

Optimizing Through Learned Errors for Accurate Sports Field Registration c'è il codice

Wei Jiang¹ Juan Camilo Gamboa Higuera² Baptiste Angles¹
Weiwei Sun¹ Mehrsan Javan³ Kwang Moo Yi¹

¹Visual Computing Group, University of Victoria ²McGill University ³SPORTLOGiQ Inc.

{jiangwei, bangles, weiweisun, kyi}@uvic.ca, {gamboa}@cim.mcgill.ca, {mehrsan}@sportlogiq.com

Abstract

We propose an *optimization-based framework to register sports field templates onto broadcast videos*. For accurate registration we go beyond the prevalent feed-forward paradigm. Instead, we propose to train a deep network that regresses the registration error, and then register images by finding the registration parameters that minimize the regressed error. We demonstrate the effectiveness of our method by applying it to real-world sports broadcast videos, outperforming the state of the art. We further apply our method on a synthetic toy example and demonstrate that our method brings significant gains even when the problem is simplified and unlimited training data is available.¹

1. Introduction

Estimating the relationship between a template and an observed image with deep learning [9, 38, 24, 46] has received much attention recently, due to the success of deep learning in many other areas in computer vision [17, 16, 40]. Registration of a sports field template onto a camera view is not an exception [6, 19, 42], where deep learning has shown promising results compared to traditional baselines. Despite the recent advancements, there is room for further improvement, especially for augmented reality and sport analytics.

For mixed and augmented reality, even the slightest inaccuracies in estimates can break immersion [31]. For sports analytics, good alignment is crucial for detecting the important events – e.g. offside in soccer.

Existing methods have also acknowledged this limitation, and have sought to improve accuracy. For example, some rely on a hierarchical strategy [38, 24]. In these methods, the refinement network is used on top of a rough pose estimator, where both are feed-forward networks. However, as we will demonstrate through our experiments, there is an alternative way to enhance performance.

In order to achieve more accurate registration, we take a different route to the commonly used feed-forward paradigm. Inspired by classic optimization-based approaches for image registration [37, 30, 28, 35], we propose optimizing to reduce the estimated registration error. Opposed to traditional methods, we rely on a deep network for estimating the error.

Specifically, as illustrated in Fig. 1, we propose a two-stage deep learning pipeline, similar to existing methods [38, 24], but with a twist on the refinement network. The first-stage network, which we refer to as the *initial registration network*, provides a rough estimate of the registration, parameterized by a homography transform. For the second-stage network, instead of a feed-forward refinement network, we train a deep neural network that regresses the error of our estimates – *registration error network*. We then use the initial registration network to provide an initial estimate, and optimize the initial estimate using the gradients provided by differentiating through the registration error network. This allows much more accurate estimates compared to the single stage feed-forward inference.

In addition, we propose *not* to train the two networks together. While end-to-end and joint training is often preferred [39, 36], we find that it is beneficial to train the two networks separately – decoupled training. We attribute this to two observations: the two networks – initial registration network and registration error network – aim to regress different things – pose and error; it is useful for the registration error network to be trained independently of the initial registration network so that it does not overfit to the mistakes the initial registration network makes while training.

We demonstrate empirically that our framework consistently outperforms feed-forward pipelines. We apply our method to sports field registration with broadcast videos. We show that not only our method outperforms feed-forward networks in a typical registration setup, it is also able to outperform the state of the art even when training data is scarce – a trait that is desirable with deep networks. We further show that our method is not limited to sport fields registration. We create a simple synthetic toy dataset

¹ Code is available at https://github.com/vcg-uvic/sportsfield_release.




Figure 1. Illustration of our framework. We train two deep networks each dedicated for a different purpose. We first obtain an initial pose estimate from a feed-forward network that regresses directly to the homography parameterization \mathbf{h} – DNN on the left in blue. We then warp our sports field template according to this initial estimate and concatenate it with the input image. We feed this concatenated image to a second deep neural network – DNN on the right in red – that estimates the error of the current warping. We then differentiate through this network to obtain the direction in which the estimated error is minimized, and optimize our estimated homography accordingly – red arrow. This optimization process is repeated multiple times until convergence. All figures in this paper are best viewed in color.

of estimating equations of a line in the image, and show that even in this simple case when unlimited train data is available, our method brings significant advantage.

To the best of our knowledge, **our method is the first method that learns to regress registration errors for optimization-based image registration.** The idea of training a deep network to regress the error and perform optimization-based inference has recently been investigated for image segmentation, multi-label classification, and object detection [13, 23]. While the general idea exists, it is non-trivial to formulate a working framework for each task. Our work is the first successful attempt for sport field registration, thanks to our two-stage pipeline and the way we train the two networks.

To summarize, our contributions are:

- we propose a **novel two-stage image registration framework** where we iteratively optimize our estimate by differentiating through a learned error surface;
- we propose a **decoupled training strategy** to train the **initial registration network** and the **registration error network**;
- our method **achieves the state-of-the-art performance, even when the training dataset size is as small as 209 images**;
- we **demonstrate the potential of our method through a generic toy example.**

2. Related Work

Sports field registration. Early attempts on registering sports field to broadcast videos [12, 11, 37] typically rely on a set of pre-calibrated reference images. These calibrated references are used to estimate a relative pose to the image of interest. To retrieve the relative pose, these methods either assume that images are of correspond to consecutive frames in a video [12, 11], or use local features, such as

SIFT [29] and MSER [32], to find correspondences [37]. These methods, however, require that the set of calibrated images contains images with similar appearance to the current image of interest, as traditional local features are weak against long-term temporal changes [47]. While learned alternatives exist [50, 36, 10], their performances in the context of pose estimation and registration remains questionable [41].

To overcome these limitations, more recent methods [19, 42, 6] focus on converting broadcast videos into images that only contain information about sports fields, *e.g.* known marker lines, then perform registration. Homayounfar *et al.* [19] perform semantic segmentation on broadcast images with a deep network, then optimize for the pose using branch and bound [27] with a Markov Random Field (MRF) formulated with geometric priors. While robust to various scenic changes, their accuracy is still limited. Sharma *et al.* [42] simplify the formulation by focusing on the edges and lines of sports fields, rather than the complex semantic segmentation setup. They use a database of edge images generated with known homographies to extract the pose, which is then temporally smoothed. Chen and Little [6] further employ an image translation network [21] in a hierarchical setup where the play-field is first segmented out, followed by sports field line extraction. They also employ a database to extract the pose, which is further optimized through Lucas-Kanade optimization [30] on distance transformed version of the edge images. The bottleneck of these two methods is the necessity of a database, which hinders their scalability.

Homography estimation between images. Traditional methods for homography estimation include sparse feature-based approaches [49] and dense direct approaches [30]. Regardless of sparse or dense, traditional approaches are mainly limited by either the quality of the local features [48], or by the robustness of the objective function used for optimization [2].

Deep learning based approaches have also been proposed

for homography estimation. In [9], the authors propose to train a network that directly regresses to the homography between two images through self supervision. Interestingly, the output of the regression network is discretized, allowing the method to be formulated as classification. Nguyen *et al.* [34] train a deep network in an unsupervised setup to learn to estimate the relative homography. The main focus of these methods, however, is on improving the inference speed, without significant improvements on accuracy when compared to traditional baselines.

Feed-forward 6 Degree-of-Freedom (DoF) pose estimators. Pose estimators are also highly related to image registration. Deep networks have also been proposed to directly regress the 6 DoF pose of cameras [51, 43, 25]. Despite being efficient to compute, these methods highly depend on their parameterization of the pose – naive parameterizations can lead to bad performance, and are known to have limited accuracy [5]. To overcome this limitation, recent works focus on regressing the 2D projection of 3D control points [46, 38, 24]. Compared with directly predicting the pose, control points based pose show improved performance due to the robust estimation of parameters. Our initial registration network follows the same idea as these methods to obtain our initial estimate.

Optimizing with learned neural networks. Incorporating optimization into deep pipelines is a current topic of interest. BA-Net [45] learns to perform Levenberg-Marquardt optimization within the network to solve dense bundle adjustment. LS-Net [7] learns to predict the directions to improve a given camera pose estimate. Han *et al.* [14] also learn to estimate the Jacobian matrix from an image pair to update the 6 DoF camera pose. In contrast to these methods, which propose learning a function to update a camera pose estimate, we propose to learn an error function that predicts how well two images are aligned. Using the error function we can obtain the update direction via differentiation. The most similar work to ours is the deep value networks [13], where they train a network to estimate intersection over union (IoU) between the input and ground truth masks regarding image segmentation. While sharing a similar idea, it is non-trivial to extend and adapt their method to image registration. For example, their method is limited to a static initial estimate, which requires a longer optimization trajectory than ours. This may become a problem when applied to sport field registration, where the broadcast view change drastically even when there is small camera rotation. We show through experiments that just having an error network is not enough, as we will show later in Table 2.

3. Method

For clarity in presentation, we first assume that our models are pre-trained and detail our overall framework at inference time.

We then provide details on the training setup and the architectural choices.

3.1. Inference

Overview. Our pipeline is depicted in Fig. 1. We assume a known planar sports field template and undistorted images, so that we can represent the image-template alignment with a homography matrix. The framework can be broken down into two stages: the first stage provides an initial estimate of the homography matrix, the second iteratively optimizes this estimate. The first stage follows a typical feed-forward paradigm [9, 46], and we utilize a deep neural network. However, any method can be used here instead, such as a database search [42, 6].

The distinctiveness of our model comes from the second stage of the pipeline. Using the first stage estimate, we warp the sports field template to the current view. We concatenate this warped image with the current observed image, and evaluate the registration error through a second neural network. We then backpropagate the estimated error to the homography parameters to obtain the gradient, which gives the direction in which the parameters should be updated to minimize the registration error. Then, using this gradient, we update the homography parameters. This process is performed iteratively until convergence or until a maximum number of iterations is met. This *inference through optimization* allows our method to be significantly more accurate than a typical feed-forward setup, provided that our error model gives reasonable error predictions.

Details – initial registration. We follow the recent trend of using projected coordinates for pose parameterization [9, 5]. In the case of homographies, this can be done with 4 points [1]. We parameterize the homography \mathbf{h} defining the relationship between the input image \mathbf{I} and the target template \mathbf{m} through the coordinate of the four control points on the current input image when warped onto the sports field template. Specifically, considering a normalized image coordinate system where the width and height of the image is set to one, and the centre of the image is at the origin, we use $(-0.5, 0.1)$, $(-0.5, 0.5)$, $(0.5, 0.5)$, and $(0.5, 0.1)$, that is, the corners of the lower three-fifths of the image as our reference control points. We write the reference control points \mathbf{h}_{ref} as,

$$\mathbf{h}_{ref} = [-0.5, 0.1, -0.5, 0.5, 0.5, 0.5, 0.5, 0.1]^\top. \quad (1)$$

We use the lower parts of the image as sports field broadcast videos are typically in a setup where the camera is looking down on the field, as shown in Fig. 2.

Let (u_k, v_k) denote the k -th control point of the current image \mathbf{I} projected onto the sports field template \mathbf{m} . We then write the homography \mathbf{h} as

$$\mathbf{h} = [u_1, v_1, u_2, v_2, u_3, v_3, u_4, v_4]^\top. \quad (2)$$




Figure 2. Illustration of control points. The yellow dots on the left are the control points we use on the normalized image coordinate, and the red dots on the right are the control points after they are transformed via the homography \mathbf{h} . Our initial registration network regresses the positions of the red dots.

We obtain the actual transformation matrix \mathbf{T} from \mathbf{h} and \mathbf{h}_{ref} through direct linear transformation [15].

Given an initial registration network $f_{\Phi}(\cdot)$, we obtain a rough homography estimate $\hat{\mathbf{h}}^{(0)}$ for image \mathbf{I} as

$$\hat{\mathbf{h}}^{(0)} = f_{\Phi}(\mathbf{I}), \quad (3)$$

where the superscript in parenthesis denote the optimization iteration.

Details – optimization. With the current homography estimate $\hat{\mathbf{h}}^{(i)}$ at optimization iteration i , we warp the play-field template \mathbf{m} to obtain an image of the template in the current view, using a bilinear sampler [22] to preserve differentiability. We concatenate the result of this warping operation $\mathcal{W}(\mathbf{m}, \hat{\mathbf{h}}^{(i)})$ and the image \mathbf{I} , and pass it as input to the model $g_{\Psi}(\cdot)$ to obtain a prediction of the registration error $\hat{\epsilon}^{(i)}$ as

$$\hat{\epsilon}^{(i)} = g_{\Psi}([\mathbf{I}; \mathcal{W}(\mathbf{m}, \hat{\mathbf{h}}^{(i)})]), \quad (4)$$

where $[\cdot; \cdot]$ denotes concatenation along the channel direction of two images. We then retrieve the gradient of $\hat{\epsilon}^{(i)}$ with respect to $\hat{\mathbf{h}}^{(i)}$ and apply this gradient to retrieve an updated estimate. In practice, we rely on Adam [26] for a stable optimization.

Note here that our registration error network is not trained to give updates. It simply regresses to the correctness of the current estimate. We show empirically in Section 4.3 that this is a much more effective than, for example learning to provide a perfect homography, or learning to correct erroneous estimates.

3.2. Training

To avoid overfitting, we propose to purposely decouple the training of two networks. We show in Section 4.3 that this is necessary in order to obtain the best performance.

Initial registration network. To train the initial registration network, we directly regress the four control points of our template warped into a given view using the ground

truth homography. With the ground truth homography \mathbf{h}_{gt} , we train our deep network to minimize

$$\mathcal{L}_{init} = \|\mathbf{h}_{gt} - \hat{\mathbf{h}}^{(0)}\|_2^2 = \|\mathbf{h}_{gt} - f_{\Phi}(\mathbf{I})\|_2^2. \quad (5)$$

Note that while we use a deep network to obtain the initial homography estimate, any other method can also be used in conjunction, such as nearest neighbor search.

Registration error network. To train the registration error network, we create random perturbations on the ground truth homography. We then warp the target template to the view using the perturbed ground truth homography, and concatenate it with the input image to be used as input data for training. The network model is trained to predict a registration error metric, e.g. the IoU. We detail our design choice of error metric in Section 4.3.

In more detail, with the ground truth homography \mathbf{h}_{gt} , we create a perturbed homography \mathbf{h}_{pert} by applying uniform noise hierarchically: one for global translation, and one for local translation of each control point. Specifically, we add a global random translation $\alpha_g \sim U(-\delta_g, \delta_g)$, where $\alpha_g \in \mathbb{R}^2$, to all control points, and add a local random translation of $\alpha_l \sim U(-\delta_l, \delta_l)$, where $\alpha_l \in \mathbb{R}^8$ individually to each control point. We then warp the target template according to the perturbed homography to create our input data for training. Thus, the input to the registration error network for training is $[\mathbf{I}; \mathcal{W}(\mathbf{m}, \mathbf{h}_{pert})]$. Then, to train the network, we minimize

$$\mathcal{L}_{error} = \|\text{Err}(\mathbf{I}, \mathcal{W}(\mathbf{m}, \mathbf{h}_{pert})) - g_{\Psi}([\mathbf{I}; \mathcal{W}(\mathbf{m}, \mathbf{h}_{pert})])\|_2^2, \quad (6)$$

where $\text{Err}(\cdot, \cdot)$ is the error metric, for example the IoU value.

4. Sports field registration results

We apply the proposed method to sports field registration. We first discuss the datasets, baselines, the metrics used for our evaluation, as well as implementation details. We then present qualitative and quantitative results of our method, compared to the state of the art. We then provide experimental insights to our method.

4.1. Experimental setup

Datasets. To validate our method, we rely on two datasets. The World Cup dataset [20] is a dataset made of broadcast videos of soccer games. It has 209 images for training and validation, and 186 images for testing. This dataset is extremely small, making it challenging to apply deep methods. The state of the art for this dataset [6] relies on learning to transfer the input image to look similar to the sports field template, then searching a database of known homographies and warped templates to retrieve the estimate. For

our method, we use 39 images from the train-valid split as validation dataset, and respect the original test split for testing. The Hockey dataset is composed of broadcast videos of NHL ice hockey games [19]. This is a larger dataset than the World Cup dataset, having 1.67M images in total. Of this large dataset, we use two sequences of 800 consecutive images as validation and testing sets. By using consecutive frames, we ensure that images from one game do not fall into different splits. See Fig. 3 for example images.

Baselines. We compare our method against three existing works for sports field registration [19, 42, 6]. As there is no publicly available implementation of the two methods [19, 42], we take the results reported on the respective papers for the World Cup dataset. For [6], we use the authors’ public implementation. For [19] with the Hockey dataset, we use the reported results as a reference².

In addition, we compare our method against feed forward baselines – single stage feed-forward network (SFF) and a two-stage feed-forward refinement network (FFR). We further explore whether the error registration network can be used alone by retrieving the initial estimate by searching a database of known poses, *e.g.* the training set, and using the example which gives the lowest error estimate. We will refer to the initial estimate obtained through nearest neighbor search as NN, and the fully optimized estimate as NNo. To do a nearest neighbor search we evaluate the registration error for the query image with all the training homographies using the trained registration error network, and return the homography with lowest estimated error. Although this method is not scalable because the computational requirement grows linearly with the size of the database, it provides insight into the capability of the trained registration error network.

Metrics. As existing literature use different metrics [19, 42, 6], IoU_{part} and $\text{IoU}_{\text{whole}}$, we report both values. IoU_{part} is the intersection over union when only the visible region is considered, while $\text{IoU}_{\text{whole}}$ is the same considering the entire sports field template.

4.2. Implementation details

Initial registration network. Following a recent trend [39, 16], we base our network on the ResNet-18 architecture [18]. Instead of the classification head, we simply replace the last fully connected layer to estimate 8 numbers which represent the homography, \mathbf{h} . We use the pretrained weights for the network trained on ImageNet [8], and fine-tune.

Registration error network. For the registration error network, we also rely on the ResNet-18 architecture, but with

² No information is provided by the authors on how the train, validation, and test splits are created, thus the results are not directly comparable.

spectral normalization [33] on all convolutional layers, and take as input a 6-channel image, that is, the concatenation of the input image and the warped target template. Spectral normalization smooths the error predictions by constraining the Lipschitz constant of the model, which limits the magnitude of its gradients. As the output of the registration error network cannot be negative, we use sigmoid function as the final activation function for the IoU-based error metrics, and squaring function for reprojection error metric. For the registration network, as the input is very different from a typical image-based network, we train from scratch.

Hyperparameters. We train our networks with the Adam [26] optimizer, with default parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$, and with a learning rate of 0.0001. We train until convergence, and use the validation dataset to perform early stopping. For the noise parameters δ_g and δ_l in Section 3.2 we empirically set $\delta_g = 0.05$ and $\delta_l = 0.02$, by observing the validation dataset results. For inference, we again use Adam, but with a learning rate of 10^{-3} . We run our optimization for 400 iterations, and return the estimate that gave the lowest estimated error predicted by the trained registration error network.

4.3. Results

Comparison against existing pipelines. Qualitative highlights are shown in Fig. 3 and Fig. 4, with quantitative results summarized in Table 1. In Table 1, for the World Cup dataset, our method performs best in all evaluation metrics. For the Hockey dataset, our method delivers near perfect results.

Comparison against feed-forward baselines. As shown in Table 1, having an additional feed-forward refinement network (FFR) only provides minor improvement over the initial estimate (SFF). This phenomenon is more obvious in the WorldCup dataset results, where training data is scarce. By contrast, our method is able to provide significant reduction in the registration error.

Effect of different target error metrics. We also compare results when different target error is used for the training of the registration error network in Table 2. We compare regressing to $\text{IoU}_{\text{whole}}$, IoU_{part} , and the average reprojection error of all pixels inside the current view (Reproj.). Interestingly, regressing to IoU_{part} does not guarantee best performance in terms of IoU_{part} . In all cases, regressing to $\text{IoU}_{\text{whole}}$ gives best performance.

Coupled training. It is a common trend to train multiple components together. However, our framework does not allow joint training, as the two networks are aiming for entirely different goals. Nonetheless, we simultaneously trained the two networks, thus allowing the registration error network to see all the mistakes that the initial registration network makes during training (Coupled). Coupled train-

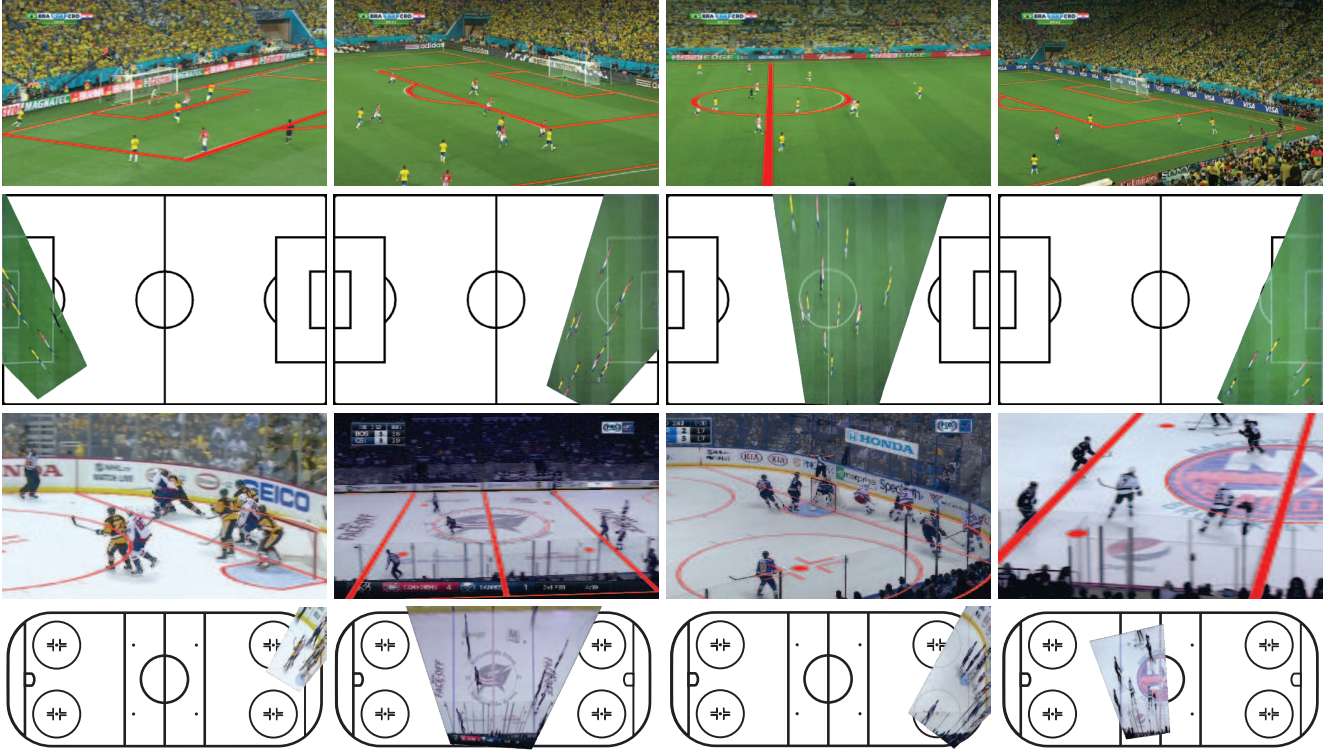


Figure 3. Qualitative highlights of our method. (Top) red lines are the sports field lines overlayed on the current view using estimated homographies. (Bottom) current view overlayed on sports field template. Our method can handle various sports fields and camera poses.



Figure 4. Qualitative example demonstrating the effect of number of optimization iterations on registration accuracy. From left to right, example registration result at iterations 0, 20, 40 and 60. Notice the misalignment near the center circle. As more optimization iterations are performed, the registration becomes more accurate.

ing, however, performs worse than decoupled training, as shown in Table 2. In case of the Hockey dataset, coupled training performs even worse than feed-forward refinement. This is because while the initial registration network is converging, it is making predictions with smaller and smaller mistakes, thus the registration error network is learning a narrow convergence basin due to the small perturbations it sees. The estimates that fall out of the convergence basin can not be optimized using the learned error. Therefore, it is necessary to have a decoupled training setup to stop this from happening.

Using only the error estimation network. The two variants, NN and NNo, provide insights into the capability of the registration error networks. Due to the limited size of the database, i.e. training data, NN provides initial es-

timates with lower accuracy than the single stage feed-forward network SFF. However, with optimization (NNo), the registration results are even comparable to the results from our full pipeline. This observation shows that the registration error network can provide a wide convergence basin, and can optimize for inaccurate initial estimates.

Note that we only test these methods on the World Cup dataset, as applying the method on Hockey dataset requires too much computation due to the larger database to search.

Inference performance We perform all experiments on an Nvidia GTX 1080Ti GPU. To optimize one frame, our method achieves 41.76 optimization iterations per second, thus 9.58 seconds per frame. Our method also supports batch inference. To optimize a batch with 64 frames, it achieves 4.66 optimization iterations per second, thus in av-

			[19]	[42]	[6]	SFF	FFR	Ours
World Cup	IoU _{whole}	mean	83	–	89.2	83.9	84.0	89.8
		median	–	–	91.0	85.7	86.2	92.9
	IoU _{part}	mean	–	91.4	94.7	90.2	90.3	95.1
		median	–	92.7	96.2	91.9	92.1	96.7
Hockey	IoU _{whole}	mean	82 ²	–	–	86.5	93.0	96.2
		median	–	–	–	87.3	94.0	97.0
	IoU _{part}	mean	–	–	–	90.4	96.0	97.6
		median	–	–	–	91.0	96.8	98.4

Table 1. Quantitative results for different methods. Best results are in bold. Our method performs best in all evaluation metrics. See text for details.

			IoU _{whole}	IoU _{part}	Reproj.	Coupled	NN	NNo
World Cup	IoU _{whole}	mean	89.8	87.9	89.1	87.3	73.8	86.3
		median	92.9	90.6	91.4	91.1	73.6	88.2
	IoU _{part}	mean	95.1	94.7	95.1	94.4	87.4	94.0
		median	96.7	96.3	96.5	96.5	89.5	95.7
Hockey	IoU _{whole}	mean	96.2	95.6	94.9	87.9	–	–
		median	97.0	96.6	95.5	89.5	–	–
	IoU _{part}	mean	97.6	97.3	97.1	93.6	–	–
		median	98.4	98.3	97.6	94.7	–	–

Table 2. Quantitative results for different variants of our method. Best results are in bold. IoU_{whole}, IoU_{part}, and Reproj. are three target error metrics we investigate. Coupled is when we couple the training of two networks. NN is when we use nearest neighbor search and NNo is when we further optimize the homography estimate with the registration error network after NN.

erage 1.36 seconds per frame.

4.4. Quality of the estimated error surface

To validate that the trained registration error network can provide a convergence basin, we visualize the average estimated error surface for translation over all test samples. To do so we create a regular grid with X from $[-0.5, 0.5]$, and Y from $[-0.5, 0.5]$ with resolution 50 by 50. For each point on the grid we warp the template with ground truth homography combined with the translation from the origin to the point location. We then pass the observed image concatenated with the warped sports field to the trained registration error network, and infer the registration error at that point on the grid. We calculate the error surface for all the test samples, and visualize the average.

As show in Fig. 5, the estimated error surface resembles the ground truth one. The error is lower towards the origin where the perturbation – translation – is smaller, and is higher towards the border where the perturbation is larger.

Most importantly, the minima of the estimated error is very close to the origin, which is the ground truth. This allows our optimization-based inference to work properly.




Figure 5. Average estimated and ground truth error surface visualization for translation. See how the estimated error surface resembles the ground truth one, including the location of the minima at the centre. This allows optimization through learned errors.

5. Toy experiment – Line fitting

Beyond sport fields registration, our method could be applied to other tasks that involve parameter regression. Here, we show briefly that, even a task as simple and generic as fitting a line equation in an image can benefit from our method. We hope to shed some light into the potentials of our method.

Inspired by the experiment from DSAC [3, 4], we validate our framework with the task of estimating the equation of a line from synthetic images, as shown in Fig. 6. Unlike DSAC, we are not learning to reject outliers via their pixel coordinates, but rather are directly regressing to the line equations given an image of a line.

5.1. Experimental setup

Initial network. We follow the same setup as our image registration task, but instead regress two parameters that define the equation of a line, that is, a – the angle and b – the intercept, where the line equation is given by $v = \tan(a)u + b$, and u and v are the image coordinates.

To create the synthetic images, we first generate random lines by selecting a random pivot point in an 64×64 image, then uniformly sample in range $[-0.4\pi, 0.4\pi]$ to obtain its angle. We draw this line with a random color. We then add a random colored ellipse with random parameters as distraction, and finally apply additive Gaussian noise. We use VGG-11 [44] as the backbone for the initial registration network.

Error network. We use the intercept error as the target error metric, that is maximum error between ground-truth and estimated intercept at $u = 0$ or $u = 63$. To generate erroneous estimates for training, we add uniform noise α_a and α_b to the ground truth a and b respectively, where




Figure 6. Estimating line equations of synthetic images. The red line represents the estimated line equation from (top row) a feed-forward network, and (bottom row) the proposed method. The other colored line in each image is the target line. Our method provides accurate estimates, shown by the high overlap with the thick white line. See Section 5 for details.

$\alpha_a \sim U[-0.1\pi, 0.1\pi]$, and $\alpha_b \sim U[-5, 5]$. To render the estimate into an image in a differentiable way, we warp the template image which is simply an image of a line, using the hypothesized line parameters as in the case of image registration. We concatenate the input image with the warp template to the error network to estimate the error, in this case the intercept error. We also use VGG-11 as the backbone for the error network.

We train both networks until convergence and optimize for 400 iterations at inference time.

5.2. Results

As shown in Fig. 6, our method estimates the line parameters more accurately than a feed-forward deep network. Quantitative results are shown in Table 3. As shown, even in this simple generic task, our method outperforms its feed-forward counterpart. As this task can be viewed as a simplified version of other computer vision tasks, it shows that our method may be applicable outside the scope of the current paper. We further highlight that this experimental setup is with unlimited labeled data. Even in such a case, our method brings significant improvement in performance.

	Feed-forward	Ours
mean error	5.1	3.0
median error	4.4	1.5

Table 3. Quantitative results for line fitting. Our method achieved better accuracy than a single stage feed-forward network. This line fitting experiment can be viewed as a general regression task.

6. Conclusions

We have proposed a two-stage pipeline for registering sports field templates to broadcast videos accurately. In contrast to existing methods that do single stage feed-

forward inference, we opted for an optimization-based inference inspired by established classic approaches. The proposed method makes use of two networks, one that provides an initial estimate for the registration homography, and one that estimates the error given the observed image and the current hypothesized homography. By optimizing through the registration error network, accurate results were obtained.

We have shown through experiments that the proposed method can be trained with very sparse data, as little as 209 images, and achieve state-of-the-art performance. We have further revealed how different design choices in our pipeline affect the final performance. Finally, we have shown that our framework can be translated into other tasks and improve upon feed-forward strategies.

As future work, since the inference is optimization-based, we can naturally embed temporal consistency by reusing the optimization state for consecutive images to register sports field for a video. We show preliminary results of doing so in our supplementary video.

References

- [1] S. Baker, A. Datta, and T. Kanade. Parameterizing Homographies. Technical report, Robotics Institute, Carnegie Mellon University, 2006. 3
- [2] S. Baker and I. Matthews. Lucas-Kanade 20 Years On: A Unifying Framework. *International Journal of Computer Vision*, pages 221–255, 2004. 2
- [3] E. Brachmann, A. Krull, S. Nowozin, J. Shotton, F. Michel, S. Gumhold, and C. Rother. Dsac-differentiable ransac for camera localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6684–6692, 2017. 7
- [4] E. Brachmann and C. Rother. Learning less is more-6D camera localization via 3D surface regression. In *CVPR*, 2018. 7

- [5] T. Bugra, S. Sudipta N, and F. Pascal. Real-time Seamless Single Shot 6d Object Pose Prediction. In *Conference on Computer Vision and Pattern Recognition*, 2018. 3
- [6] J. Chen and J. J. Little. Sports Camera Calibration via Synthetic Data. *Conference on Computer Vision and Pattern Recognition Workshops*, 2019. 1, 2, 3, 4, 5, 7
- [7] R. Clark, M. Bloesch, J. Czarowski, S. Leutenegger, and A. J. Davison. LS-Net: Learning to Solve Nonlinear Least Squares for Monocular Stereo. *arXiv Preprint*, 2018. 3
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A Large-Scale Hierarchical Image Database. In *Conference on Computer Vision and Pattern Recognition*, 2009. 5
- [9] D. DeTone, T. Malisiewicz, and A. Rabinovich. Deep image homography estimation. In *RSS Workshop on Limits and Potentials of Deep Learning in Robotics*, 2016. 1, 3
- [10] D. Detone, T. Malisiewicz, and A. Rabinovich. Superpoint: Self-Supervised Interest Point Detection and Description. *CVPR Workshop on Deep Learning for Visual SLAM*, 2018. 2
- [11] B. Ghanem, T. Zhang, and N. Ahuja. Robust Video Registration Applied to Field-sports Video Analysis. In *International Conference on Acoustics, Speech, and Signal Processing*, 2012. 2
- [12] A. Gupta, J. J. Little, and R. Woodham. Using Line and Ellipse Features for Rectification of Broadcast Hockey Video. In *Canadian Conference on Computer and Robot Vision*, 2011. 2
- [13] M. Gygli, M. Norouzi, and A. Angelova. Deep value networks learn to evaluate and iteratively refine structured outputs. 2017. 2, 3
- [14] L. Han, M. Ji, L. Fang, and M. Nießner. RegNet: Learning the Optimization of Direct Image-to-Image Pose Registration. *arXiv Preprint*, 2018. 3
- [15] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000. 4
- [16] K. He, G. Gkioxari, P. Dollar, and R. Girshick. Mask R-CNN. In *International Conference on Computer Vision*, 2017. 1, 5
- [17] K. He, X. Zhang, R. Ren, and J. Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on Imagenet Classification. In *International Conference on Computer Vision*, 2015. 1
- [18] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *Conference on Computer Vision and Pattern Recognition*, 2016. 5
- [19] N. Homayounfar, S. Fidler, and R. Urtasun. Sports Field Localization via Deep Structured Models. In *Conference on Computer Vision and Pattern Recognition*, 2017. 1, 2, 5, 7
- [20] N. Homayounfar, S. Fidler, and R. Urtasun. Sports Field Localization via Deep Structured Models. In *Conference on Computer Vision and Pattern Recognition*, 2017. 4
- [21] P. Isola, J. Zhu, T. Zhou, and A. Efros. Image-To-Image Translation with Conditional Adversarial Networks. *Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [22] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu. Spatial Transformer Networks. In *Advances in Neural Information Processing Systems*, 2015. 4
- [23] B. Jiang, R. Luo, J. Mao, T. Xiao, , and Y. Jiang. Acquisition of Localization Confidence for Accurate Object Detection. In *European Conference on Computer Vision*, 2018. 2
- [24] W. Kehl, F. Manhardt, F. Tombari, S. Ilic, and N. Navab. SSD-6D: Making Rgb-Based 3D Detection and 6D Pose Estimation Great Again. In *International Conference on Computer Vision*, 2017. 1, 3
- [25] A. Kendall, M. Grimes, and R. Cipolla. Posenet: A Convolutional Network for Real-Time 6-DOF Camera Relocalization. In *International Conference on Computer Vision*, 2015. 3
- [26] D. Kingma and J. Ba. Adam: A Method for Stochastic Optimisation. In *International Conference on Learning Representations*, 2015. 4, 5
- [27] C. Lampert, M. Blaschko, and T. Hofmann. Efficient Sub-window Search: A Branch and Bound Framework for Object Localization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31:2129–2142, 2009. 2
- [28] V. Lepetit, F. Moreno-noguer, and P. Fua. EPnP: An Accurate O(n) Solution to the PnP Problem. *International Journal of Computer Vision*, 81(2), 2009. 1
- [29] D. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, 20(2):91–110, 2004. 2
- [30] B. Lucas and T. Kanade. An Iterative Image Registration Technique with an Application to Stereo Vision. In *International Joint Conference on Artificial Intelligence*, 1981. 1, 2
- [31] E. Marchand, H. Uchiyama, and F. Spindler. Pose Estimation for Augmented Reality: a Hands-on Survey. *IEEE Transactions on Visualization and Computer Graphics*, 2016. 1
- [32] J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust Wide-Baseline Stereo from Maximally Stable Extremal Regions. *Image and Vision Computing*, 22(10):761–767, 2004. 2
- [33] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral Normalization for Generative Adversarial Networks. In *International Conference on Learning Representations*, 2018. 5
- [34] T. Nguyen, S. W. Chen, S. S. Shivakumar, C. J. Taylor, and V. Kumar. Unsupervised Deep Homography: A Fast and Robust Homography Estimation Model. *IEEE Robotics and Automation Letters*, 2018. 3
- [35] D. Oberkamp, D. DeMenthon, and L. Davis. Iterative Pose Estimation Using Coplanar Feature Points. *Computer Vision, Graphics, and Image Processing*, 63(3):495–511, 1996. 1
- [36] Y. Ono, E. Trulls, P. Fua, and K. M. Yi. Lf-Net: Learning Local Features from Images. In *Advances in Neural Information Processing Systems*, 2018. 1, 2
- [37] J. Puwein, R. Ziegler, J. Vogel, and M. Pollefeys. Robust Multi-view Camera Calibration for Wide-baseline Camera Networks. In *IEEE Winter Conference on Applications of Computer Vision*, 2011. 1, 2
- [38] M. Rad and V. Lepetit. Bb8: A Scalable, Accurate, Robust to Partial Occlusion Method for Predicting the 3D Poses of Challenging Objects Without Using Depth. In *International Conference on Computer Vision*, 2017. 1, 3

- [39] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Advances in Neural Information Processing Systems*, 2015. 1, 5
- [40] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Conference on Medical Image Computing and Computer Assisted Intervention*, 2015. 1
- [41] J. Schönberger, H. Hardmeier, T. Sattler, and M. Pollefeys. Comparative Evaluation of Hand-Crafted and Learned Local Features. In *Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [42] R. A. Sharma, B. Bhat, V. Gandhi, and C. V. Jawahar. Automated Top View Registration of Broadcast Football Videos. In *IEEE Winter Conference on Applications of Computer Vision*, 2018. 1, 2, 3, 5, 7
- [43] Siddharth Mahendran and Haider Ali and Rene Vidal. 3d pose regression using convolutional neural networks. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, 2017. 3
- [44] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *International Conference on Learning Representations*, 2015. 7
- [45] C. Tang and P. Tan. Ba-Net: Dense Bundle Adjustment Network. In *International Conference on Learning Representations*, 2019. 3
- [46] B. Tekin, S. Sinha, and P. Fua. Real-Time Seamless Single Shot 6D Object Pose Prediction. In *Conference on Computer Vision and Pattern Recognition*, 2018. 1, 3
- [47] Y. Verdie, K. M. Yi, P. Fua, and V. Lepetit. TILDE: A Temporally Invariant Learned DETector. In *Conference on Computer Vision and Pattern Recognition*, 2015. 2
- [48] F. Wu and F. Xiangyong. An Improved RANSAC Homography Algorithm for Feature Based Image Mosaic. In *Proceedings of the 7th WSEAS International Conference on Signal Processing, Computational Geometry & Artificial Vision*, 2007. 2
- [49] Q. Yan, Y. Xu, X. Yang, and T. Nguyen. HEASK: Robust Homography Estimation Based on Appearance Similarity and Keypoint Correspondences. *Pattern Recognition*, 2014. 2
- [50] K. M. Yi, E. Trulls, V. Lepetit, and P. Fua. LIFT: Learned Invariant Feature Transform. In *European Conference on Computer Vision*, 2016. 2
- [51] X. Yu, T. Schmidt, V. Narayanan, and F. Dieter. PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes. *Robotics: Science and Systems (RSS)*, 2018. 3