values of k . Thus, whenever there is a peak of a wrong class close to 1, this will also cause vanishing gradients.

We instead follow Häger *et al.* [12], applying softplus function followed by an l_1 normalization, which the authors found beneficial for regression-by-classification. Our experiments in Sec. 4.3 show that simply switching to softplus will slightly improve uncertainty estimation and regression.

3.3. Hinge-Wasserstein loss

Most of the datasets used in computer vision are annotated with unimodal ground truth, *i.e.*, instead of full conditional density annotations, only the most likely output is provided as the training target. When we use a loss that rewards output of only a single mode at the annotation, we effectively discourage the output to represent aleatoric uncertainty, which in turn causes overconfidence. In the extreme case, when there is no evidence to support *any* hypothesis, the output in (1) should be uniformly distributed.¹

An intuitive way to use unimodal ground truth to learn probability density outputs is to discount the loss for all bins by a margin γ_W , and not penalize bins that are below this level. This will allow aleatoric uncertainty in the input to be represented in the output. This is similar to the hinge losses from support vector machines and to the triplet loss used in contrastive learning [30], where it only matters if the distance to a negative sample is larger up to a certain point. Similarly, we allow incorrect bins to be non-zero,

¹Note that what we expect to see is the dataset prior. The bins, however, are usually chosen as quantiles of the training set, which results in a uniform distribution over the bins.

as long as they are sufficiently below the level of the main mode. In detail, the predicted probability density will be reduced by the threshold, passed through a ReLU and then re-normalized to sum to one,

$$\tilde{p}[k] = \max(p[k] - \gamma_W, 0) , \quad (8)$$

$$\bar{p}[k] = \frac{\tilde{p}[k]}{\sum_k \tilde{p}[k]} . \quad (9)$$

The loss is then defined as the Wasserstein distance between the the renormalized probability densities and ground truths,

$$\mathcal{L}(p, q) = W_1(\bar{p}, q) . \quad (10)$$

We call this new loss the *hinge-Wasserstein* loss (abbreviated hinge- W_1). Note that (10) works for both implementations of W_1 as described by (3) and (4), i.e., it allows both a Dirac training target and a full distribution target. The parameter γ_W depends on the total number of output bins, and is normally set as $\gamma_W = 1/K$. This can be interpreted as the hinge-Wasserstein loss allowing a random guess (which would correspond to a uniform distribution).

Proper scoring rule A *scoring rule* $S(p_\theta, q)$ [10] is a function that evaluates the quality of a predictive distribution $p_\theta(y|\mathbf{x})$, with respect to a true distribution q . A *proper scoring rule* [10] should satisfy: $S(p_\theta, q) \leq S(q, q), \forall p_\theta$, with equality if $p_\theta = q$. If the inequality is strict for all $p_\theta \neq q$, S is *strictly proper*. It has been shown that W_m is not proper for finite samples [32]. However, we use discrete distributions with $m = 1$, which can be proved to be proper.

Proposition: *Discrete $-W_1$ is strictly proper.*

Proof by contradiction: First we note that $-W_1(q, q) = 0$ for any PDF q , see (3). Thus, setting $S = -W_1$ results in the requirement: $W_1(p_\theta, q) \geq 0$ with equality iff $p_\theta = q$.

Assumption: There exists a discrete distribution p that is different from q , for which $W_1(p, q) = 0$. In p , at least two bin values must be different from those in q , as $\sum_k p[k] = 1$ must hold. Denote their bin indices by m_1, m_2 , where $m_1 < m_2$, without loss of generality.

Contradiction: Consider (3) is a sum of non-negative values (since the absolute value). According to the assumption we have $p[m_1] \neq q[m_1]$, and thus, $D[m_1] > 0$ and as a consequence $W_1(p, q) > 0$. This is contradictory to the assumption. For cases where more than two bin values are different, we will get more non-zero terms $D[k]$ in S contributing to the contradiction. This proves the equality uniqueness. \square

Based on the proposition above, it can be easily shown that *hinge-Wasserstein is a proper scoring rule but not strictly proper* because with the hinge mechanism p_θ that achieves the optimal w_1 distance is not unique.

3.4. Predictive uncertainty evaluation

We evaluate uncertainty estimation capabilities using *sparsification plots* and *CRPS*.

Sparsification Plots: A *sparsification plot* [1, 16] assesses how well the predicted uncertainty coincides with the output error. This requires a scalar uncertainty value u , which is a function of the network probability density output, $\hat{p}_y = Z(\mathbf{x})$,

$$u = f(\hat{p}_y) . \quad (11)$$

A *sparsification curve* is a plot of the mean absolute error (MAE) as a function of a fraction p of samples that have been removed. The removed samples are those with the highest predicted uncertainty according to u in (11). Similarly, an *oracle curve* is created by instead removing the fraction p of samples with the highest absolute errors. The oracle curve as a monotonically decreasing function indicates the lower bound of sparsification curves. Both curves are normalized by the MAE on all the test set samples, i.e. they start at $(0, 1)$. See Fig. 3 in [16] for an example of sparsification plots. By plotting the vertical distance between the two curves, defined as the sparsification error, against p , we get the sparsification error plot as in Fig. 4. We use *area under sparsification error* (AUSE) to quantitatively evaluate uncertainty estimation to fairly compare different approaches with different oracles. A common choice for u , is the Shannon entropy [12] [26] [29] of \hat{p}_y ,

$$u_H(\hat{p}_y) = H(\hat{p}_y) = - \sum_{k=1}^K \hat{p}_y[k] \log \hat{p}_y[k] . \quad (12)$$

We show in the supplement the comparison between entropy and other uncertainty measures, e.g., variance, maximum value of the predicted histogram, and (12) returns best AUSE. Hence, we show experiment results with entropy.

CRPS: While AUSE measures the correlation between uncertainty magnitude and prediction error, it disregards the shape of each predicted probability distribution. As the goal is to predict multimodal distributions, we can use a *proper scoring rule* to assess this. A proper scoring rule addresses calibration and sharpness simultaneously [11]. Here we extend *continuous ranked probability score* (CRPS) to allow multimodal ground truth. Specifically, the step function in CRPS is replaced with the cumulative distribution of equally-weighted Dirac mixtures: $q(y) = \frac{1}{N} \sum_i^N \delta(y_i^*)$,

$$\text{CRPS}(p_\theta, q) = \int_{-\infty}^{\infty} |p_\theta(y) - q(y)|^2 dy . \quad (13)$$

Note that this metric only works when multimodal ground truth distribution is available.

4. Experiments

We analyze the performance of the proposed hinge- W_1 loss under multimodal aleatoric uncertainty. First, we construct a synthetic dataset, where we control the presence of aleatoric uncertainty, to show that using the nega-

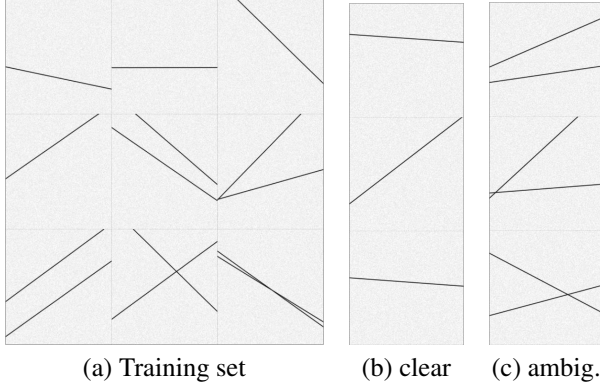


Figure 3. Example images from the synthetic dataset with controllable aleatoric uncertainty: (a) training set, one or two lines per image, one line in the annotation for unimodal training; both lines in the annotation for multimodal training; (b) test set 1, one line per image; and (c) test set 2, two lines per image.

tive log-likelihood or Wasserstein losses lead to overconfidence. Our proposed hinge- W_1 , in contrast, reports uncertainty whenever uncertainty is present. Next, we evaluate our hinge- W_1 loss on two separate regression tasks: horizon line regression on *Horizon Lines in the Wild* (HLW) (Sec. 4.2), and stereo disparity estimation on *Scene Flow Datasets* (Sec. 4.3).

4.1. Synthetic dataset

Dataset: To create a controlled environment where we can analyze the behaviour of our proposed approach, we choose a simplistic yet representative regression problem – parameter estimation for lines in noisy images. Aleatoric uncertainty is introduced in the form of multiple lines in the images as shown in Fig. 3. One or more randomly generated lines are rendered on a background with Gaussian noise. These image-line pairs can be used for supervised training. We can easily control the aleatoric uncertainty amount by varying the fraction of images with multiple lines.

The training set contains 2000 clear images with one line and 2000 ambiguous images with two lines. We create two test sets: The clear set shown in Fig. 3 (b), contains 500 images with one line per image; the ambiguous set shown in Fig. 3 (c), contains 500 images with two lines per image. To imitate real-world tasks, a randomly-picked line in the ambiguous images is labeled as the unimodal ground truth. By labeling both lines in the ambiguous images, we also create multimodal ground truth distributions for α and ρ (called MM GND in the experiments).

Implementation details: A line is often parameterized as a slope and an offset (α, ρ) , which define the set of pixel coordinates (x, y) as $(-\sin \alpha \quad \cos \alpha \quad -\rho)^T (x \quad y \quad 1) = 0$, $(x, y) \in \Omega$, where Ω is the set of image coordinates. To make the representation unique, we also restrict the parameters to $\alpha \in [-\pi/2, \pi/2]$ and $\rho \in [0, \infty)$. Then α and ρ are discretized into $K=100$ bins. The first 99 bins are linearly-

spaced from the minimum value of training samples to the maximum, and the last bin is from the maximum to infinity as in [34]. All models use a *Resnet18* [13] as the backbone network, with the last layer replaced by a linear layer followed by two prediction heads, each of which consists of a softplus and normalization as in Sec. 3.2, for the slope and the offset respectively. We assume that \hat{p}_α and \hat{p}_ρ are independent. We conduct an ablation study on NLL, W_1 and hinge- W_1 losses with both (i) unimodal and (ii) multimodal ground truths. For computation, we follow [12] and apply small Gaussian smoothing around the unimodal target. Note that due to the use of an argmax decoding (inherited from [34]) there always exists a quantization error given by the bin sizes, and thus an upper bound on the AUC. This error could be reduced by a more advanced decoding (see [6]) or by adding an offset prediction branch (see [9]).

Metrics: To evaluate line regression performance, we employ a metric that is commonly used in the similar task of horizon line detection, the *horizon detection error*, proposed by Barinova *et al.* [2]. It is calculated as the maximum vertical distance between ground truth and predicted lines in the image, normalized by the height of the image. The cumulative histogram of the horizon detection error is often used to assess the error distribution for the test set, and the area under the curve (AUC) is commonly reported as a summary. For uncertainty evaluation, we apply AUSE and CRPS (Sec. 3.4) separately on the line parameters.

Results: Fig. 5 shows qualitative examples of inferences. The hinge- W_1 captures the multimodality in the input in most cases, except for Fig. 5 (b), the *metameric case* [6], where the modes are too close and interfere. Table. 1 shows the quantitative results of both line regression performance and uncertainty evaluation on the two-line test set. The model trained with plain Wasserstein and unimodal targets shows overconfidence with high AUSE and CRPS, which indicates plain W_1 severely omits secondary modes. Though usually unavailable, adding multimodal targets to training largely improves the uncertainty scale (AUSE) and predicted distribution (CRPS) while slightly improves regression. While increasing hinge value γ_W , AUSE and CRPS keep improving until the γ_W is too large. The AUSE of hinge Wasserstein with $\gamma_W=0.015$ are close to those with multimodal targets. A large hinge damages the training because there are few bins in the predicted density above the hinge at the beginning, thus no gradient to backpropagate.

4.2. Horizon lines in the wild

Dataset: To show that the findings in Sec. 4.1 for the synthetic dataset transfer to a related task in the real world, we evaluate our approach on the challenging *Horizon Lines in the Wild* (HLW) [34] benchmark. HLW is a large dataset of real-world images captured in a diverse set of environments, with horizon lines annotated using *structure from motion*. It




Figure 4. Sparsification error curves for the HLW task (lower is better). (a) α entropy and absolute error as the oracle. (b) Same setting for ρ entropy. See Table 3 for AUSE.

contains 100553 training images and 2018 test images.

Implementation Details: Labels for the horizon line slope and offset are continuous in HLW. They need to be discretized into K bins respectively and use the one hot vector of bin index as a ground truth q_y . The $K = 100$ bins are chosen to be approximately equally likely to occur, by linearly interpolating the cumulative distribution function of the corresponding parameter over the training set. We follow the training procedure in [34] and replace the backbone with Resnet18. We apply hinge- W_1 as the training loss and compare with the original NLL loss, an ensemble of 5 instances (as in [20]) trained with NLL loss and Plain W_1 .

Metrics: We use the same regression metric (AUC) and uncertainty evaluation metric (AUSE) as in Sec. 4.1 and Sec. 3.4. The only exception is that CRPS is not computed because multimodal ground truth targets are unavailable.

Results: Fig. 4 and Tab. 3 show quantitative results for different configurations of our method, compared to the baseline implementation of the NLL loss [34]. Using plain W_1 matches the baseline in terms of the AUC. However, it causes an even more severe overconfidence problem compared to NLL, as shown by the large AUSE score. Hinge- W_1 beats the baseline by a large margin in terms of AUSE for α and ρ yet only shows a minor decrease in AUC. The sparsification error plots in Fig. 4 indicates that NLL often leads to overconfident predictions for ρ with small error as seen from the huge peak in the sparsification error curve on the right. By contrast, hinge- W_1 can mitigate such overconfidence. In Tab. 3 we report the results for an ensemble of 5 NLL trained networks. As can be seen, the ensemble obtains the best AUC and slightly improves the α AUSE, however at the price of $5\times$ more computation. The ensemble is actually complementary to our approach, and ensembles could potentially be combined with RbC, if a larger compute budget is available.

4.3. Stereo disparity

We also study stereo disparity estimation. Given a pair of rectified stereo images, the task is to predict the disparity between their x -coordinates in every pixel.

Dataset: Following [9] we use the synthetic Scene Flow




Figure 5. Density prediction for α with a model trained using hinge- W_1 with $\gamma_W = 0.01$, and inference on both test sets. Note: Only unimodal ground truth was used during training. Blue shows Gaussian-smoothed ground truth, and orange shows predicted densities. (a) and (c) show examples of where two output peaks overlap the ground truth; (b) shows that the model cannot distinguish two peaks if they are too close; (d) shows the model working well with unimodal ground truth.

dataset [25]. It contains 35k training and 4k test image pairs in different domains, and forms a challenging benchmark for the disparity task. All image pairs are accompanied by unimodal ground truth disparity.

Metrics: To evaluate the regression performance, we measure the end-point-error (EPE), which is the absolute pixel difference between the predicted and true disparity, averaged over all pixels and images. We further measure the k -pixel threshold error (PE), which is the percentage of pixels where the error is more than k pixels. We evaluate PE at 1-pixel and 3-pixel thresholds, denoted 1PE and 3PE. To evaluate the uncertainty estimation, we use AUSE, as described in Sec. 3.4. To further investigate, we compute all metrics on edge pixels, as these represent regions of high uncertainty, where uncertainty estimation plays a larger role.

Implementation details: We train a Continuous Disparity Network (CDN) [8] on top of the PSMNet [4] backbone, which predicts bin center offsets in addition to the bin probabilities. This allows for continuous sub-pixel disparity predictions. We use the same settings as in [8], *i.e.* we train for 10 epochs with batch size 8, and use a constant learning rate of 0.001 with the Adam optimizer. The disparity is discretized into 96 bins of two pixels each, allowing for predictions in $[0, 192]$. Furthermore, we follow [8] in training the models in a multimodal setting, with synthetic multimodal ground truth based on neighboring pixels. In this setting, the true distribution is chosen as a Dirac mixture with disparity values based on a $k \times k$ neighborhood. The weight of the center pixel disparity is set to $\alpha = 0.8$, and the rest are set with equal weight.

Results: Table 2 shows the regression performance and uncertainty evaluation on the disparity task. We compare the results with softmax normalization, softplus normalization, and in multimodal training. We find that increasing the hinge γ_W improves the uncertainty estimation, while retaining a high regression performance, in all three settings. The fact that our proposed loss yields improvements in both unimodal and multimodal settings demonstrates that it is well suited for a wide range of conditions. We find that both softplus normalization and synthetic multimodal disparity

Table 1. Ablation study on the synthetic dataset. MM GND means training with multimodal ground truth (typically unavailable in real-world application); for all other results we use unimodal ground truth. Standard error is computed over five randomly initialized models.

Loss	AUC \uparrow	α AUSE \downarrow	ρ AUSE \downarrow	α CRPS \downarrow	ρ CRPS \downarrow
Plain W_1 with MM GND	46.88 \pm 0.23	19.89 \pm 1.29	19.49 \pm 0.95	4.71 \pm 0.02	5.20 \pm 0.02
Plain W_1	46.36 \pm 0.17	90.49 \pm 3.58	69.67 \pm 3.83	8.72 \pm 0.05	9.42 \pm 0.05
<i>hinge</i> - W_1 $\gamma_W = 0.005$	48.52 \pm 0.12	70.17 \pm 2.36	53.26 \pm 2.62	8.02 \pm 0.04	8.42 \pm 0.09
<i>hinge</i> - W_1 $\gamma_W = 0.01$	49.00 \pm 0.09	63.47 \pm 3.17	52.52 \pm 2.83	6.80 \pm 0.14	7.04 \pm 0.03
<i>hinge</i> - W_1 $\gamma_W = 0.015$	39.80 \pm 0.14	21.97 \pm 1.54	28.38 \pm 5.15	9.52 \pm 0.19	9.27 \pm 0.13

Table 2. Stereo disparity results on Scene Flow. Regression performance in terms of EPE, IPE, and 3PE, and uncertainty evaluation in terms of entropy-based AUSE. MM denotes multimodal training with $k = 5$, and standard error is reported over five runs.

Setting	Loss	All pixels				Edge pixels			
		EPE \downarrow	IPE \downarrow	3PE \downarrow	AUSE \downarrow	EPE \downarrow	IPE \downarrow	3PE \downarrow	AUSE \downarrow
Softmax	Plain W_1 [8]	0.98 \pm 0.01	9.44 \pm 0.06	4.04 \pm 0.03	19.4 \pm 0.37	3.05 \pm 0.03	17.4 \pm 0.12	10.1 \pm 0.10	27.5 \pm 0.70
Softmax	<i>hinge</i> - W_1 , $\gamma_W = 0.01$ (Ours)	0.99 \pm 0.01	9.62 \pm 0.06	4.08 \pm 0.03	18.7 \pm 0.43	3.05 \pm 0.03	17.6 \pm 0.12	10.1 \pm 0.06	26.4 \pm 0.31
Softplus	Plain W_1 [8]	1.00 \pm 0.01	9.74 \pm 0.07	4.12 \pm 0.03	18.1 \pm 0.89	3.05 \pm 0.01	17.5 \pm 0.09	10.1 \pm 0.07	27.2 \pm 1.64
Softplus	<i>hinge</i> - W_1 , $\gamma_W = 0.0025$ (Ours)	0.97 \pm 0.02	9.35 \pm 0.16	3.97 \pm 0.06	16.5 \pm 0.55	2.98 \pm 0.03	17.1 \pm 0.16	9.80 \pm 0.07	23.6 \pm 0.63
Softplus	<i>hinge</i> - W_1 , $\gamma_W = 0.005$ (Ours)	0.96 \pm 0.01	9.31 \pm 0.05	3.96 \pm 0.03	16.0 \pm 0.39	3.00 \pm 0.03	17.1 \pm 0.11	9.84 \pm 0.12	23.5 \pm 0.58
Softplus	<i>hinge</i> - W_1 , $\gamma_W = 0.01$ (Ours)	0.98 \pm 0.01	9.48 \pm 0.06	4.05 \pm 0.03	16.4 \pm 0.48	3.04 \pm 0.02	17.2 \pm 0.16	9.99 \pm 0.10	23.1 \pm 0.77
Softplus	<i>hinge</i> - W_1 , $\gamma_W = 0.015$ (Ours)	0.97 \pm 0.02	9.38 \pm 0.13	3.98 \pm 0.06	16.1 \pm 0.60	3.01 \pm 0.03	17.2 \pm 0.21	9.88 \pm 0.11	22.9 \pm 0.80
Softplus	<i>hinge</i> - W_1 , $\gamma_W = 0.02$ (Ours)	0.98 \pm 0.01	9.46 \pm 0.09	4.04 \pm 0.03	15.5 \pm 0.19	3.03 \pm 0.03	17.2 \pm 0.13	9.94 \pm 0.09	21.8 \pm 0.37
Softplus	<i>hinge</i> - W_1 , $\gamma_W = 0.04$ (Ours)	1.02 \pm 0.01	9.75 \pm 0.09	4.17 \pm 0.05	15.1 \pm 0.26	3.13 \pm 0.02	17.7 \pm 0.14	10.3 \pm 0.08	21.3 \pm 0.49
Softplus, MM	Plain W_1 [8]	1.00 \pm 0.03	9.61 \pm 0.25	4.15 \pm 0.16	14.1 \pm 1.46	3.15 \pm 0.11	17.59 \pm 0.38	10.3 \pm 0.26	19.8 \pm 2.08
Softplus, MM	<i>hinge</i> - W_1 , $\gamma_W = 0.0075$ (Ours)	0.96 \pm 0.01	9.27 \pm 0.11	3.94 \pm 0.03	13.0 \pm 0.32	3.00 \pm 0.03	17.00 \pm 0.17	9.79 \pm 0.13	17.2 \pm 0.01
Softplus, MM	<i>hinge</i> - W_1 , $\gamma_W = 0.01$ (Ours)	0.97 \pm 0.03	9.40 \pm 0.21	4.01 \pm 0.08	12.6 \pm 0.01	3.04 \pm 0.03	17.20 \pm 0.08	9.98 \pm 0.03	16.3 \pm 0.90

Table 3. Test results on *Horizon Lines in the Wild*. For *hinge*- W_1 and *plain Wasserstein*, Gaussian smoothed ($\sigma = 4$) training target apply. Both metrics are multiplied by 100.

Loss	AUC \uparrow	α AUSE \downarrow	ρ AUSE \downarrow
NLL[34]	64.13 \pm 0.04	44.30 \pm 0.86	51.47 \pm 2.56
Ensemble NLL	66.83	39.50	48.00
Plain W_1	64.40 \pm 0.22	100.40 \pm 2.76	153.70 \pm 3.33
<i>hinge</i> - W_1 , $\gamma_W=0.0100$	66.60 \pm 0.11	49.88 \pm 1.73	78.49 \pm 1.61
<i>hinge</i> - W_1 , $\gamma_W=0.0150$	64.24 \pm 0.41	29.18 \pm 0.75	69.41 \pm 5.18
<i>hinge</i> - W_1 , $\gamma_W=0.0200$	62.72 \pm 0.14	26.97 \pm 1.19	30.82 \pm 1.92
<i>hinge</i> - W_1 , $\gamma_W=0.0250$	62.00 \pm 0.09	32.45 \pm 1.36	31.11 \pm 6.11

labels further improve the uncertainty estimation. In the results for edge pixels, the same improvements can be seen.

5. Concluding remarks

Unavailability of multimodal groundtruth: We generate multimodal ground truth for training on the synthetic dataset in Sec. 4.1 and on the disparity task in Sec. 4.3. Experiments show that MM ground truth improves uncertainty evaluation and regression for plain W_1 . When we have prior knowledge about the full multimodal distribution, synthetic MM groundtruth is a way to improve uncertainty. On top of this, *hinge*- W_1 can bring further improvements. However, in most real-world datasets, we have neither the full ground truth distribution nor the prior knowledge about it. Our proposed *hinge*- W_1 can mitigate the absence of multimodal ground truth by reducing the penalty for secondary modes in the prediction. A network trained with *Hinge*- W_1 and unimodal ground truth has an AUSE that is close to that

from plain- W_1 and multimodal ground truth.

Choice of hinge: All the experiments on different tasks reveal that the uncertainty in terms of AUSE keeps improving as the hinge value increases from 0 until that hinge is too large to maintain stable training. Though not optimal, the hinge value $1/K$ (K is the number of bins) already significantly improves AUSE. It is worth noticing that the optimal hinge is usually smaller in terms of CRPS than AUSE. This is because CRPS, as a proper scoring rule, requires sharpness of the predictive distributions. The most common use of uncertainty is to detect problematic outputs, which is more closely related to AUSE (AUSE measures how well the uncertainty measure is able to sort the test set samples).

We have analyzed the behaviour of the Wasserstein loss on a synthetic dataset, and shown that the absence of full ground truth distributions leads to highly overconfident unimodal predictions. We have provided a solution in the form of an added hinge, and demonstrated that this modification mitigates overconfidence, when training on datasets where full ground truth distributions are unavailable. In the future we are interested in addressing overconfidence also in higher dimensional regression tasks, and in combining RbC with ensemble methods.

Acknowledgement: This work was funded by Swedish national strategic research environment ELLIIT, grant C08, and partially by the Wallenberg Artificial Intelligence, Autonomous Systems, and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. Computational resources were provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS).

References

- [1] Oisín Mac Aodha, Ahmad Humayun, Marc Pollefeys, and Gabriel J. Brostow. Learning a confidence measure for optical flow. *IEEE Transactions on Pattern Recognition and Machine Intelligence (TPAMI)*, 35:1107–1120, 2013. [5](#)
- [2] Olga Barinova, Victor Lempitsky, Elena Tretyak, and Pushmeet Kohli. Geometric image parsing in man-made environments. In *Computer Vision—ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5–11, 2010, Proceedings, Part II 11*, pages 57–70. Springer, 2010. [6](#)
- [3] Eric Brachmann and Carsten Rother. Neural-guided RANSAC: Learning where to sample model hypotheses. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4322–4331, 2019. [3](#)
- [4] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5410–5418, 2018. [7](#)
- [5] Jianlin Cheng, Zheng Wang, and Gianluca Pollastri. A neural network approach to ordinal regression. In *2008 IEEE international joint conference on neural networks (IEEE world congress on computational intelligence)*, pages 1279–1284. IEEE, 2008. [2](#)
- [6] Per-Erik Forssén, Björn Johansson, and Gösta Granlund. Channel associative networks for multiple valued mappings. In *2nd International Cognitive Vision Workshop*, pages 4–11, Graz, Austria, 2006. [2](#), [3](#), [6](#)
- [7] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pages 1050–1059. PMLR, 2016. [2](#)
- [8] Divyansh Garg, Yan Wang, Bharath Hariharan, Mark Campbell, Kilian Weinberger, and Wei-Lun Chao. Wasserstein distances for stereo disparity estimation. In *NeurIPS*, 2020. [2](#), [3](#), [7](#), [8](#)
- [9] Divyansh Garg, Yan Wang, Bharath Hariharan, Mark Campbell, Kilian Q. Weinberger, and Wei-Lun Chao. Wasserstein distances for stereo disparity estimation. In *NeurIPS*, 2020. [4](#), [6](#), [7](#)
- [10] Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477), 2007. [5](#)
- [11] Tilmann Gneiting, Fadoua Balabdaoui, and Adrian E Raftery. Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 69(2):243–268, 2007. [5](#)
- [12] Gustav Häger, Mikael Persson, and Michael Felsberg. Predicting disparity distributions. In *IEEE International Conference on Robotics and Automation (ICRA’21)*, pages 4363–4369, 2021. [2](#), [3](#), [4](#), [5](#), [6](#)
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. [6](#)
- [14] 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. [2](#)
- [15] Georg Hess, Christoffer Petersson, and Lennart Svensson. Object detection as probabilistic set prediction. In *17th European Conference on Computer Vision (ECCV)*, pages 550–566. Springer, 2022. [2](#)
- [16] Eddy Ilg, Ozgun Cicek, Silvio Galesso, Aaron Klein, Osama Makansi, Frank Hutter, and Thomas Brox. Uncertainty estimates and multi-hypotheses networks for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 652–667, 2018. [2](#), [5](#)
- [17] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30, 2017. [1](#), [3](#)
- [18] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan, Kennedy Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *International Conference on Computer Vision (ICCV)*, pages 66–75, 2017. [4](#)
- [19] Peter Kotschieder, Madalina Fiterau, Antonio Criminisi, and Samuel Rota Buló. Deep neural decision forests. In *Proceedings of the IEEE international conference on computer vision*, pages 1467–1475, 2015. [2](#)
- [20] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017. [2](#), [7](#)
- [21] Jun-Tae Lee, Han-Ul Kim, Chul Lee, and Chang-Su Kim. Semantic line detection and its applications. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3229–3237, 2017. [3](#)
- [22] Xiaofeng Liu, Yang Zou, Tong Che, Peng Ding, Ping Jia, Jane You, and B.V.K. Vijaya Kumar. Conservative wasserstein training for pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. [2](#), [4](#)
- [23] Wei-Yin Loh. Classification and regression trees. *Wiley interdisciplinary reviews: data mining and knowledge discovery*, 1(1):14–23, 2011. [2](#)
- [24] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. In *4th International Conference on Learning Representations, ICLR 2016*, 2016. [1](#)
- [25] N. Mayer, E. Ilg, P. Häusser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. arXiv:1512.02134. [7](#)
- [26] Alireza Namdari and Zhaojun Li. A review of entropy measures for uncertainty quantification of stochastic processes. *Advances in Mechanical Engineering*, 11(6): 1687814019857350, 2019. [5](#)
- [27] Zhenxing Niu, Mo Zhou, Le Wang, Xinbo Gao, and Gang Hua. Ordinal regression with multiple output cnn for age estimation. In *Proceedings of the IEEE conference on*

- computer vision and pattern recognition*, pages 4920–4928, 2016. [2](#)
- [28] Nicolas Papernot and Patrick Mcdaniel. Deep k-nearest neighbors: Towards confident, interpretable and robust deep learning. *ArXiv*, abs/1803.04765, 2018. [2](#)
- [29] Lukasz Rudnicki. Shannon entropy as a measure of uncertainty in positions and momenta. *Journal of Russian Laser Research*, 32:393–399, 2011. [5](#)
- [30] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823, 2015. [4](#)
- [31] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. [3](#)
- [32] Thordis L. Thorarinsdottir, Tilmann Gneiting, and Nadine Gissibl. Using proper divergence functions to evaluate climate models. *SIAM/ASA Journal on Uncertainty Quantification*, 1(1), 2013. [4](#), [5](#)
- [33] M. Thorpe. Introduction to optimal transport. Lecture Notes, 2018. [4](#)
- [34] Scott Workman, Menghua Zhai, and Nathan Jacobs. Horizon lines in the wild. In *British Machine Vision Conference (BMVC)*, pages 20.1–20.12, 2016. Acceptance rate: 39.4%. [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [35] Fan Zhang and Lauren J. O’Donnell. Chapter 7 - support vector regression. In *Machine Learning*, pages 123–140. Academic Press, 2020. [2](#)