

Towards Multimodal Depth Estimation from Light Fields

Titus Leistner, Radek Mackowiak, Lynton Ardizzone, Ullrich Köthe, Carsten Rother
Visual Learning Lab, Heidelberg University

first.last@iwr.uni-heidelberg.de

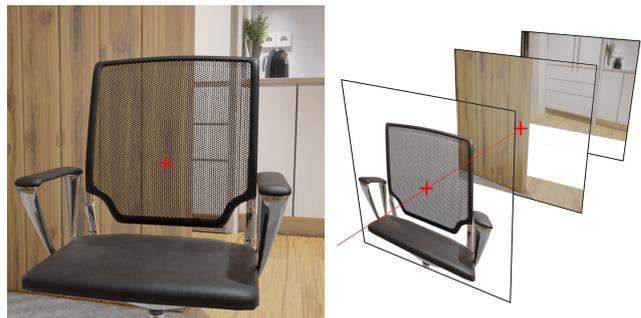
Abstract

Light field applications, especially light field rendering and depth estimation, developed rapidly in recent years. While state-of-the-art light field rendering methods handle semi-transparent and reflective objects well, depth estimation methods either ignore these cases altogether or only deliver a weak performance. We argue that this is due current methods only considering a single “true” depth, even when multiple objects at different depths contributed to the color of a single pixel. Based on the simple idea of outputting a posterior depth distribution instead of only a single estimate, we develop and explore several different deep-learning-based approaches to the problem. Additionally, we contribute the first “multimodal light field depth dataset” that contains the depths of all objects which contribute to the color of a pixel. This allows us to supervise the multimodal depth prediction and also validate all methods by measuring the KL divergence of the predicted posteriors. With our thorough analysis and novel dataset, we aim to start a new line of depth estimation research that overcomes some of the long-standing limitations of this field.

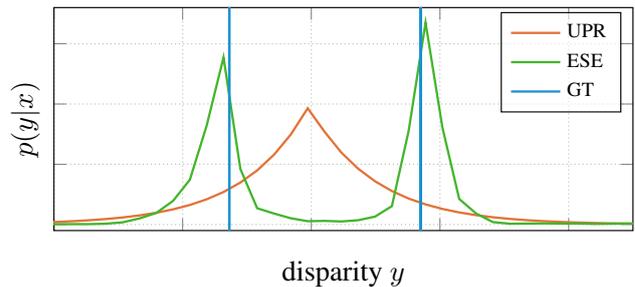
1. Introduction

Light field recordings and their applications, like real time rendering for virtual reality or highly accurate depth estimation, have improved vastly in recent years. However, while light field rendering methods handle transparent and reflective objects well, current depth estimation methods still perform poorly in those areas. State-of-the-art depth estimation methods mainly fail in three corner cases: at objects edges, semi-transparent and reflective surfaces. All three are caused by multiple objects at different depths contributing to the projected color of a single pixel on the camera sensor. Most existing models fundamentally ignore these cases and assume only one “true” depth for each pixel.

Instead, we propose a series of deep-learning based methods to perform multimodal depth estimation, and depth estimation with uncertainty estimates. For this, we start with the basic idea of outputting Bayesian posteriors,



(a) Rendered scene (left) containing multiple disparity layers (right)



(b) Disparity posterior distributions, predicted by different methods

Figure 1. **Comparison of disparity posterior distributions.** Synthetic scene (a) containing overlapping objects at different depths. (b) shows disparity posterior distributions, estimated by different methods, for a single pixel (red crosses in (a)). This pixel captured two disparity modes (mesh material of chair (foreground) and wooden wall (background)). Note, that the Unimodal Posterior Regression (UPR) network, which outputs the mean and width of a Laplacian distribution, makes a wrong and uncertain prediction. The EPI-Shift Ensemble (ESE) detects both valid modes near the ground truth (GT).

whereas standard regression models just produce a single estimate. From this idea, and by using a simple and well-founded maximum likelihood training framework, we develop three different light field depth estimation methods, all three of which are able to infer uncertainty estimates, and use multiple ground truth values during training. Two of the proposed methods are also able to predict multiple distinct depth values per pixel at test time.

To train our methods, we propose to utilize a multimodal dataset that contains the exact depth, color and opacity of all depth planes that are visible in an image. This is, in contrast to other current datasets which only contain a single “true” depth value. Our multimodal dataset consists of randomly generated synthetic scenes with a significant proportion of occlusion and transparent objects. This, for the first time, enables supervised training of multimodal depth estimation.

Our main contributions are as follows:

- An exploration of our three novel deep learning methods for light field based depth estimation being able to handle multiple depth modes: (i) Unimodal Posterior Regression (UPR); (ii) EPI-Shift-Ensemble (ESE); (iii) Discrete Posterior Prediction (DPP)
- The release of the first multimodal light field dataset, containing the depth of all objects and their contribution to the color of each pixel in an image
- A thorough evaluation of the predicted depth posterior distributions. We observe that the more restrictive UPR method works well when the unimodal depth assumption of traditional methods holds. In cases where it does not, the model is able to express a high uncertainty. In general, the discrete DPP method is superior to UPR and ESE.

2. Related Work

In this section, we first discuss the related methods for depth estimation from light fields. We secondly introduce different posterior prediction methods mainly inspired by works on uncertainty quantification.

2.1. Light Field and Multimodal Depth Estimation

Sinha *et al.* [20] separate albedo and reflections in multi-view stereo recordings. Their method discretizes the disparity space and utilizes a plane-sweep framework to compute an error volume based on pairwise normalized cross correlation. Up to two distinct layers are extracted from this volume using a modified semi global-matching algorithm.

Johannsen *et al.* [11] introduced a similar method for multimodal light field depth estimation. The method uses a dictionary of small EPI patches to encode light fields and estimate the depth at each pixel. In addition, a mask that separates unimodal from bimodal pixels is computed. Both aforementioned methods are able to estimate multiple depth modes from light fields but can not compete with the overall performance of state-of-the-art deep-learning based frameworks. We refer to Appendix D for a comparison.

Heber *et al.* [6] introduced the first deep learning based approach for light field depth estimation. A neural network predicts the slope of local 2D per-pixel hyper planes. In a

second optimization step, a disparity map is inferred from the hyperplane parameters. Due to noise and inaccurate results in untextured areas, an additional post-processing optimization is required. The approach is a first step in the direction of learned depth from lightfields, but the results suffer from artifacts and blur which are addressed in their follow-up work [7, 8].

Shin *et al.* [19] use a fully-convolutional architecture for direct disparity regression, consisting of two parts: A multi-stream network comprised of four input networks for the horizontal, vertical and both diagonal light field view stacks. The output features of those streams are concatenated and refined by a fully-convolutional head that directly outputs disparities. EPI-Net achieves state-of-the-art performance on the HCI 4D Light Field Benchmark [9] but is limited to a small disparity interval due to its small receptive field.

Leistner *et al.* [13] addressed the problem of light field depth estimation for high resolution and wide baseline light fields utilizing neural networks. Instead of increasing the receptive field of the network which would yield worse generalization, they proposed to transform the input images with a shear transformation, called EPI-Shift. The method applies a certain number of forward passes with shifted input EPIs and joins the result into a single prediction.

The most recent works by Tsai *et al.* [22] and Chen *et al.* [1] aim for a better utilization of the light field data. In contrast to previous methods, they utilize all light field views but fuse them in an early stage to reduce redundancy.

With Neural Radiance Fields (NeRF) Mildenhall *et al.* [15] started a new line of research that represents the plenoptic function of a given scene as a neural network. The network is trained using images recorded from different view points. At first glance, the radiance along a ray in the NeRF framework also models a depth posterior distribution. However, we argue that this assumption does not hold in general. Zhang *et al.* [23] observed an inherent shape-radiance ambiguity in NeRF models. This is caused by training with a photometric loss only instead of supervised learning using ground truth depth. As a consequence, there exists a family of radiance fields which perfectly explain all training images, even if the shape is incorrect.

Tosi *et al.* [21] addressed the problem of smoothness bias at depth discontinuities for stereo depth estimation methods. Their method predicts a bimodal Laplacian mixture distribution for each pixel and always chooses the mode with the largest probability density to retain sharp edges.

2.2. Prediction of Regression Posteriors

Traditional depth estimation methods predict a single depth value per pixel. In the following, we therefore highlight works, focused on other tasks, that estimate a whole posterior instead. MacKay [14] introduced the idea of Bayesian Neural Networks (BNNs) by proposing a prob-

abilistic interpretation of the training process. Instead of predicting only one output, the proposed network models the likelihood of possible outputs given the input. The posterior probability of network weights, given the dataset can be inferred by maximizing this likelihood. Now, Bayes rule is applied using maximum likelihood and some prior distribution over the network weights. However, for large networks, this cannot be computed efficiently.

Neal [16] introduced the Markov Chain Monte Carlo (MCMC) method to approximate the posterior over the model weights. In addition, he analyzed the importance of the chosen prior over a high number of model weights. This methods enabled an efficient approximation and therefore the training of BNNs. However, even recent MCMC approaches are only applicable to a limited number of dimensions.

Kendall and Gal [12] analyzed different types of uncertainty relevant for Computer Vision. They propose to capture aleatoric uncertainty (uncertainty inherent in the data) by training the network to predict a variance for its output. Epistemic uncertainty (uncertainty inherent in the model), which can be explained away using an infinite amount of data, is inferred using the Monte-Carlo Dropout technique [3]. The authors propose a single model to infer both types of uncertainties.

Ilg *et al.* [10] compared different uncertainty quantification methods for optical flow. They compare aleatoric uncertainty learned by a single network to an ensemble of such networks and further, to a single network with multiple heads trained by the Winner-Takes-it-All loss. In comparison, the authors find that Monte-Carlo Dropout does not work well for regression tasks and that ensembling also does not provide significantly better uncertainties than single networks. However, we argue that this is caused by the lack of epistemic uncertainty for optical flow, due to the utilized synthetic datasets being quite large. As shown by [12], Monte Carlo dropout ensembles capture epistemic uncertainty, which can be explained away given enough data. Because epistemic uncertainty should be close to zero for huge synthetic datasets as used by [10], there is no room for improvement over a single network in the first place.

3. Method

Before we introduce our posterior estimation methods, we explain the basics of light field depth estimation. A 4D light field is typically recorded by a 2D array of cameras, aligned on a regular grid. Each camera is assigned a pair of so-called view coordinates (u, v) , *e.g.* $(1, 1)$ for the top left camera. Each pixel in an image is then also assigned a pair of image coordinates (s, t) . A slice along one axis in image space with the corresponding axis in view space, *e.g.* fixing u and s , forms a so-called Epipolar Plane Image (EPI). Each pixel in the central camera view is visible as a line structure

in the EPI. The slope of this line is equal to the negative inverse disparity of the pixel. Thus, the task of depth estimation from light fields is to robustly detect this slope. However, this approach assumes opaque, smooth and lambertian surfaces. Non-textured, specular or semi-transparent regions and depth edges are ambiguous and therefore challenging even for state-of-the-art methods. In this work we make progress for these cases by estimating the full depth posterior distribution. This is especially useful for pixels with more than one valid depth mode, caused by either semi-transparency or the point spread at depth edges. Unlike previous works that only predict a single depth, we are able to find those modes.

3.1. Posterior Estimation

From here on, we will simply refer to the input of a depth estimation network as x , and to disparity as y . In practice, the input to a depth estimation network is a concatenation of horizontal, vertical and diagonal EPIs. Standard regression models usually output a single guess for the disparity, $\hat{y} = f_w(x)$, where f_w could be EPI-Net [19] with network weights denoted as w . Instead, our goal is to estimate the posterior distribution $p(y | x)$ of the disparity y given an input light field x . In the following, we present four different approaches that all model such a posterior distribution.

To supervise more complex posterior distributions that can represent more than one mode, we created our own multimodal depth dataset. Unlike common datasets that only contain a single ground truth disparity y_i for each pixel i , we include the disparities of multiple depth modes y_{ij} for transparent objects and depth edges. For each disparity we also include the amount of color η_{ij} that it contributed to the pixel, *i.e.* the perceived opacity of that object in the pixel. From a Bayesian perspective, we interpret η_{ij} as the probability $p(y_{ij})$ of this disparity. This choice is justified in both an intuitive and methodological sense: η_{ij} corresponds to the fraction of the pixel’s area that is taken up by the object at depth y_{ij} . Lacking any prior knowledge, this is also equal to the probability that the depth at any subpixel position corresponds to that object. This equality between the opacity η_{ij} and probability $p(y_{ij})$ is valid both at edges as well as fine structures such as grids or woven meshes. We also extend the definition to apply to semi-transparent materials such as printed glass as a simplifying assumption. In Appendix B, we give a more rigorous examination of the connection between η_{ij} and $p(y_{ij})$.

3.1.1 Unimodal Posterior Regression

The most common approach for learning distributions is Maximum-Likelihood (ML) learning, which most loss functions can be reformulated as. The ML objective aims to find the model parameters w which maximize the log likeli-

hood of the training data $\{(x_i, y_i)\}_{i=0}^N$ under the estimated posterior distribution. In practice, we minimize the *negative* log likelihood instead:

$$\mathcal{L}_{\text{ML}} = -\frac{1}{N} \sum_i \log p(y_i | x_i, w). \quad (1)$$

It can be shown that this objective minimizes the Kullback-Leibler divergence (KLD) between $p(y | x, w)$ and the true posterior $p^*(y | x)$.

Previous regression based approaches [19] simply use the \mathcal{L}_1 loss to make a single prediction:

$$\mathcal{L}_1 = \frac{1}{N} \sum_i |y_i - f_w(x_i)|. \quad (2)$$

We see that this is equal to the ML objective when the posterior is assumed to be a Laplace distribution $p(y | x, w) \propto \exp(-|y - \mu|/b)/2b$ with the network output $\mu = f_w(x)$ and a fixed value of $b = 1$. This motivates the following simple extension, which is an adaptation of the Dawid-Sebastiani-score [2], later popularized in [12], except that the \mathcal{L}_2 -loss corresponding to a Gaussian posterior was used instead: we allow the network to change the width b of the posterior. With this, it becomes

$$p(y | x, w) = \frac{1}{2b} \exp\left(-\frac{|y - \mu|}{b}\right), \text{ with } [b, \mu] = f_w(x). \quad (3)$$

Putting this back into the ML objective, we get the following loss function for the predictive uncertainty:

$$\mathcal{L}_{\text{UPR}} = \frac{1}{N} \sum_i \frac{|\mu_i - y_i|}{b_i} + \log b_i, \text{ with } [\mu_i, b_i] = f_w(x_i). \quad (4)$$

This loss can be understood intuitively: If the network struggles to predict y_i , the \mathcal{L}_1 loss term can be down-weighted by increasing the scale parameter b for this pixel. To avoid the trivial solution $b_i \rightarrow \infty$ for any input, high b are penalized by a regularization term $\log b$. In practice, we let the network predict $\log b$ instead of b to improve numerical stability.

This approach gives us a measure of aleatoric uncertainty [12] for each pixel which is already helpful for many downstream applications. However, the implicit assumption for the method to work well is that the true posterior is also Laplacian. Needless to say, this is certainly not true in multi-modal cases, which cannot be modeled by the Laplace distribution. With multiple ground-truth depth modes y_{ij} as opposed to only y_i , as in our dataset, the loss for a Laplace distribution becomes

$$\mathcal{L}_{\text{UPR}}^{\text{MM}} = \frac{1}{N} \sum_i \sum_j p(y_{ij}) \frac{|y_{ij} - f_w(x_i)|}{\log b_i} + \log b_i \quad (5)$$

and can be applied to the \mathcal{L}_1 loss respectively:

$$\mathcal{L}_1^{\text{MM}} = \frac{1}{N} \sum_i \sum_j p(y_{ij}) |y_{ij} - f_w(x_i)| \quad (6)$$

However, in any case, those networks will focus on a single mode, or lie in between, and compensate for its wrong prediction by expressing a very high uncertainty like in Fig. 1.

3.1.2 EPI-Shift-Ensemble

Commonly, one way to circumvent this exact issue is to use an ensemble of networks. Instead of just estimating a single posterior, M networks predict M different posteriors, which are then averaged:

$$p(y|x) = \frac{1}{M} \sum_k p(y|x, w_k) \quad (7)$$

over all networks with learned weights w_k , $k = 1 \dots M$. It has been shown that ensembles deliver some of the best uncertainty estimates among existing methods [17]. Their main limitation is the high computational cost, especially for training. Various approaches try to avoid this and train only a single model. For instance, Monte Carlo dropout exhibits similar characteristics as a true ensemble [5]. Instead, we propose a new scheme which uniquely exploits the nature of light field data, which we term EPI-Shift-Ensemble (ESE). To motivate this, we take note of the technique from [13], where the EPI is sheared in such a way that a global offset Δy is added to the disparity of the EPI. This is advantageous, because it enables inference on large disparities *e.g.* for wide-baseline light field cameras. We denote this operation as *shift*($x, \Delta y$) and extend it to arbitrary sub-pixel steps. To form our EPI-Shift-Ensemble, we successively shift the input M times in steps of Δy . This gives us M different augmented inputs, where each has a different artificial disparity offset. Each input is fed into the network, and the resulting posteriors are shifted back by the same offset. In this way, we get M different estimates by the same network, each from a different input containing the same information. The overall posterior can be expressed as

$$p_{\text{ESE}}(y | x, w) = \frac{1}{M} \sum_k p\left(y - k \cdot \Delta y \mid \text{shift}(x, k \cdot \Delta y), w\right). \quad (8)$$

In the summation, $k = -\lfloor M/2 \rfloor \dots \lfloor M/2 \rfloor$. For the shape of the individual posteriors, we use the same Laplacian as before. Note that the network weights w are shared for each summand, only the shift is different. However, this operation alone does not prevent the problems seen for the single mode approach: the network will try to average out bi-modal solutions, or collapse into one mode. As a result, we see little to no diversity in the EPI-Shift-Ensemble, and no

multi-modal posteriors. To prevent the collapse, we mask the loss during training so that it only applies to pixels with a shifted $y' = y - k \cdot \Delta y$ within the range of one step $|y'| < \Delta y/2$. In all other cases the output will have a large uncertainty:

$$\mathcal{L}_{\text{ESE}} = \frac{1}{N} \sum_i \begin{cases} \frac{|\mu_i - y_i|}{b_i} + \log b_i & \text{if } |y'_i| < \frac{\Delta y}{2} \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

We extend this to our multimodal dataset similarly to our unimodal networks:

$$\mathcal{L}_{\text{ESE}}^{\text{MM}} = \frac{1}{N} \sum_i \sum_j p(y_{ij}) \begin{cases} \frac{|\mu_i - y_{ij}|}{b_i} + \log b_i & \text{if } |y'_{ij}| < \frac{\Delta y}{2} \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

After training with this loss, the network will only be confident (narrow posterior) if it estimates that the input disparity is $y \in [-\Delta y/2, \Delta y/2]$, and the predicted posterior will always be centered in this range. When using this network in the EPI-Shift-Ensemble, we see that each term k will only contribute a narrow posterior if a plausible disparity lies between $(k - 1/2)\Delta y$ and $(k + 1/2)\Delta y$, thus ensuring diverse outputs and the possibility of multi-modal predictions. Three details should be noted: First, the model is not trained as an ensemble, a single model is trained just as before with the modified \mathcal{L}_{ESE} -loss. The ensembling operation is only performed at inference time. Second, the masked loss does not reduce the effective size of our training set, as we also apply random EPI-shifts as a part of the data augmentation process. This way, all pixels will randomly fulfill $|y_i| < \Delta y/2$ at some point. Lastly, the inference time is M times longer, as M forward passes have to be performed to compute the EPI-Shift-Ensemble.

3.1.3 Discrete Posterior Prediction

The approach of discretizing regression tasks has been successful for stereo depth estimation in the past, and promises to model more expressive posteriors. Specifically, by discretizing the range of disparities, a softmax output can be used to represent the posterior. The posterior is then a step function consisting of these discrete probabilities. If y_j are the discretization steps, we can write

$$p(y_j|x) \propto \text{softmax}(f_w(x_i))_j := \frac{\exp(f_w(x_i)_j)}{\sum_k \exp(f_w(x_i)_k)} \quad (11)$$

If multiple modes are used for the training, we also discretize the distribution $p(y_j)$ over these modes. Maximum likelihood training is then simply equivalent to the categorical cross-entropy (CE) loss, where the correct ‘‘class’’ is the bin j that the training example y_i lands in:

$$\mathcal{L}_{\text{CE}} = \frac{1}{N} \sum_i -\log(\text{softmax}(f_w(x_i))_j). \quad (12)$$

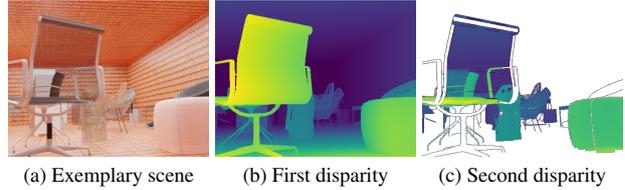


Figure 2. An exemplary scene (a) from our randomly generated dataset. (b) and (c) show different disparity modes. Each pixel has at least one disparity mode. Behind semi-transparent objects and at depth edges, a second disparity mode (c) exists.

For the multimodal dataset, we can compute the cross-entropy accordingly:

$$\mathcal{L}_{\text{CE}}^{\text{MM}} = \frac{1}{N} \sum_i \sum_j -p(y_j) \log(\text{softmax}(f_w(x_i))_j). \quad (13)$$

Note, that the two correspond, for a unimodal dataset, we simply have $p(y_j) = 1$ and $p(y_{l \neq j}) = 0$, simplifying to the cross-entropy in Eq. 12.

Although the discrete posterior prediction gives us an uncertainty estimation at much lower computational cost than the EPI-Shift-Ensemble and can represent more flexible posteriors compared to simple Laplacians, softmax probabilities are generally known to be overconfident [4] and also make wrong but confident predictions in ambiguous or unseen cases. Different techniques for post-calibration of uncertainties exist, outlined in [4]. However, while they may prevent overconfidence going from the training to a test set, they do not make the uncertainties more reliable in general [17], e.g. for ambiguous inputs.

3.2. Dataset Generation

To train and validate all methods above, ground truth multimodal depth data is required. Because all previous light field datasets only contain a single depth per pixel, we generated a novel multimodal depth light field dataset containing 110 randomly generated indoor scenes (see Fig. 2). To improve the training performance, the dataset generator follows four goals: (i) relatively photorealistic appearance, (ii) high diversity to improve generalization of trained models, (iii) many occlusions and depth edges to improve the performance at object edges and, (iv) a large proportion of pixels with multiple valid depths. Note, that we decided to not include reflections explicitly. However, as [20] demonstrates, reflections at flat surfaces behave similar to alpha-transparency. Therefore, the methods trained with our dataset should also be usable to estimate multimodal depths caused by reflections. We refer to Appendix C for more information on the dataset generation process.

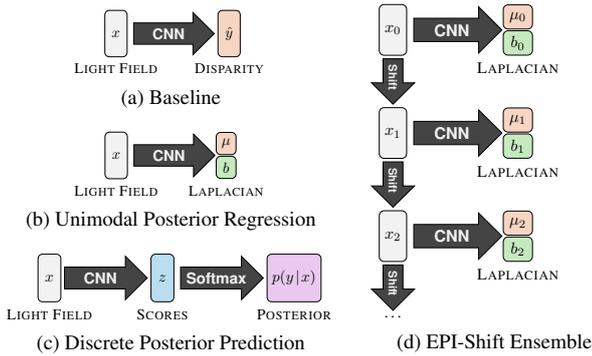


Figure 3. **Network architecture overview:** Our baseline is a simple feed-forward network trained to only predict disparity (a). We compare three posterior-regression methods: Laplacian prediction using learned loss attenuation (b), an ensemble on shifted inputs (d) and a discrete softmax classification network (c).

4. Experiments

In this section, we describe our training procedure and analyze the predicted posterior distributions. We therefore distinguish two possible applications: unimodal prediction with uncertainty and multimodal prediction.

4.1. Architectures and Training Details

To ensure a fair comparison, we chose the state-of-the-art method EPI-Net [19] with minor modifications as our backbone network architecture for all models. We describe the exact network architectures in more detail in Appendix A. The network consists of a total number of eight blocks, containing two convolutional layers, followed by a Rectified Linear Unit (ReLU) each and one Batch Normalization layer. For UPR, we added an additional output layer to also predict the negative log width $-\log b$. The EPI-Shift Ensemble is a modification of UPR with the only difference being the loss functions. We experimented with different Δy for our ESE method, resulting in different M s. $\Delta y = 0.1$ performed best and gives $M = 70$ for the disparity range of our data. Our extended EPI-Shift transformation, based on [13], is described in more detail in Appendix A.1. For the discrete DPP method, we chose a number of 108 “classes”. This is motivated by the common BadPix007 metric that considers a pixel as “correct” if it is closer than $0.07px$ to the ground truth. A number of 108 classes in a disparity range of $-3.5 \dots 3.5$ leads to a bin size of ≈ 0.065 which is slightly below this threshold. We implemented our framework using PyTorch [18].

We trained all networks using the loss functions described in Sec. 3.1. The unimodal loss functions, denoted as \mathcal{L}_x are always applied to the closest disparity. All multimodal loss functions, denoted as \mathcal{L}_x^{MM} are applied to all disparity modes.

In addition, we reimplemented EPI-Net [19] in our own

framework as a baseline for a fair comparison. The learning rate was set to 10^{-3} for all models, using a batch size of 512 and the Adam optimizer. We trained on randomly cropped patches ($96px \times 96px$) from the set of 100 training scenes in our dataset. To further improve the diversity of our dataset, we make use of a number of data augmentation operations: We apply a random sub-pixel EPI-Shift between $-2px$ and $2px$ [13]. In addition, we randomly rotate the light field by a multiple of 90° , randomly rotate the colors in RGB-space and randomly change brightness and contrast.

4.2. Posterior Evaluation

Depth estimation methods are usually evaluated by measuring the pixel-wise error to a ground truth disparity map. However, to correctly measure the quality of estimates in areas with multiple valid depths, a different set of metrics is required.

We consider two application scenarios: (i) the estimation of just a single disparity, but with an additional confidence measure to ensure that the estimate can be “trusted”. This may be required, *e.g.* in industrial and robotics applications where decisions are based on the estimated depth; (ii) the estimation of multiple depths in areas with transparent objects or at object edges. This is typically required by Computer Graphics applications that aim to render the recorded scene from a different angle. In the following, we introduce metrics for both cases.

4.2.1 Unimodal Prediction with Uncertainty

Previous methods and datasets [9] always consider the disparity of the closest object as “true”, even when this object is transparent. However, estimation methods oftentimes output the disparity of the background object in those cases, which may lead to severe issues in downstream applications. In addition, more ambiguities usually occur in non-textured areas which cannot be estimated correctly. Most optimization-based methods fill in those ambiguous regions by interpolation between adjacent pixels with confident predictions. Due to the limited receptive field of neural network based methods, this is only possible to some extent. In any case, a confidence measure is extremely useful for downstream applications in order to decide if an estimate can be trusted. To achieve this, the overall variance σ^2 of the predicted posterior distribution can be used as an uncertainty measure. We aim for a consistently high uncertainty in regions with ambiguous predictions. To evaluate the quality of the estimated posteriors, we remove the $s\%$ of pixels with the highest posterior variance (uncertainty). With these ambiguous and uncertain cases filtered out, the BadPix of the remaining pixels is lower, which can be plotted as a sparsification curve (see Fig. D.3a). The optimal curve can be computed by removing those $s\%$ of pixels

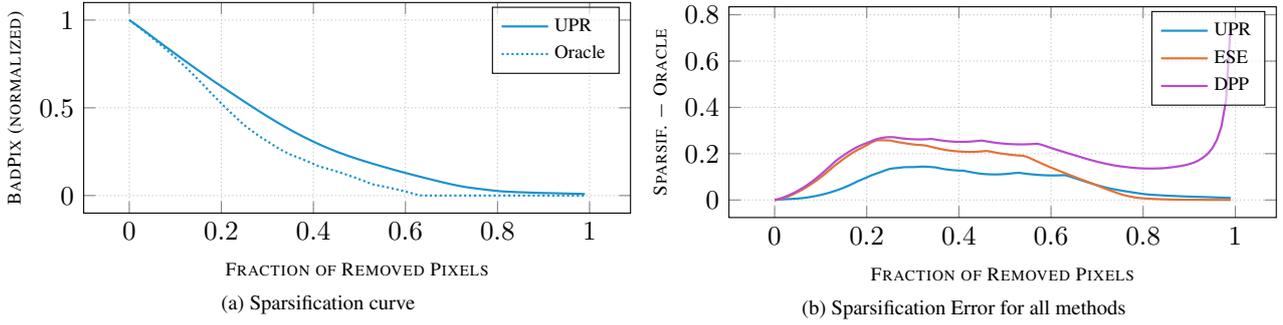


Figure 4. **Unimodal uncertainty quantification:** Sparsification results of analyzed methods with respect to the disparity BadPix007. By removing a certain fraction with the highest predicted uncertainty (a), the error decreases. The “Oracle” is a lower bound, created by removal of truly worst pixels. We also compute the difference between the predicted sparsification and its “Oracle”, denoted as Sparsification Error (SE) (b), for a comparison of the three methods (trained on all depth modes)

with the largest ground truth error. We call this the “oracle” curve, which represents the lower bound of what is achievable. This method is commonly used to evaluate uncertainties for regression tasks [10]. In order to compare all methods, we compute the Sparsification Error (SE) by subtracting the oracle curve from the sparsification curve. The Area under the Sparsification Error (AuSE) quantifies the uncertainty quality of each method with a single number.

4.2.2 Multimodal Prediction

For many applications, including re-rendering of a recorded scene from different angles, estimating multiple depths for pixels at object edges and transparent surfaces is desirable. In addition to the disparities of all modes, the contribution of each mode to the color of the pixel is also important. To evaluate both, we measure the Kullback-Leibler divergence (KLD)

$$\mathcal{D}_{\text{KL}} = \int p(y_i) \log \left(\frac{p(y_i)}{p(y_i|x_i)} \right) \quad (14)$$

between the predicted disparity posterior $p(y_i|x_i)$ and the true disparity distribution $p(y_i)$ at a pixel i . Intuitively, the KLD will be minimal if the posterior assigns a high probability density to each true disparity mode and a low density to disparities that are not present at a pixel i . In other words: the KLD is lowest, if each mode is recovered sharply, instead of a broad uncertainty covering both. E.g. in Fig. 1b, the green and orange curves have the same uncertainty, but the green has a much lower KLD. Optimally, the density at each disparity mode corresponds to the contribution η_i of this mode to the resulting pixel color. As the KLD is only well defined between two continuous or two discrete distributions and we compare continuous as well as discrete methods, we chose to discretize all distributions. We therefore assign each ground truth disparity to one out of

$K = 108$ bins with a width of $h \approx 0.065\text{px}$ each.

$$p(y_k) = \sum_j p(y_j) \forall j \text{ with } |y_j - y_k| < \frac{h}{2} \quad (15)$$

This is, again, motivated by the well-established BadPix007 metric [9] which considers a pixel as correct if the \mathcal{L}_1 distance to the ground truth is below 0.07px . For the baseline method that only outputs one disparity, we simply set the probability of the bin that contains this disparity to 1. All continuous posterior distributions are discretized by integrating over the interval of each bin:

$$p(y_k|x) = \int_{(k-0.5)h}^{(k+0.5)h} p(y|x) dy \quad (16)$$

We now average the discrete KLD

$$\mathcal{D}_{\text{KL}} = \frac{1}{NK} \sum_i \sum_k p(y_{ik}) \log \left(\frac{p(y_{ik})}{p(y_{ik}|x_i)} \right) \quad (17)$$

over all pixels in all validation scenes. In addition, we also compute the KLD over all unimodal and all multimodal pixels separately. A pixel is considered multimodal if it has at least two modes j with $p(y_j) > 0.3$.

4.2.3 Results

Table 1 compares the unimodal, multimodal and sparsification performance of all methods. In the following, we will interpret our results with respect to the aforementioned applications: unimodal disparity estimation with uncertainty and multimodal disparity estimation.

With respect to pure unimodal performance, our baseline method and DPP perform best. The higher MSE for DPP is caused by small discretization errors due to the discrete number of bins. Those small errors are well below

Method	Unimodal Metrics		KL Divergence			AuSE ↓	Time ↓ (in sec)
	MSE ↓	BadPix ↓	Unimodal ↓	Multimodal ↓	Overall ↓		
BASE (uni)	0.374	0.229	4.720	7.876	5.421	-	2.188
BASE (multi)	0.563	0.307	5.259	8.514	6.025	-	2.211
UPR (uni)	0.439	0.235	1.719	3.381	1.879	0.071	2.260
UPR (multi)	0.676	0.285	1.987	3.156	2.114	0.072	2.287
ESE (uni)	1.269	0.223	4.164	3.628	4.160	0.099	17.492
ESE (multi)	1.850	0.229	4.283	3.719	4.277	0.121	16.902
DPP (uni)	0.765	0.209	1.631	3.057	1.734	0.272	4.348
DPP (multi)	0.686	0.231	1.824	2.987	1.914	0.197	4.382

Table 1. **Evaluation**, from left to right: Mean Squared Error and the common BadPix007 score (percentage of pixels with $|y_i - \hat{y}_i| > 0.07$), Kullback-Leibler divergence on unimodal, multimodal and all pixels, Area under Sparsification Error (AuSE), runtime of one forward pass. Our methods were trained with both losses \mathcal{L}_x (uni) and $\mathcal{L}_x^{\text{MM}}$ (multi) respectively. Lower is better

a threshold of 0.07px and therefore ignored by the BadPix metric which shows that DPP indeed predicts 2% more pixels correctly compared to the baseline. UPR performs only slightly worse than both methods overall. However, due to the uncertainty being directly supervised by \mathcal{L}_{UPR} this method is superior in terms of sparsification. This means that its uncertainty metric reflects most accurately whether a prediction is correct. We conclude that, if the application requires only a single disparity and confidence is important, UPR should be considered.

With respect to the accuracy of the predicted posterior distribution, DPP performs best in unimodal and also multimodal areas. However, as most softmax prediction methods, it is overconfident, as reflected by the sparsification error in Fig. D.3b. Despite the popularity of ensemble-based models for uncertainty estimation, ESE cannot compete with the other two methods. It especially performs worse in highly non-textured areas with a large and noisy uncertainty output b_i for all “ensemble members”. To decide for a unimodal disparity estimate, the member with the lowest uncertainty is still used, but due to the noise, this choice becomes arbitrary. The chosen member can only output a disparity from a small interval within the potentially large disparity range. On average, this causes a high deviation from the ground truth for ESE. In contrast, our other methods are able to “smooth” the disparity in uncertain areas from adjacent, more certain pixels. This is also reflected by the relatively high MSE but low BadPix error: the amount of “correctly” predicted pixels is on par with other methods, but the deviation of “wrongly” predicted pixels is generally higher. In addition, as all members contribute slightly to the mixture of Laplacians, the density of the posterior is higher along the whole disparity interval, which leads to a worse multimodal KLD compared to UPR and DPP. We therefore generally recommend DPP for multimodal predictions in small, narrow-baseline light fields. However, due to its shift-operation, ESE can, unlike other

methods, operate on arbitrary large disparity ranges and is therefore still advisable for high-resolution or wide-baseline light field cameras. In addition, it performs also relatively well in terms of sparsification. Comparing the methods trained on only one mode with the same methods trained on multiple modes shows that multimodal training leads to a slightly better multimodal performance for UPR and DPP, but always comes at a cost in unimodal areas. Our baseline method cannot efficiently represent multimodal posteriors as it only predicts a single disparity.

We conclude that the exact model and training method should be carefully chosen based on the intended application. We refer to Appendix D for examples of estimated disparity posterior distributions, an evaluation of our methods on the common HCI 4D Light Field Dataset [9], additional qualitative results as well as a comparison to Sinha *et al.* [20] and Johannsen *et al.* [11].

5. Conclusion

To summarize: we investigated the problem of posterior estimation for dense regression on the exemplary task of depth estimation from light fields. We therefore contributed the first light field dataset with multimodal depth ground truth. Additionally, we introduced and compared novel approaches for multimodal light field depth estimation, building on common uncertainty quantification tools. We observe that methods assuming a single valid depth work well if this assumption holds. DPP, which predicts arbitrary posterior distributions, works best in general. Our ESE method does not achieve the same performance, but estimates accurate confidence measures even for wide-baseline light fields. We hope that our insights lay the foundations for a new line of depth estimation research that overcomes some longstanding limitations of the field.

Acknowledgement: We thank the Center for Information Services and High Performance Computing (ZIH) at TU Dresden for generous allocations of computer time.

References

- [1] JX Chen, Shuo Zhang, and YF Lin. Attention-based multi-level fusion network for light field depth estimation. In *Proc AAAI Conf Artif Intell*, volume 35, pages 1009–1017, 2021. [2](#)
- [2] A Philip Dawid and Paola Sebastiani. Coherent dispersion criteria for optimal experimental design. *Annals of Statistics*, pages 65–81, 1999. [4](#)
- [3] 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. [3](#)
- [4] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR, 2017. [5](#)
- [5] Kazuyuki Hara, Daisuke Saitoh, and Hayaru Shouno. Analysis of dropout learning regarded as ensemble learning. In *International Conference on Artificial Neural Networks*, pages 72–79. Springer, 2016. [4](#)
- [6] Stefan Heber and Thomas Pock. Convolutional networks for shape from light field. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3746–3754, 2016. [2](#)
- [7] Stefan Heber, Wei Yu, and Thomas Pock. U-shaped networks for shape from light field. In *BMVC*, volume 3, page 5, 2016. [2](#)
- [8] Stefan Heber, Wei Yu, and Thomas Pock. Neural epi-volume networks for shape from light field. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2252–2260, 2017. [2](#)
- [9] Katrin Honauer, Ole Johannsen, Daniel Kondermann, and Bastian Goldluecke. A dataset and evaluation methodology for depth estimation on 4d light fields. In *Asian Conference on Computer Vision*, pages 19–34. Springer, 2016. [2](#), [6](#), [7](#), [8](#), [12](#), [15](#)
- [10] 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. [3](#), [7](#)
- [11] Ole Johannsen, Antonin Sulc, and Bastian Goldluecke. What sparse light field coding reveals about scene structure. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3262–3270, 2016. [2](#), [8](#), [12](#), [13](#)
- [12] 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. [3](#), [4](#)
- [13] Titus Leistner, Hendrik Schilling, Radek Mackowiak, Stefan Gumhold, and Carsten Rother. Learning to think outside the box: Wide-baseline light field depth estimation with epi-shift. In *2019 International Conference on 3D Vision (3DV)*, pages 249–257. IEEE, 2019. [2](#), [4](#), [6](#), [11](#)
- [14] David JC MacKay. A practical bayesian framework for back-propagation networks. *Neural computation*, 4(3):448–472, 1992. [2](#)
- [15] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020. [2](#)
- [16] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012. [3](#)
- [17] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, 32, 2019. [4](#), [5](#)
- [18] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. [6](#)
- [19] Changha Shin, Hae-Gon Jeon, Youngjin Yoon, In So Kweon, and Seon Joo Kim. EpiNet: A fully-convolutional neural network using epipolar geometry for depth from light field images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4748–4757, 2018. [2](#), [3](#), [4](#), [6](#), [10](#), [15](#)
- [20] Sudipta N Sinha, Johannes Kopf, Michael Goesele, Daniel Scharstein, and Richard Szeliski. Image-based rendering for scenes with reflections. *ACM Transactions on Graphics (TOG)*, 31(4):1–10, 2012. [2](#), [5](#), [8](#), [12](#), [13](#)
- [21] Fabio Tosi, Yiyi Liao, Carolin Schmitt, and Andreas Geiger. Smd-nets: Stereo mixture density networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8942–8952, 2021. [2](#)
- [22] Yu-Ju Tsai, Yu-Lun Liu, Ming Ouhyoung, and Yung-Yu Chuang. Attention-based view selection networks for light-field disparity estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12095–12103, 2020. [2](#)
- [23] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. [2](#)