Uncertainty of Optical flow: Regression calibration insights

Xuanlong Yu[1], Gianni Franchi[2], Emanuel Aldea[1]
1: SATIE, Université Paris-Saclay, 2: ENSTA Paris, Institut polytechnique de Paris
xuanlong.yu,emanuel.aldea @universite-paris-saclay.fr; gianni.franchi@ensta-paris.fr

Context and objectives

Reliable optical flow is crucial for many vision tasks related to navigation or planning. Deep neural networks achieve state-of-the-art performance on the problem of optical flow estimation. We would like to have a tool to indicate the reliability of optical flow which is the uncertainty estimator for optical flow.

- In this work we aim to quantify epistemic uncertainty for optical flow by using existing approaches.
- We use two different evaluation tools to measure the performance of uncertainty estimators, one is for the uncertainty ordering, one is for regression calibration degree.
- Our work gives insights on optical flow uncertainty calibration problem and the trade-offs which occur depending on the considered uncertainty estimation frameworks.

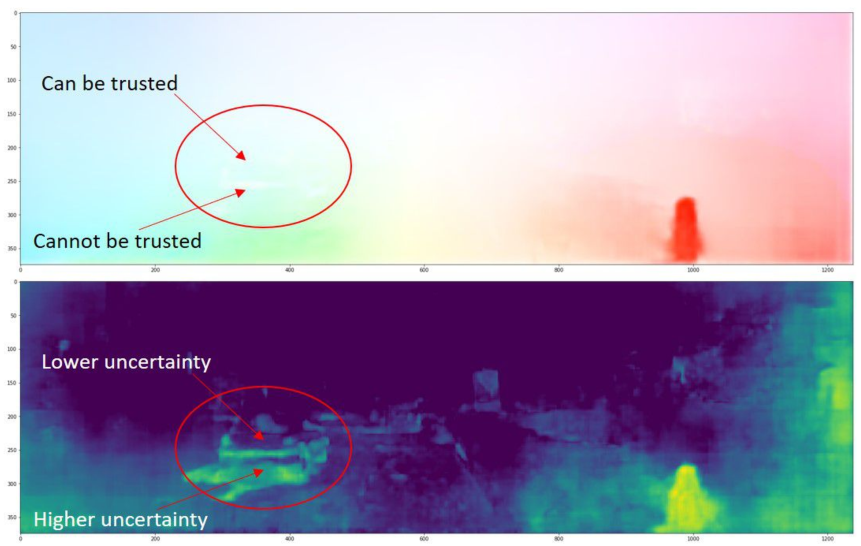

Fig. 1: Demonstration of optical flow prediction and one of its uncertainty estimation. We use FlowNetS [Dos+15] trained on the FlyingChairs [Dos+15] dataset and tested on KITTI [MG15] training set. The uncertainty map is made by MC Dropout [GG16] as an example.

Deep neural network and uncertainty estimation

- **MC Dropout** [GG16]: aim to find the posterior distribution of the parameters given the training dataset $P(\Theta \mid \mathcal{D})$, not only the values corresponding to the MAP. To make a prediction y on a new sample $\mathbf{x}$ the BNN compute : $P(y \mid \mathbf{x}, \mathcal{D}) = \int P(y \mid \mathbf{x}, \Theta) P(\Theta \mid \mathcal{D}) d\Theta$. We set Dropout layers in the model and do multiple forward propagations and calculate the variance among the inferences.
- **Deep Ensembles**[LPB17]: Train multiple deep neural networks to have access to their uncertainty. Each of the networks optimizes a heteroscedastic uncertainty loss [KG17] which considers that the output of optical flow model is a Gaussian distribution, the mean is the flow prediction and the variance is the uncertainty. Finally we combine the variances to get the final uncertainty.
- **Kullback-Leibler divergence**: The ground-truth that we want to match is a Gaussian distribution of mean the optical flow and of variance the square error of the optical flow. This loss aims to see if KL is more suitable than the previous heteroscedastic loss used in deep ensembles.
- **L2 regression**: a simple strategy using regression with means squared error loss. The trained variance is targeting the square error of the flow prediction.

Uncertainty evaluation methods

- **Sparsification Curve** [Ilg+18]: is a kind of Accuracy-Rejection Curve. It is a relative quality indicator: it reports the correctness of relative uncertainty ordering of the observed pixels with respect to the ideal ordering.
- **Calibration Curve**[KFE18]: has not yet been employed for optical flow analysis, is an absolute quality indicator, and it underlines the correlation between the expected confidence level taking into account the ground-truth and the confidence level reported by the variance estimator. The uncertainty estimation should be wider enough to cover the error zone confidently but also as sharp as possible.

Sequential training for uncertainty

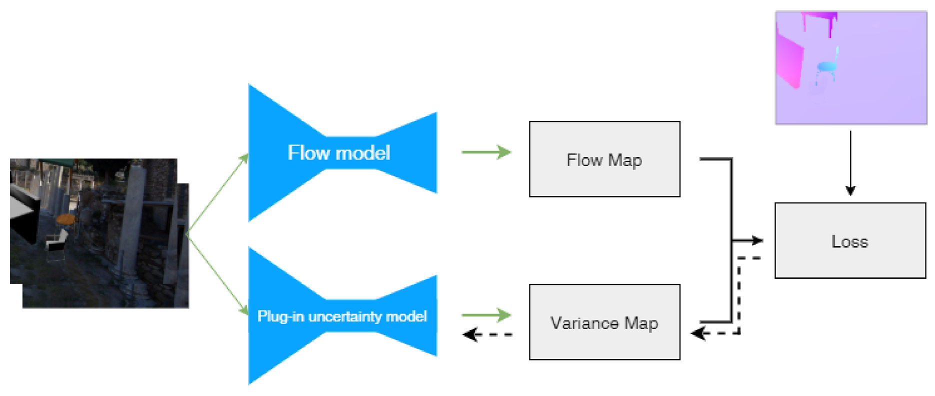

Fig. 2: We use a pre-trained and frozen optical flow model to make flow estimations and feed them to the loss to train the uncertainty estimator. Solid line: forward propagate; Dotted line: backward propagate; Black line: Procedure only during training; Green line: Procedure during inferring and training.

Experimental results

We train a FlowNetS [Dos+15] on FlyingChairs [Dos+15] for its flow and uncertainty predictions in combining with different uncertainty estimation approaches (MC: MC Dropout; $L2$: L2 regression targeting square error; KL: KL-divergence; EDE_i/PDE_i: empirical/predictive deep ensembles with i samples) by using sequential training. Then we evaluate the performance of different approaches on KITTI[MG15] without fine-tuning on it.

Curve Type	Calibration Error				Sparsification Error				Run time (ms)
Noise Type	Gaussian Noise		Motion Blur		Gaussian Noise		Motion Blur		
Method Name	AUC	RC	AUC	RC	AUC	RC	AUC	RC	
MC	1.0191	0.0171	0.9763	0.0206	**0.5102**	0.0406	**0.4232**	**0.0137**	196
$L2$	**0.8158**	**0.0059**	**0.8241**	**0.0122**	1.5801	0.2759	1.7393	0.3513	32
KL	1.0139	0.0148	1.0139	0.0327	0.9577	**0.0398**	0.8051	0.0696	32
PDE_1	0.9719	0.0111	0.9559	0.0150	1.4238	0.0536	1.2272	0.0366	32
EDE_2	1.0604	0.0192	0.9689	0.0264	1.0397	0.1044	1.1384	0.1331	69
PDE_2	0.9704	0.0209	0.9182	0.0189	0.8721	0.0512	0.8432	0.0438	66
EDE_3	1.0330	0.0128	0.9520	0.0304	0.7112	0.0868	0.7016	0.0815	99
PDE_3	0.9529	0.0143	0.9146	0.0146	0.7763	0.0465	0.7001	0.0183	101
EDE_4	1.0133	0.0126	0.9262	0.0265	0.6370	0.0751	0.5860	0.0598	131
PDE_4	0.9449	0.0117	0.8985	0.0185	0.7264	0.0431	0.6370	0.0216	136
EDE_5	1.0297	0.0153	0.9357	0.0297	0.8901	0.1549	0.7961	0.1343	165
PDE_5	0.9691	0.0157	0.9106	0.0150	0.8826	0.1021	0.8241	0.0894	172

Table. 1: Area under the curve (AUC) and rate of change of the curve (RC) for the calibration error and sparsification error curves. Subscript number of the method: the number of models used in ensembles; Bold values: the best ones; Red values: the worst ones. Run time: time consumption for outputting mean and variance, the test is executed on 1 NVIDIA TITAN RTX.

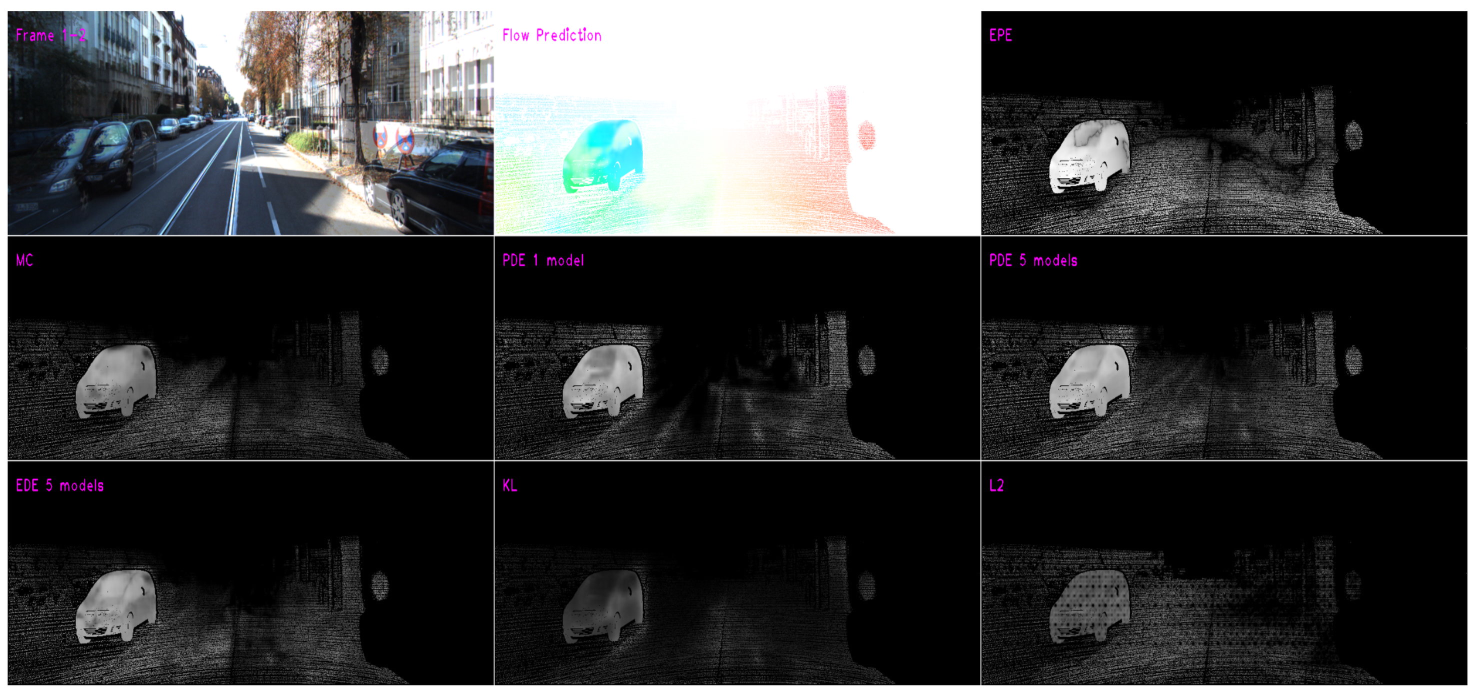

Fig. 3: Uncertainty visualizations for different estimation methods. To visualize the uncertainties we combine the variance from two channels based on our hypothesis, and we use the entropy of a Gaussian distribution. For the joint visualization, all entropies are re-scaled using the maximum value among them.

Conclusions

- In-depth characterization of current most popular techniques for uncertainty estimation applied to optical flow.
- Proposal of a new plug-in training technique, loss and criteria.
- Support for transferring a model trained on synthetic data and applied in real-world, diverse settings.

References

[Dos+15] Alexey Dosovitskiy et al. "Flownet: Learning optical flow with convolutional networks". In: *Proceedings of the IEEE international conference on computer vision*. 2015, pp. 2758–2766.

[GG16] Yarin Gal and Zoubin Ghahramani. "Dropout as a bayesian approximation: Representing model uncertainty in deep learning". In: *international conference on machine learning*. 2016, pp. 1050–1059.

[Ilg+18] Eddy Ilg et al. "Uncertainty estimates and multi-hypotheses networks for optical flow". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2018, pp. 652–667.

[KFE18] Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. "Accurate uncertainties for deep learning using calibrated regression". In: *arXiv preprint arXiv:1807.00263* (2018).

[KG17] Alex Kendall and Yarin Gal. "What uncertainties do we need in bayesian deep learning for computer vision?" In: *Advances in neural information processing systems*. 2017, pp. 5574–5584.

[LPB17] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. "Simple and scalable predictive uncertainty estimation using deep ensembles". In: *Advances in neural information processing systems*. 2017, pp. 6402–6413.

[MG15] Moritz Menze and Andreas Geiger. "Object scene flow for autonomous vehicles". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2015, pp. 3061–3070.