

CNN-based Patch Matching for Optical Flow with Thresholded Hinge Embedding Loss

Christian Bailer¹

Kiran Varanasi¹

Didier Stricker^{1,2}

Christian.Bailer@dfki.de

Kiran.Varanasi@dfki.de

Didier.Stricker@dfki.de

¹German Research Center for Artificial Intelligence (DFKI), ²University of Kaiserslautern

Abstract

Learning based approaches have not yet achieved their full potential in optical flow estimation, where their performance still trails heuristic approaches. In this paper, we present a CNN based patch matching approach for optical flow estimation. An important contribution of our approach is a novel thresholded loss for Siamese networks. We demonstrate that our loss performs clearly better than existing losses. It also allows to speed up training by a factor of 2 in our tests. Furthermore, we present a novel way for calculating CNN based features for different image scales, which performs better than existing methods. We also discuss new ways of evaluating the robustness of trained features for the application of patch matching for optical flow. An interesting discovery in our paper is that low-pass filtering of feature maps can increase the robustness of features created by CNNs. We proved the competitive performance of our approach by submitting it to the KITTI 2012, KITTI 2015 and MPI-Sintel evaluation portals where we obtained state-of-the-art results on all three datasets.

1. Introduction

In recent years, variants of the PatchMatch [5] approach showed not only to be useful for nearest neighbor field estimation, but also for the more challenging problem of large displacement optical flow estimation. So far, most top performing methods like Deep Matching [32] or Flow Fields [3] strongly rely on robust multi-scale matching strategies, while they still use engineered features (data terms) like SIFTFlow [22] for the actual matching.

On the other hand, works like [30, 34] demonstrated the effectiveness of features based on Convolutional Neural Network (CNNs) for matching patches. However, these works did not validate the performance of their features using an actual patch matching approach like *PatchMatch* or *Flow Fields* that matches all pixels between image pairs. Instead, they simply treat matching patches as a classification

problem between a predefined set of patches.

This ignores many practical issues. For instance, it is important that CNN based features are not only able to distinguish between different patch positions, but the position should also be determined accurately. Furthermore, the top performing CNN architectures are very slow when used for patch matching as it requires matching several patches for every pixel in the reference image. While Siamese networks with L_2 distance [30] are reasonably fast at testing time and still outperform engineered features regarding classification, we found that they are usually underperforming engineered features regarding (multi-scale) patch matching.

We think that this has among other things (see Section 4) to do with the convolutional structure of CNNs: as neighboring patches share intermediate layer outputs it is much easier for CNNs to learn matches of neighboring patches than non neighboring patches. However, due to propagation [5] correctly matched patches close to each other usually contribute less for patch matching than patches far apart from each other. Classification does not differentiate here.

A first solution to succeed in CNN based patch matching is to use pixel-wise batch normalization [12]. While it weakens the unwanted convolutional structure, it is computationally expensive at test time. Thus, we do not use it. Instead, we improve the CNN features themselves to a level that allows us to outperform existing approaches.

Our first contribution is a novel loss function for the Siamese architecture with L_2 distance [30]. We show that the hinge embedding loss [30] which is commonly used for Siamese architectures and variants of it have an important design flaw: they try to decrease the L_2 distance unlimitedly for correct matches, although very small distances for patches that differ due to effects like illumination changes or partial occlusion are not only very costly but also unnecessary, as long as false matches have larger L_2 distances. We demonstrate that we can significantly increase the matching quality by relaxing this flaw.

Furthermore, we present a novel way to calculate CNN based features for the scales of Flow Fields [3], which clearly outperforms the original multi-scale feature creation

approach, with respect to CNN based features. Doing so, an important finding is that low-pass filtering CNN based feature maps robustly improves the matching quality.

Moreover, we introduce a novel matching robustness measure that is tailored for binary decision problems like patch matching (while ROC and PR are tailored for classification problems). By plotting the measure over different displacements and distances between a wrong patch and the correct one we can reveal interesting properties of different loss functions and scales. Our main contributions are:

1. A novel loss function, that clearly outperforms other state-of-the-art losses in our tests and allows to speed up training by a factor of around two.
2. A novel multi-scale feature creation approach tailored for CNN features for optical flow.
3. New evaluation measure of matching robustness for optical flow and corresponding plots.
4. We show that low-pass filtering the feature maps created by CNNs improves matching robustness.
5. We demonstrate the effectiveness of our approach by obtaining a top performance on all three major evaluation portals KITTI 2012 [14], 2015 [25] and MPI-Sintel [8]. Former learning based approaches always trailed heuristic approaches on at least one of them.

2. Related Work

While regularized optical flow estimation goes back to Horn and Schunck [18], randomized patch matching [5] is a relatively new field, first successfully applied in approximate nearest neighbor estimation where the data term is well-defined. The success in optical flow estimation (where the data term is not well-defined) started with publications like [4, 10]. One of the most recent works is Flow Fields [3], which showed that with proper multi-scale patch matching, top performing optical flow results can be achieved.

Regarding patch or descriptor matching with learned data terms, there exists a fair amount of literature [17, 30, 34, 31]. These approaches treat matching at an abstract level and do not present a pipeline to solve a problem like optical flow estimation or 3D reconstruction, although many of them use 3D reconstruction datasets for evaluation. Zagoruyko and Komodakis [34] compared different architectures to compare patches. Simo-Serra et al. [30] used the Siamese architecture [6] with L_2 distance. They argued that it is the most useful one for practical applications.

Recently, several successful CNN based approaches for stereo matching appeared [35, 23, 24]. However, so far there are still few approaches that successfully use learning to compute optical flow. Worth mentioning is FlowNet [11]. They tried to solve the optical flow problem as a whole with CNNs, having the images as CNN input and the optical flow

Layer	1	2	3	4	5	6	7	8
Type	Conv	MaxPool	Conv	Conv	MaxPool	Conv	Conv	Conv
Input size	56x56	52x52	26x26	22x22	18x18	9x9	5x5	1x1
Kernel size	5x5	2x2	5x5	5x5	2x2	5x5	5x5	1x1
Out. channels	64	64	80	160	160	256	512	256
Stride	1	2	1	1	2	1	1	1
Nonlinearity	Tanh	-	Tanh	Tanh	-	Tanh	Tanh	Tanh

Table 1. The CNN architecture used in our experiments.

as output. While the results are good regarding runtime, they are still not state-of-the-art quality. Also, the network is tailored for a specific image resolution and to our knowledge training for large images of several megapixel is still beyond today's computational capacity.

A first approach using patch matching with CNN based features is PatchBatch [12]. They managed to obtain state-of-the-art results on the KITTI dataset [14], due to pixel-wise batch normalization and a loss that includes batch statistics. However, pixel-wise batch normalization is computationally expensive at test time. Furthermore, even with pixel-wise normalization their approach trails heuristic approaches on MPI-Sintel [8]. A recent approach is Deep-DiscreteFlow [15] which uses DiscreteFlow [26] as basis instead of patch matching. Despite using recently invented dilated convolutions [23] (we do not use them, yet) they also trail the original DiscreteFlow approach on some datasets.

3. Our Approach

Our approach is based on a Siamese architecture [6]. The aim of Siamese networks is to learn to calculate a meaningful feature vector $D(p)$ for each image patch p . During training the L_2 distance between feature vectors of matching patches ($p_1 \equiv p_2^+$) is reduced, while the L_2 distance between feature vectors of non-matching patches ($p_1 \neq p_2^-$) is increased (see [30] for a more detailed description).

Siamese architectures can be strongly speed up at testing time as neighboring patches in the image share convolutions. Details on how the speedup works are described in our supplementary material. The network that we used for our experiments is shown in Table 1. Similar to [7], we use Tanh nonlinearity layers as we also have found them to outperform ReLU for Siamese based patch feature creation.

3.1. Loss Function and Batch Selection

The most common loss function for Siamese network based feature creation is the hinge embedding loss:

$$l_h(p_1, p_2) = \begin{cases} L_2(p_1, p_2), & p_1 \equiv p_2 \\ \max(0, m - L_2(p_1, p_2)), & p_1 \neq p_2 \end{cases} \quad (1)$$

$$L_2(p_1, p_2) = \|D(p_1) - D(p_2)\|_2 \quad (2)$$

It tries to minimize the L_2 distance of matching patches and to increase the L_2 distance of non-matching patches




Figure 1. If a sample is pushed (blue arrow), although it is clearly on the correct side of the decision boundary other samples also move due to weight change. If most samples are classified correctly beforehand, this creates more false decision boundary crossings than correct ones. l_h performs the unnecessary push, l_t not.

above m . An architectural flaw which is not or only poorly treated by existing loss functions is the fact that the loss pushes feature distances between matching patches unlimitedly to zero ($L_2(p_1, p_2^+) \rightarrow 0$). We think that training up to very small L_2 distances for patches that differ due to effects like rotation or motion blur is very costly – it has to come at the cost of failure for other pairs of patches. A possible explanation for this cost is shown in Figure 1. As a result, we introduce a modified hinge embedding loss with threshold t that stops the network from minimizing L_2 distances too much:

$$l_t(p_1, p_2) = \begin{cases} \max(0, L_2(p_1, p_2) - t), & p_1 \equiv p_2 \\ \max(0, m - (L_2(p_1, p_2) - t)), & p_1 \neq p_2 \end{cases} \quad (3)$$

We add t also to the second equation to keep the “virtual decision boundary” at $m/2$. This is not necessary but makes comparison between different t values fairer.

As our goal is a network that creates features with the property $L_2(p_1, p_2^+) < L_2(p_1, p_2^-)$ one might argue that it is better to train this property directly. A known function to do this is a gap based loss [17, 33], that only keeps a gap in the L_2 distance between matching and non-matching pairs:

$$l_g(p_1, p_2^+) = \max(0, L_2(p_1, p_2^+) - L_2(p_1, p_2^-) + g), \quad (4)$$

$$p_1 \equiv p_2^+ \cap p_1 \neq p_2^-$$

$l_g(p_1, p_2^-)$ is set to $-l_g(p_1, p_2^+)$ (reverse gradient). While l_g intuitively seems to be better suited for the given problem than l_t , we will show in Section 4 why this is not the case. There we will also compare l_t to further loss functions.

The given loss functions have in common that the loss gradient is sometimes zero. Ordinary approaches still back propagate a zero gradient. This not only makes the approach slower than necessary, but also leads to a variable effective batch size of training samples, that are actually back propagated. This is a limited issue for the hinge embedding loss l_h , where only $\approx 25\%$ of the training samples obtain a zero gradient in our tests. However, with l_t (and suitable t) more than 80% of the samples obtain a zero gradient.

As a result, we only add training samples with a non-zero loss to a batch. All other samples are rejected without back

propagation. This not only increases the training speed by a factor of around two in our tests, but also improves the training quality by avoiding variable effective batch sizes.

3.2. Training

Our training set consists of several pairs of images ($I_1, I_2 \in I_{all}$) with known optical flow displacement between their pixels. We first subtract the mean from each image and divide it by its standard derivation. To create training samples, we randomly extract patches $p_1 \in I_1$ and their corresponding matching patches $p_2^+ \in I_2, p_1 \equiv p_2^+$ for positive training samples. For each p_1 , we also extract one non-matching patch $p_2^- \in I_2, p_1 \neq p_2^-$ for negative training samples. Negative samples p_2^- are sampled from a distribution $N(p_2^+)$ that prefers patches close to the matching patch p_2^+ , with a minimum distance to it of 2 pixels, but it also allows to sample patches that are far from p_2^+ . The exact distribution can be found in our supplementary material.

We only train with pairs of patches where the center pixel of p_1 is not occluded in the matching patch p_2^+ . Otherwise, the network would train the occluding object as a positive match. However, if the patch center is visible we expect the network to be able to deal with a partial occlusion. We use a learning rate between 0.004 and 0.0004 that decreases linearly in exponential space after each batch i.e. $learnRate(t) = e^{-x_t} \rightarrow learnRate(t+1) = e^{-(x_t+\epsilon)}$.

3.3. Multi-scale matching

The *Flow Fields* approach [3], which we use as basis for our optical flow pipeline compares patches at different scales using scale spaces [21], i.e. all scales have the full image resolution. It creates feature maps for different scales by low-pass filtering the feature map of the highest scale (Figure 2 left). For SIFTFlow [22] features used in [3], low-pass filtering features (i.e. feature \rightarrow low-pass = feature \rightarrow downsample \rightarrow upsample) performs better than recalculating features for each scale on a different resolution (i.e. downsample \rightarrow feature \rightarrow upsample).

We observed the same effect for CNN based features – even if the CNN is also trained on the lower resolutions. However, with our modifications shown in Figure 2 right (that are further motivated in Section 4), it is possible to obtain better results by recalculating features on different resolutions. We use a CNN trained and applied only on the highest image resolution for the highest and second highest scale. Furthermore, we use a CNN trained on 3 resolutions (100%, 50% and 25%) to calculate the feature maps for the third and fourth scale applied at 50% and 25% resolution, respectively. For the multi-resolution CNN, the probability to select a patch on a lower resolution for training is set to be 60% of the probability for the respective next higher resolution. For lower resolutions, we also use the distribution $N(p_2^+)$. This leads to a more wide spread distribution with

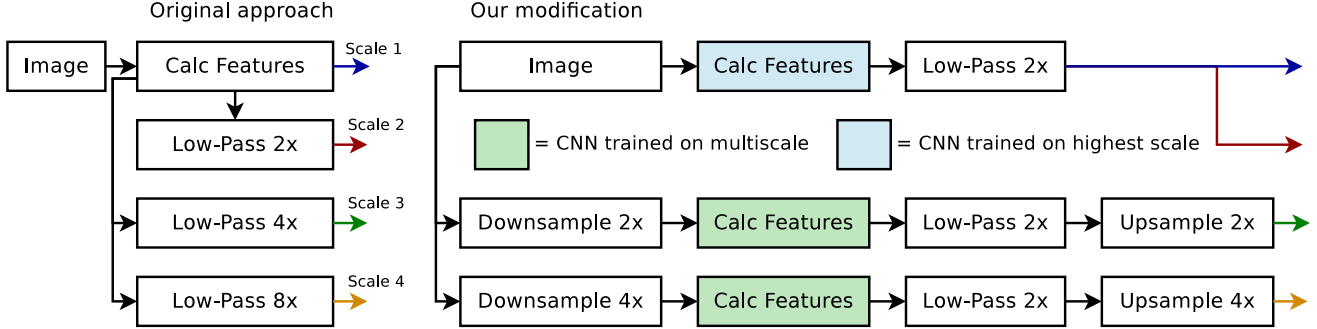


Figure 2. Our modification of feature creation of the *Flow Fields* approach [3] for clearly better CNN performance. Note that *Flow Fields* expects feature maps of all scales in the full image resolution (See [3] for details). Reasons of design decision can be found in Section 4.1.

respect to the full image resolution.

Feature maps created by our CNNs are not used directly. Instead, we perform a 2x low-pass filter on them, before using them. Low-pass filtering image data creates matching invariance while increasing ambiguity (by removing high frequency information). Assuming that CNNs are unable to create perfect matching invariance, we can expect a similar effect on feature maps created by CNNs. In fact, a small low-pass filter clearly increases the matching robustness.

The *Flow Fields* approach [3] uses a secondary consistency check with different patch size. With our approach, this would require to train and execute two additional CNNs. To keep it simple, we perform the secondary check with the same features. This is possible due to the fact that *Flow Fields* is a randomized approach. Still, our tests with the original features show that a real secondary consistency check performs better. The reasoning for our design decisions in Figure 2 can be found in Section 4.1.

3.4. Evaluation Methodology for Patch Matching

In previous works, the evaluation of the matching robustness of (learning based) features was performed by evaluation methods commonly used in classification problems like ROC in [7, 34] or PR in [30]. However, patch matching is not a classification problem, but a binary decision problem. While one can freely label data in classification problems, patch matching requires to choose, at each iteration, out of two proposal patches p_2, p_2^* the one that fits better to p_1 . The only exception from this rule is outlier filtering. This is not really an issue, as there are better approaches for outlier filtering, like the forward backward consistency check [3], which is more robust than matching-error based outlier filtering¹. In our evaluation, the *matching robustness* r of a network is determined as the probability that a wrong patch p_2^- is not confused with the correct patch p_2^+ :

$$r = \sum_{(I_1, I_2) \in S} \sum_{p_1 \in I_1} P_{p_1}^r(p_2^-) / (|I_1| |S|) \quad (5)$$

¹Even if outlier filtering would be performed by matching error, the actual matching remains a decision problem.

$$P_{p_1}^r(p_2^-) = P(L_2(p_1, p_2^+) < L_2(p_1, p_2^-)), \quad (6)$$

$$p_1 \equiv p_2^+ \in I_2, p_1 \neq p_2^- \in N(p_2^+),$$

where S is a set of considered image pairs (I_1, I_2) , $|S|$ the number of image pairs and $|I_1|$ the number of pixels in I_1 . As r is a single value we can plot it for different cases:

1. The curve for different spatial distances between p_2^+ and p_2^- (r_{dist}).
2. The curve for different optical flow displacements between p_1 and p_2^+ (r_{flow}).

r_{dist} and r_{flow} vary strongly for different locations. This makes differences between different networks hard to visualize. For better visualization, we plot the *relative matching robustness errors* E_{dist} and E_{flow} , computed with respect to a pre-selected network $net1$. E is defined as:

$$E(net1, net2) = (1 - r(net2)) / (1 - r(net1)) \quad (7)$$

4. Evaluation

We examine our approach on the KITTI 2012 training set [14] as it is one of the few datasets that contains ground truth for non-synthetic large displacement optical flow estimation. We use patches taken from 130 of the 194 images of the set for training and patches from the remaining 64 images for validation. Each tested network is trained with 10 million negative and 10 million positive samples in total. Furthermore, we publicly validate the performance of our approach by submitting our results to the KITTI 2012, the recently published KITTI 2015 [25] and MPI-Sintel evaluation portals (with networks trained on the respective training set). We use the original parameters of the *Flow Fields* approach [3] except for the outlier filter distance ϵ and the random search distance R . ϵ is set to the best value for each network (with accuracy ± 0.25 , mostly: $\epsilon = 1.5$). The random search distance R is set to 2 for four iterations and to $R = 1$ for two additional iterations to increase accuracy. The batch size is set to 100 and m to 1.

To evaluate the quality of our optical flow results we calculate the endpoint error (EPE) for non-occluded areas

Approach	EPE >3 px noc.	EPE >3 px all	EPE noc.	EPE all	Approach	EPE >3 px noc.	EPE >3 px all	EPE noc.	EPE all
ours	4.95%	11.89%	1.10 px	2.60 px	all resolutions	5.66%	13.01%	1.27 px	2.98 px
original ([3]+CNN)	5.48%	12.59%	1.28 px	3.08 px	nowlpass	<u>5.21%</u>	<u>12.21%</u>	<u>1.19 px</u>	<u>2.80 px</u>
ms res 1	5.17%	12.10%	1.17 px	2.80 px	ms res 2+	5.18%	12.12%	1.21 px	2.84 px

Table 2. Comparison of CNN based multi-scale feature creation approaches. See text for details.

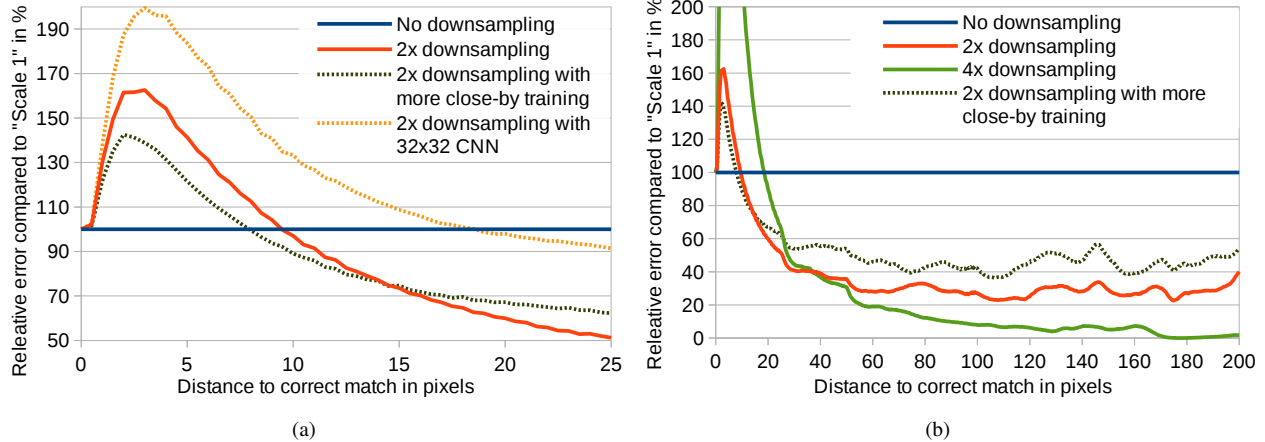


Figure 3. Relative matching robustness errors E_{dist} (“No Downsampling”, X). Features created on lower resolutions are more accurate for large distances but less accurate for small ones. *No downsampling* is on the horizontal line as results are normalized for it. Details in text.

(noc) as well as occluded + non-occluded areas (all). (noc) is a more direct measure as CNNs are only trained here. However, the interpolation into occluded areas (like *Flow Fields* we use EpicFlow [28] for that) also depends on good matches close to the occlusion boundary, where matching is especially difficult due to partial occlusions of patches. Furthermore, like [14], we measure the percentage of pixels with an EPE above a threshold in pixels (px).

4.1. Comparison of CNN based Multi-Scale Feature Map Approaches

In Table 2, we compare the *original* feature creation approach (Figure 2 left) with *our* approach (Figure 2 right), with respect to our CNN features. We also examine two variants of our approach in the table: *nowlpass* which does not contain the “Low-Pass 2x” blocks and *all resolutions* which uses 1x,2x,4x,8x up/downsampling for the four scales (instead of 1x,1x,2x,4x in Figure 2 right). The reason why *all resolutions* does not work well is demonstrated in Figure 3 (a). Starting from a distance between p_2^+ and p_2^- of 9 pixels, CNN based features created on a 2x down-sampled image match more robustly than CNN based features created on the full image resolution. This is insufficient as the random search distance on scale 2 is only $2R = 4$ pixels. Thus, we use it for scale 3 (with random search distance $4R = 8 \approx 9$ pixels).

One can argue that by training the CNN with more close-by samples $N_{close}(p_2^+)$ more accuracy could be gained. But

Approach/ Loss	EPE >3 px noc.	EPE >3 px all	EPE noc.	EPE all	robust- ness r
L_h	7.26%	14.78%	1.46 px	3.33 px	98.63%
$L_t, t = 0.2$	6.17%	13.51%	1.37 px	3.10 px	99.15%
$L_t, t = 0.3$	4.95%	11.89%	1.10 px	2.60 px	99.34%
$L_t, t = 0.4$	<u>5.18%</u>	<u>12.10%</u>	1.25 px	3.14 px	99.41%
$L_g, g = 0.2$	5.92%	13.17%	1.41 px	3.37 px	99.15%
$L_g, g = 0.4$	5.89%	13.23%	1.41 px	3.36 px	99.31%
$L_g, g = 0.6$	6.37%	13.74%	1.51 px	3.40 px	99.08%
Hard Mining x2 [30]	6.03%	13.34%	1.35 px	2.99 px	99.07%
DrLIM [16]	5.36%	12.40%	1.18 px	2.79 px	99.15%
CENT. [12]	6.32%	13.90%	1.46 px	3.37 px	98.72%
SIFTFlow [22]	11.52%	19.79%	1.99 px	4.33 px	97.31%
SIFTFlow* [22]	5.85%	12.90%	1.52 px	3.56 px	97.31%

Table 3. Results on KITTI 2012 [14] validation set. Best result is bold, 2. best underlined. *SIFTFlow* uses our pipeline tailored for CNNs. *SIFTFlow** uses the original pipeline [3] (Figure 2 left).

raising extremely the amount of close-by samples only reduces the accuracy threshold from 9 to 8 pixels. Using a CNN with smaller 32x32 patches instead of 56x56 patches does not raise the accuracy either– it even clearly decreases it. Figure 3 (b) shows that downsampling decreases the matching robustness error significantly for larger distances. In fact, for a distance above 170 pixels, the relative error of *4x downsampling* is reduced by nearly 100% compared to *No downsampling* – which is remarkable.

Multi-resolution network training We examine three variants of training our multi-resolution network (green boxes in Figure 2): training it on 100%, 50% and 25% resolution although it is only used for 50% and 25% resolution, at testing time (*ours* in Table 2), training it on 50% and 25% resolutions, where it is used for at testing time (*ms res 2+*) and training it only on 100% resolution (*ms res 1*). As can be seen in Table 2 training on all resolutions (*ours*) clearly performs best. Likely, mixed training data performs best as samples of the highest resolution provide the largest entropy while samples of lower resolutions fit better to the problem. However, training samples of lower resolutions seem to harm training for higher resolutions. Therefore, we use an extra CNN for the highest resolution.

4.2. Loss Functions and Mining

We compare our loss l_t to other state-of-the-art losses and Hard Mining [30] in Figure 4 and Table 3. As shown in the table, our thresholded loss l_t with $t = 0.3$ clearly outperforms all other losses. DrLIM [16] reduces the mentioned flaw in the hinge loss, by training samples with small hinge loss less. While this clearly reduces the error compared to hinge, it cannot compete with our thresholded loss l_t . Furthermore, no speedup during training is possible like with our approach. CENT. (CENTRIFUGE) [12] is a variant of DrLIM which performs worse than DrLIM in our tests.

Hard Mining [30] only trains the hardest samples with the largest hinge loss and thus also speeds up training. However, the percentage of samples trained in each batch is fixed and does not adapt to the requirements of the training data like in our approach. With our data, Hard Mining becomes unstable with a mining factor above 2 i.e. the loss of negative samples becomes much larger than the loss of positive samples. This leads to poor performance ($r = 96.61\%$ for Hard Mining x4). We think this has to do with the fact that the hardest of our negative samples are much harder to train than the hardest positive samples. Some patches are e.g. fully white due to overexposure (negative training has no effect here). Also, many of our negative samples have, in contrast to the samples of [30], a very small spatial distance to their positive counterpart. This makes their training even harder (We report most failures for small distances, see supplementary material), while positive samples do not change.

To make sure that our dynamic loss based mining approach (L_t with $t = 0.3$) cannot become unstable towards much larger negative loss values we tested it to an extreme: we randomly removed 80% of the negative training samples while keeping all positive. Doing so, it not only stayed stable, but it even used a smaller positive/negative sample mining ratio than the approach with all training samples – possibly it can choose harder positive samples which contribute more to training. Even with the removal of 80% (8 million) of possible samples we achieved a matching ro-




Figure 5. The distribution of L_2 errors for different for L_t and L_g for positive samples p_2^+ and negative samples p_2^- with distance of 10 pixels to the corresponding positive sample.

bustness r of 99.18%.

L_g performed best for $g = 0.4$ which corresponds to a gap of $L_t, t = 0.3$ ($g_{L_t} = 1 - 2t$). However, even with the best g , L_g performs significantly worse than L_t . This is probably due to the fact that the variance $\text{Var}(L_2(p_1, p_2))$ is much larger for L_g than for L_t . As shown in Figure 5, this is the case for both positive (p_2^+) as well as negative (p_2^-) samples. We think this affects the test set negatively as follows: if we assume that p_1, p_2^+, p_2^- are unlearned test set patches it is clear that the condition $L_2(p_1, p_2^+) < L_2(p_1, p_2^-)$ is more likely violated if $\text{Var}(L_2(p_1, p_2^+))$ and $\text{Var}(L_2(p_1, p_2^-))$ are large compared to the learned gap. Only with L_t it is possible to force the network to keep the variance small compared to the gap. With L_g it is only possible to control the gap but not the variance, while l_h keeps the variance small but cannot limit the gap.

Matching Robustness plots Some loss functions perform worse than others although they have a larger matching robustness r . This mostly can be explained by the fact that they perform poorly for large displacements (as shown in Figure 4 (b)). Here, correct matches are usually more important as missing matches lead to larger endpoint errors. An averaged r over all pixels does not consider this.

Figure 4 also shows the effect of parameter t in L_t . Up to $t \approx 0.3$, all distances and flow displacements are improved, while small distances and displacements benefit more and up to a larger $t \approx 0.4$. The improvement happens as unnecessary destructive training is avoided (see Section 3.1). Patches with small distances benefit more from larger t , likely as the *real gap* $g_{real} = |L_2(p_1, p_2^-) - L_2(p_1, p_2^+)|$ is smaller here (as p_2^- and p_2^+ are very similar for small distances). For large displacements patches get more chaotic (due to more motion blur, occlusions etc.), which forces larger variances of the L_2 distances and thus a larger gap is required to counter the larger variance.

L_g performs worse than L_t mainly at small distances and large displacements. Likely, the larger variance is more destructive for small distances, as the *real gap* g_{real} is smaller




Figure 4. Relative matching robustness errors $E(L_t, t = 0.3, X)$ for different loss functions plotted for different distances (a) and displacements (b). Note that the plot for $L_t, t = 0.3$ is on the horizontal line, as E is normalized for it. See text for details.

(more sensitive) here. Figure 4 also shows that low-pass filtering the feature map increases the matching robustness for all distances and displacements. In our tests, a $2.25\times$ low-pass performed the best (tested with ± 0.25). Engineered SIFTFlow features can benefit from much larger low-pass filters which makes the original pipeline (Figure 2 left) extremely efficient for them. However, using them with our pipeline (which recalculates features on different resolutions) shows that their low matching robustness is justified (see Table 3). SIFTFlow also performs better in outlier filtering. Due to such effects that can so far not directly be trained, it is still challenging to beat well designed purely heuristic approaches with learning. In fact, existing CNN based approaches often still underperform purely heuristic approaches – even direct predecessors (see Section 4.3).

4.3. Public Results

Our public results on the KITTI 2012 [14], 2015 [25] and MPI-Sintel [8] evaluation portals are shown in Table 4, 5 and 6. For the public results we used 4 extra iterations with $R = 1$ for best possible subpixel accuracy and for similar runtime to *Flow Fields* [3]. t is set to 0.3. On KITTI 2012 our approach is the best in all measures, although we use a smaller patch size than *PatchBatch* (71×71) [12]. *PatchBatch* (51×51) with a patch size more similar to ours performs even worse. *PatchBatch**(51×51) which is like our work without pixel-wise batch normalization even trails purely heuristic methods like *Flow Fields*.

On KITTI 2015 our approach also clearly outperforms *PatchBatch* and all other general optical flow methods including *DeepDiscreteFlow* [15] that, despite using CNNs, trails its engineered predecessor *DiscreteFlow* [26] in many measures. The only methods that outperform our approach are the rigid segmentation based methods *SDF* [1], *JFS* [20]

and *SOF* [29]. These require segmentable rigid objects moving in front of rigid background and are thus not suited for scenes that contain non-rigid objects (like MPI-Sintel) or objects which are not easily segmentable. Despite not making any such assumptions our approach outperforms two of them in the challenging foreground (moving cars with reflections, deformations etc.). Furthermore, our approach is clearly the fastest of all top performing methods although there is still optimization potential (see below). Especially, the segmentation based methods are very slow.

On the non rigid MPI-Sintel datasets our approach is the best in the non-occluded areas, which can be matched by our features. Interpolation into occluded areas with *EpicFlow* [28] works less well, which is no surprise as aspects like good outlier filtering which are important for occluded areas are not learned by our approach. Still, we obtained the best overall result on the more challenging final set that contains motion blur. In contrast, *PatchBatch* lags far behind on MPI-Sintel, while *DeepDiscreteFlow* again clearly trails its predecessor *DiscreteFlow* on the clean set, but not the final set. Our approach never trails on the relevant matchable (non-occluded) part.

Our detailed runtime is 4.5s for CNNs (GPU) + 16.5s patch matching (CPU) + 2s for up/downsampling and low-pass (CPU). The CPU parts of our approach likely can be significantly sped up using GPU versions like a GPU based propagation scheme [2, 13] for patch matching. This is contrary to *PatchBatch* where the GPU based CNN already takes the majority of time (due to pixel-wise normalization). Also, in final tests (after submitting to evaluation portals) we were able to improve our CNN architecture (see supplementary material) so that it only needs 2.5s with only a marginal change in quality on our validation set.

Method	EPE >3 px noc.	EPE >5 px noc.	EPE >3 px all	EPE >5 px all	EPE noc.	EPE all	runtime
Ours (56x56)	4.89 %	3.04 %	13.01 %	9.06 %	1.2 px	3.0 px	23s
PatchBatch (71x71) [12]	<u>4.92 %</u>	<u>3.31 %</u>	<u>13.40 %</u>	<u>10.18 %</u>	1.2 px	<u>3.3 px</u>	60s
PatchBatch (51x51) [12]	5.29 %	3.52 %	14.17 %	10.36 %	1.3 px	3.3 px	<u>50s</u>
Flow Fields [3]	5.77 %	3.95 %	14.01 %	10.21%	1.4 px	3.5 px	23s
PatchBatch*(51x51)[12]	5.94% [12]	-	-	-	-	-	25.5s [12]

Table 4. Results on KITTI 2012 [14] test set. Numbers in brackets show the patch size for learning based methods. Best result for published methods is bold, 2. best is underlined. PatchBatch* is PatchBatch without pixel-wise batch normalization.

Type	Method	background		foreground (cars)		total		runtime
		EPE >3 px noc.	EPE >3 px all	EPE >3 px noc.	EPE >3 px all	EPE >3 px noc.	EPE >3 px all	
Rigid Segmentation based methods	SDF [1]	5.75%	8.61%	22.28%	26.69%	8.75%	11.62%	unknown
	JFS [20]	7.85%	15.90%	18.66%	22.92%	9.81%	17.07%	13 min
	SOF [29]	8.11%	14.63%	23.28%	27.73%	10.86%	16.81%	6 min
General methods	Ours (56x56)	8.91%	18.33%	20.78%	24.96%	11.06%	19.44%	23s
	PatchBatch (51x51) [12]	10.06%	<u>19.98%</u>	26.21%	30.24%	12.99%	21.69%	<u>50s</u>
	DiscreteFlow [26]	<u>9.96%</u>	21.53 %	22.17%	<u>26.68 %</u>	<u>12.18%</u>	22.38%	3 min
	DeepDiscreteFlow [15]	10.44%	20.36 %	25.86%	29.69 %	13.23%	21.92%	1 min

Table 5. Results on KITTI 2015 [25] test set. Numbers in brackets shows the used patch size for learning based methods. Best result for all published general optical flow methods is bold, 2. best underlined. Bold for segmentation based method shows that the result is better than the best general method. Rigid segmentation based methods were designed for urban street scenes and similar containing only segmentable rigid objects and rigid background (and are usually very slow), while general methods work for all optical flow problems.

5. Conclusion and Future Work

In this paper, we presented a novel extension to the hinge embedding loss that not only outperforms other losses in learning robust patch representations, but also allows to increase the training speed and to be robust with respect to unbalanced training data. We presented a new multi-scale feature creation approach for CNNs and proposed new evaluation measures by plotting matching robustness with respect to patch distance and motion displacement. Furthermore, we showed that low-pass filtering feature maps created by CNNs improves the matching result. All together, we proved the effectiveness of our approach by submitting it to the KITTI 2012, KITTI 2015 and MPI-Sintel evaluation portals where we, as the first learning based approach, achieved state-of-the-art results on all three datasets. Our results also show the transferability of our contribution, as our findings made in Section 4.1 and 4.2 (on which our architecture is based on) are solely based on KITTI 2012 validation set, but still work unchanged on KITTI 2015 and MPI-Sintel test sets, as well.

In future work, we want to improve our network architecture (Table 1) by using techniques like (non pixel-wise) batch normalization and dilated convolutions [23]. Furthermore, we want to find out if low-pass filtering invariance also helps in other application, like sliding window object detection [27]. We want to further improve our loss function L_t e.g. by a dynamic t that depends on the properties of training samples. So far, we just tested a patch size of 56x56 pixels, although [12] showed that larger patch sizes

Method(final)	EPE all	EPE not occl.	EPE occluded
Ours	5.363	2.303	<u>30.313</u>
DeepDiscreteFlow[15]	<u>5.728</u>	2.623	31.042
FlowFields [3]	5.810	<u>2.621</u>	31.799
CPM-Flow [19]	5.960	2.990	30.177
DiscreteFlow [26]	6.077	2.937	31.685
PatchBatch [12]	6.783	3.507	33.498
Method(clean)	EPE all	EPE not occl.	EPE occluded
CPM-Flow [19]	3.557	1.189	22.889
DiscreteFlow [26]	<u>3.567</u>	1.108	23.626
FullFlow [9]	3.601	1.296	<u>22.424</u>
FlowFields [3]	3.748	<u>1.056</u>	25.700
Ours	3.778	0.996	26.469
DeepDiscreteFlow[15]	3.863	1.296	24.820
PatchBatch [12]	5.789	2.743	30.599

Table 6. Results on MPI-Sintel [8]. Best result for all published methods is bold, second best is underlined.

can perform even better. It might be interesting to find out which is the largest beneficial patch size. Frames of MPI-Sintel with very large optical flow showed to be especially challenging. They lack training data due to rarity, but still have a large impact on the average EPE (due to huge EPE). We want to create training data tailored for such frames and examine if learning based approaches benefit from it.

Acknowledgments

This work was funded by the BMBF project DYNAMICS (01IW15003).

References

- [1] M. Bai, W. Luo, K. Kundu, and R. Urtasun. Exploiting semantic information and deep matching for optical flow. In *European Conference on Computer Vision (ECCV)*, 2016. 7, 8
- [2] C. Bailer, M. Finckh, and H. P. Lensch. Scale robust multi view stereo. In *European Conference on Computer Vision (ECCV)*, 2012. 7
- [3] C. Bailer, B. Taetz, and D. Stricker. Flow fields: Dense correspondence fields for highly accurate large displacement optical flow estimation. In *International Conference on Computer Vision (ICCV)*, 2015. 1, 2, 3, 4, 5, 7, 8
- [4] L. Bao, Q. Yang, and H. Jin. Fast edge-preserving patch-match for large displacement optical flow. In *Computer Vision and Pattern Recognition (CVPR)*, 2014. 2
- [5] C. Barnes, E. Shechtman, A. Finkelstein, and D. Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics-TOG*, 2009. 1, 2
- [6] J. Bromley, J. W. Bentz, L. Bottou, I. Guyon, Y. LeCun, C. Moore, E. Säcker, and R. Shah. Signature verification using a siamese time delay neural network. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(04):669–688, 1993. 2
- [7] M. Brown, G. Hua, and S. Winder. Discriminative learning of local image descriptors. *Pattern Analysis and Machine Intelligence (PAMI)*, 33(1):43–57, 2011. 2, 4
- [8] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *European Conference on Computer Vision (ECCV)*, 2012. <http://sintel.is.tue.mpg.de/results>. 2, 7, 8
- [9] Q. Chen and V. Koltun. Full flow: Optical flow estimation by global optimization over regular grids. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 8
- [10] Z. Chen, H. Jin, Z. Lin, S. Cohen, and Y. Wu. Large displacement optical flow from nearest neighbor fields. In *Computer Vision and Pattern Recognition (CVPR)*, 2013. 2
- [11] P. Fischer, A. Dosovitskiy, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [12] D. Gadot and L. Wolf. Patchbatch: a batch augmented loss for optical flow. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 2, 5, 6, 7, 8
- [13] S. Galliani, K. Lasinger, and K. Schindler. Massively parallel multiview stereopsis by surface normal diffusion. In *International Conference on Computer Vision (ICCV)*, 2015. 7
- [14] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 2013. http://www.cvlibs.net/datasets/kitti/eval_stereo_flow.php?benchmark=flow. 2, 4, 5, 7, 8
- [15] F. Güney and A. Geiger. Deep discrete flow. In *Asian Conference on Computer Vision (ACCV)*, 2016. 2, 7, 8
- [16] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *Computer Vision and Pattern Recognition (CVPR)*, 2006. 5, 6
- [17] X. Han, T. Leung, Y. Jia, R. Sukthankar, and A. C. Berg. Matchnet: unifying feature and metric learning for patch-based matching. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. 2, 3
- [18] B. K. Horn and B. G. Schunck. Determining optical flow. In *Technical symposium east*, pages 319–331. International Society for Optics and Photonics, 1981. 2
- [19] Y. Hu, R. Song, and Y. Li. Efficient coarse-to-fine patch-match for large displacement optical flow. 8
- [20] J. Hur and S. Roth. Joint optical flow and temporally consistent semantic segmentation. In *European Conference on Computer Vision (ECCV)*, 2016. 7, 8
- [21] T. Lindeberg. Scale-space theory: A basic tool for analyzing structures at different scales. *Journal of applied statistics*, 21(1-2):225–270, 1994. 3
- [22] C. Liu, J. Yuen, A. Torralba, J. Sivic, and W. T. Freeman. Sift flow: Dense correspondence across different scenes. In *European Conference on Computer Vision (ECCV)*. 2008. 1, 3, 5
- [23] W. Luo, A. G. Schwing, and R. Urtasun. Efficient deep learning for stereo matching. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 8
- [24] 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 *Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [25] M. Menze and A. Geiger. Object scene flow for autonomous vehicles. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. http://www.cvlibs.net/datasets/kitti/eval_scene_flow.php?benchmark=flow. 2, 4, 7, 8
- [26] M. Menze, C. Heipke, and A. Geiger. Discrete optimization for optical flow. In *German Conference on Pattern Recognition (GCPR)*, 2015. 2, 7, 8
- [27] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Neural Information Processing Systems (NIPS)*, 2015. 8
- [28] J. Revaud, P. Weinzaepfel, Z. Harchaoui, and C. Schmid. Epicflow: Edge-preserving interpolation of correspondences for optical flow. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. 5, 7
- [29] L. Sevilla-Lara, D. Sun, V. Jampani, and M. J. Black. Optical flow with semantic segmentation and localized layers. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 7, 8
- [30] E. Simo-Serra, E. Trulls, L. Ferraz, I. Kokkinos, P. Fua, and F. Moreno-Noguer. Discriminative learning of deep convolutional feature point descriptors. In *International Conference on Computer Vision (ICCV)*, 2015. 1, 2, 4, 5, 6
- [31] K. Simonyan, A. Vedaldi, and A. Zisserman. Learning local feature descriptors using convex optimisation. *Pattern Analysis and Machine Intelligence (PAMI)*, 36(8):1573–1585, 2014. 2

- [32] P. Weinzaepfel, J. Revaud, Z. Harchaoui, and C. Schmid. Deepflow: Large displacement optical flow with deep matching. In *International Conference on Computer Vision (ICCV)*, 2013. [1](#)
- [33] P. Wohlhart and V. Lepetit. Learning descriptors for object recognition and 3d pose estimation. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. [3](#)
- [34] S. Zagoruyko and N. Komodakis. Learning to compare image patches via convolutional neural networks. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. [1](#), [2](#), [4](#)
- [35] J. Zbontar and Y. LeCun. Stereo matching by training a convolutional neural network to compare image patches. *Journal of Machine Learning Research*, 17:1–32, 2016. [2](#)