

There and Back Again: Self-supervised Multispectral Correspondence Estimation

Celyn Walters¹, Oscar Mendez¹, Mark Johnson, Richard Bowden¹

Abstract—Across a wide range of applications, from autonomous vehicles to medical imaging, multi-spectral images provide an opportunity to extract additional information not present in color images. One of the most important steps in making this information readily available is the accurate estimation of dense correspondences between different spectra.

Due to the nature of cross-spectral images, most correspondence solving techniques for the visual domain are simply not applicable. Furthermore, most cross-spectral techniques utilize spectra-specific characteristics to perform the alignment. In this work, we aim to address the dense correspondence estimation problem in a way that generalizes to more than one spectrum. We do this by introducing a novel cycle-consistency metric that allows us to self-supervise. This, combined with our spectrum-agnostic loss functions, allows us to train the same network across multiple spectra.

We demonstrate our approach on the challenging task of dense RGB-FIR correspondence estimation. We also show the performance of our unmodified network on the cases of RGB-NIR and RGB-RGB, where we achieve higher accuracy than similar self-supervised approaches. Our work shows that cross-spectral correspondence estimation can be solved in a common framework that learns to generalize alignment across spectra.

I. INTRODUCTION

Solving the correspondence problem between two images is a fundamental problem in computer vision. Its applications are widespread, including 3D reconstruction [1], motion estimation [2] and image registration [3]. Correspondence estimation from RGB to RGB is well understood, with many solutions using correlation [4], optimisation [5], hand-crafted feature descriptors [3], [6], [7] or machine learning [8]–[11]. However, in some use-cases, relying on the visible spectrum alone is insufficient. For example, in areas such as autonomous navigation and visual surveillance, approaches using RGB cameras often fail at night, in poor weather, or due to extreme variability in lighting. Using alternative spectra, such as Infrared (IR), is a commonly used technique to address these concerns. Near Infra-Red (NIR), considering its similarities to RGB, may utilize vision algorithms developed for the visible spectrum. Far Infra-Red (FIR) gives much stronger thermal cues but requires specialized techniques. Bridging this gap will lead to more capable applications.

Thermal sensors tend to be low resolution and lack fine detail at the far range, and high resolution thermal sensors are expensive. Multispectral fusion can overcome the deficiencies of each individual sensor by combining their complementary properties. However, most traditional applications require the images to be registered.

Relative to RGB, cross-spectral correspondence estimation approaches are scarce in the literature. This is partially due to the availability and cost of sensors, in addition to the complications of solving the correspondence problem when photometric consistency between sensors does not hold true. Most multispectral datasets/methods purposefully avoid tackling stereo disparity, by either focusing on scenes at long range, where disparity is assumed to be negligible [12], or by using a beam splitter to ensure coaxial camera centres [13]–[15]. Unfortunately both approaches ignore possible stereo cues. We argue that there is benefit in correctly modelling disparity as it provides more accurate sensor fusion while providing additional stereo cues. Furthermore, neither high-resolution thermal cameras nor beam splitters are commodity items, and are therefore unlikely to be incorporated into consumer products.

Perhaps one of the most limiting factors behind research into multispectral correspondence is the poor availability of datasets. As cross-spectral research is relatively unexplored, few annotated datasets are available and ground truth annotation is time-consuming and expensive. Automatic annotation is an unsolved problem, and requires expensive precision equipment such as a laser scanner. Due to unfamiliarity with the appearance differences, human annotators find annotation difficult. Therefore, most provide sparse point annotation or weak labels such as bounding boxes [13], [16]–[19]. This motivates our proposed self-supervised approach which learns how to recover a dense flow field between RGB and IR, and vice-versa. We make the following contributions:

- 1) **Dense flow fields between spectra:** We present a spectrum-agnostic method to obtain cross-spectral flow fields at full image resolution as a way to solve the correspondence problem between different spectra.
- 2) **Self-supervised training:** We use a dual-spectrum siamese-like structure, utilizing cycle-consistency to avoid the need for ground truth correspondence. This provides scalability, allowing us to significantly increase the training data seen by the system.
- 3) **Application to RGB-FIR:** We demonstrate RGB-FIR correspondence estimation, which is seldom tackled.
- 4) **Application to RGB-RGB and RGB-NIR:** We further demonstrate our approach with competitive results for both RGB-NIR and RGB-RGB correspondence.

II. RELATED WORK

Traditional solutions to the correspondence problem are to solve it sparsely, using optimisation, correlation, or feature-matching. However, most approaches rely upon photometric

¹Celyn Walters *et al.* with Centre for Vision, Speech and Signal Processing (CVSSP), University of Surrey, UK c.walters@surrey.ac.uk

consistency, and this assumption does not hold well for cross-spectral matching. For wide-baselines, feature-matching approaches using hand-crafted feature descriptors are common (e.g. SIFT [3]). Although other features have been used for dense matching [6], [7], most approaches rely on features optimized for RGB and as a result, performance is lower when applied to other spectra [20]. Work that has attempted to adapt traditional techniques to RGB-NIR [21], [22] tends to have drawbacks, e.g. they scale poorly when applied densely over an entire image. Kim *et al.* proposed DASC [23], a dense descriptor which finds illumination/sharpness differences between both RGB and NIR pairs. Considering the closeness of NIR to the visible spectrum and its similarities with greyscale, NIR-RGB correspondence is perhaps easier than for other spectra like FIR. RGB-NIR-specific approaches do not transfer well to RGB-FIR.

Same-spectrum correspondence techniques can be adapted from RGB to other spectra. For instance, thermal depth images [24] and thermal stereo odometry [25]. For RGB-FIR, Li and Stevenson use a straight line matching scheme to register stereo images [26], but a reliance on straight edges is likely to fail in less structured scenes. By contrast, our approach uses feature losses to ensure we are robust to cross-spectral images and less structured scenes.

Mutual Information (MI) does not rely upon photometric consistency, and is widely used for multispectral matching [27]. However, MI is not differentiable and cannot be used to supervise a neural network. While MI can be approximated (e.g. MINE [28]), training a dense correspondence network using this approximation is inefficient and time consuming. Our network structure, combined with feature losses and cycle-consistency, is the best-of-both-worlds, being both spectrum-agnostic and computationally efficient.

A. Multispectral datasets

A current limitation for training and evaluating multispectral stereo correspondence techniques is the lack of suitable datasets. The majority of RGB-FIR datasets consist of pre-aligned image pairs, and therefore do not exhibit disparity. For distant scenes, which approximate orthographic projection, alignment can be made with a simple homography registration. Alternatively, images can be recorded with the same optical axis through the use of a beam-splitter. This removes the need for alignment/correspondence estimation and applies to RGB/NIR datasets EFPL [21], RANUS [29], and also RGB/FIR datasets KAIST [13], Coaxials [14], CAMEL [15].

Obtaining ground truth annotations for unaligned datasets is challenging, and human annotators rely on their familiarity with the visible spectrum. As a result, it is not always possible to distinguish objects, particularly in the case of FIR. The KAIST [13], LITIV [30] and PittsStereo [31] datasets have weak bounding box annotations or very sparse point correspondences only. The VAP dataset [32] provides synchronized RGB, FIR and depth images, with pedestrian segmentation in each modality. To our knowledge, the CATS dataset [33] is the only with dense RGB-FIR ground truth correspondence. However, the annotations are poorly registered,

making accurate quantitative evaluation impossible.

SODA from Li *et al.* uses image-to-image translation to synthesize FIR images from RGB, enabling the use of existing semantic labels [34]. However, these are limited to NIR only and exhibit artefacts. Image-to-image translation with FIR is more difficult, and it is impractical to train a correspondence network on the output. For example, a hot or cold car may have the same appearance in the visible spectrum. Given an RGB image, it is not always possible to assign a correct temperature using image-to-image translation alone.

The limited availability of annotated data has driven us to pursue a self-supervised approach. Data capture is much simpler without the requirement of ground truth, and furthermore the stereo pair does not need to be prealigned.

B. Machine learning and self-supervised training

FlowNet [8] from Dosovitskiy *et al.* is a supervised end-to-end Convolutional Neural Network (CNN) trained to estimate dense optical flow between RGB images. It was succeeded by FlowNet 2.0 [9] from Ilg *et al.*, who improve the inference speed and accuracy by stacking sub-networks and having a more detailed training regime. Although they are trained on synthetic datasets, they generalise well to real data such as KITTI [35]. Sun *et al.* make a direct comparison to FlowNet2 with PWC-Net [10], another supervised network, which is smaller and easier to train. Meister *et al.* present UnFlow, a self-supervised approach [11]. The authors use data losses between the warped image and the original, as well a consistency check between flow field directions. Wang *et al.* propose another self-supervised approach, UnDepthflow, which uses PWC-Net modules to isolate camera and scene motion [36]. UnFlow and UnDepthflow both achieve competitive accuracy on the KITTI dataset with other supervised networks. Although the use-case is similar to ours, the above approaches cannot work with images of different spectra.

A major factor in the success of self-supervised approaches is the use of cycle-consistency [37], [38]. In a cross-domain approach, Chen *et al.* employ adversarial losses to for bidirectional domain transfer [39]. This is similar to our approach in that the architecture consists of two halves which have mirrored operations. However, Chen *et al.* transfer between synthetic and real modalities (both RGB) as opposed to different spectra. Aguilera *et al.* learn a similarity measure for RGB to NIR using siamese networks [40]. Similarly, WILDCAT uses pseudo-siamese encoders to generate a shared latent space, to allow patch comparison between RGB and FIR [41]. Both of these require supervision to be trained, which limits their ability to generalise to different environments.

Generative approaches for RGB and NIR use a cycleGAN to match across generated stereo pairs for both spectra [42], [43]. Performance of these image-to-image translation approaches are subject to spectral similarity. Jeong *et al.* show that for unsynchronized pairs, most approaches score lower on FIR compared to NIR when making feature-based comparisons [43]. We use synchronized pairs and assume all geometric differences are caused by the camera viewpoints.

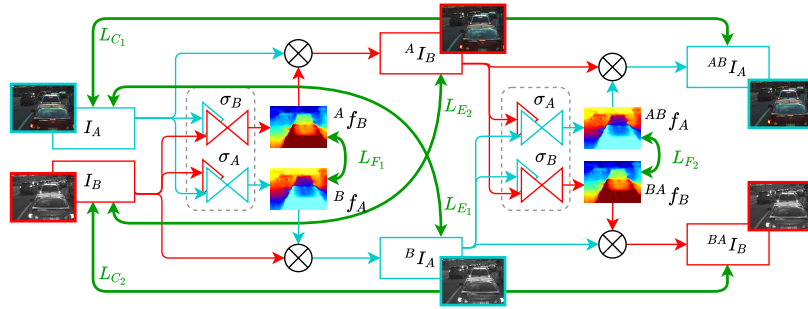


Fig. 1: Network overview. Red borders and arrows relate I_A 's point of view, blue relates to I_B 's. Green lines represent different losses (labeled). The \otimes operation represents warping an image I with a flow field f . Networks σ_A and σ_B relate to fig. 2; the same architecture but differently trained weights.

III. METHODOLOGY

Our goal is to enable dense correspondence between images with different spectra. We estimate 2D flow fields between image pairs in both directions. The data flow in our architecture is specifically designed to enable self-supervision.

Data flow. An overview of the losses and data flow is shown in Figure 1. For simplicity, we only describe the operations for one half of the full forward pass, as the other is identical but inverted. Modules σ_A and σ_B are flow estimation networks with the same architecture. They have two encoder arms, one for each spectrum without sharing the weights. We begin with two images, I_A and I_B . They describe the same scene but from different viewpoints and with different spectra. Both are provided to modules σ_A and σ_B . Flow estimation module σ_B estimates how I_A should be warped to align with I_B and returns a 2D flow field, which we refer to as ${}^A f_B$. Given ${}^A f_B$, the input image is warped using a differentiable sampling operation, producing ${}^A I_B$. This shares I_A 's spectrum, but its structure should align with I_B . We formalize this as

$$\begin{aligned} {}^A f_B &= \sigma_B(I_A, I_B) & {}^B f_A &= \sigma_A(I_B, I_A) \\ {}^A I_B &= I_A \otimes {}^A f_B & {}^B I_A &= I_B \otimes {}^B f_A, \end{aligned} \quad (1)$$

where we use the \otimes to represent a warping operation. We supervise the flow estimates by warping back to the original input, and evaluating cycle-consistency. ${}^A I_B$ is provided to other module σ_A , whose task is to estimate the flow field, ${}^{AB} f_A$, to effectively undo the previous warping operation. Similar to eq. (1),

$$\begin{aligned} {}^{AB} f_A &= \sigma_A({}^A I_B, {}^B I_A) & {}^{BA} f_B &= \sigma_B({}^B I_A, {}^A I_B) \\ {}^{AB} I_A &= {}^A I_B \otimes {}^{AB} f_A & {}^{BA} I_B &= {}^B I_A \otimes {}^{BA} f_B. \end{aligned} \quad (2)$$

${}^{AB} I_A$ should now be a perfect reconstruction of I_A and may be compared directly. Occlusions are handled implicitly by sampling interpolation, which is suitable for almost all cases.

Flow estimation module. The following relates to the module network structure in fig. 2. The modules σ simultaneously encode images to 6-layer pyramids. At each level, the features from one image are warped with the current optical flow estimate. These warped features are combined with the features from the other image into a shared cost

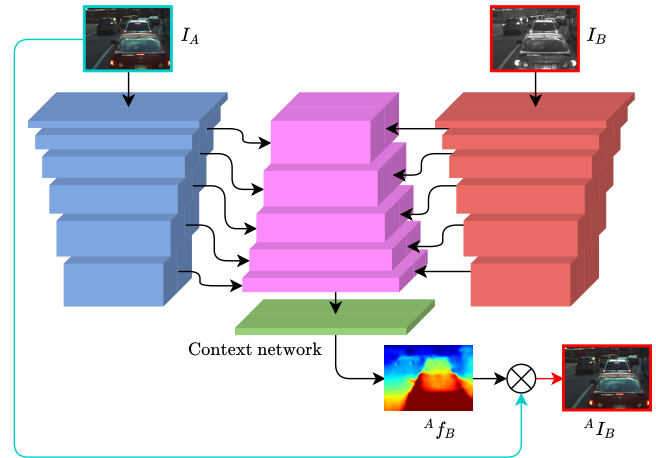


Fig. 2: Flow estimation module σ architecture. Each red/blue block consists of 3 convolution layers, downsizing by a factor of 2. Purple blocks consist of warping, cost volume, optical flow estimation, and transpose convolution layers. In contrast to PWC-Net, we separate encoder weights (red & blue).

volume. Each cost volume layer is fed through an optical flow estimator which passes the upscaled flow to the next layer, without sharing parameters between layers. Basing on PWC-Net [10], we make the following enhancements: (1) We separate the encoder parameters so that each encoder may adapt to a different spectrum, (2) We use suitable padding to overcome the restriction of tensor dimensions to multiples of 2^n convolutional layers, (3) We use strided transpose convolutions in the decoder in place of unweighted bilinear upsampling.

One forward pass of our full architecture uses each flow estimation module σ twice. The tasks of each encoder in σ are very similar. However, since our network is not targeted at any specific modality, the level of difference between aligned images is unknown. Hence, sharing weights between each encoder may impair learning. In fig. 1 we distinguish σ_A and σ_B as separate to indicate that the encoders do not share weights, but a shared decoder is used. In each instance we swap the encoders at runtime to accommodate the input and target spectra, enabling parallel but separate learning.

Cycle-consistency loss. If an image can be warped away and then back, and remain equal to its original, then both flow fields agree. The forward operation passes each image through both modules. σ_A always warps from I_B 's spectrum to I_A 's,

and σ_B from I_A 's to I_B 's. By enforcing cycle-consistency, σ_A and σ_B must work together. This encourages learned features to have a common representation between spectra. Direct photometric image comparison also enables a high degree of reconstruction accuracy. We introduce a simple mean-squared error over each pixel i for cycle-consistent supervision,

$$\begin{aligned} L_{C_1} &= \sum_{i=0}^{n-1} ({}^{AB}I_{A_i} - I_{A_i})^2 \\ L_{C_2} &= \sum_{i=0}^{n-1} ({}^{BA}I_{B_i} - I_{B_i})^2. \end{aligned} \quad (3)$$

Each cycle-consistency loss backpropagates through both modules and maintains balance between them.

Bidirectional flow-field loss. Cycle-consistency is one way to measure if two flow fields cancel out, but it only applies to the serial flow fields, e.g. ${}^A f_B$ with ${}^{AB} f_A$. Parallel flow fields, such as ${}^A f_B$ and ${}^B f_A$, should also cancel each other out. We introduce a loss on the flow-fields directly to also supervise these cases in tandem with the cycle-consistency loss. This is formalized as

$$L_{B_1} = |{}^A f_B + ({}^B f_A \otimes {}^A f_B)|, \quad (4)$$

where \otimes represents the warping function as an application of a flow field. The similar operations for its inverse, and for the second stage (${}^{AB} f_A$ and ${}^{BA} f_B$), have been omitted for conciseness.

Feature loss. The trivial solution to minimising both the cycle-consistency loss L_C (eq. (3)) and the bidirectional flow-field loss L_B (eq. (4)), is to generate flow fields which are entirely zero. In this case, the loss would be zero and ${}^{AB} I_A$ would perfectly match the original I_A . We use a cross-spectral feature loss after the first warping stage to discourage this.

${}^A I_B$ should align with I_B , but their spectra are different. To verify that the structure aligns while trying to ignore spectrum differences, we use a perceptual loss [44],

$$\begin{aligned} L_{F_1} &= \frac{1}{CHW} \|\phi({}^A I_B) - \phi(I_B)\|_2^2 \\ L_{F_2} &= \frac{1}{CHW} \|\phi({}^B I_A) - \phi(I_A)\|_2^2, \end{aligned} \quad (5)$$

where $\phi(x)$ are the layer activations of image x , and $C \times H \times W$ are the dimensions of that convolutional layer. In the original work, Johnson *et al.* defines features ϕ at layer `relu3_3` as suitable for matching content, with all layers suitable for style. Roughly speaking, the deeper the layer at which latent features are extracted, the less the features will represent the original spectrum which is an important consideration for us. However, due to the receptive field of the network, the deep layers can decrease precision. It is important to consider the significant reduction in dimensions when trying to precisely align a scene. For our problem, the best layer to use depends on the proximity of the spectra being compared. For RGB-RGB correspondence, we found the activations of shallow layers suitable ($\phi \in \text{relu2}_2$).

Regularisation loss. Correct estimated flow fields naturally

warp some pixels out of the image boundaries. In the absence of ground truth, gradients from the other losses cannot correct for pixels outside the image boundaries. Flow fields may be generated with extreme values leading to an unrecoverable situation. A low learning rate discourages this possibility but impairs the training, increasing the likelihood to get stuck in local minima. To optimize the learning rate and achieve stability, we introduce a Huber loss L_R to act against unreasonable flow field magnitudes. For flow fields ${}^A f_B$ and ${}^B f_A$:

$$L_{R_f} = \begin{cases} \frac{1}{2} x_f^2 & \text{for } |x_f| < 1, \\ |x_f| - \frac{1}{2} & \text{otherwise} \end{cases} \quad (6)$$

$$x_f = \max(0, f - m) \quad \forall f \in \{{}^A f_B, {}^B f_A\},$$

where m is defined as the maximum disparity in pixels which can be expected. This is set to the disparity of the closest object in the dataset. The regularisation loss only starts to affect flow field values above m , which we typically set to 10% of the image width. The squared term at flow field, f , lower than, x , does not have a strong effect and permits a small amount of overshoot. At higher values, the loss increases linearly to avoid exploding gradients. In summary, pixels moved far from the image boundaries incur a high loss. We only apply this on the flow fields produced at the first warping stage, as ${}^{AB} f_A$ and ${}^{BA} f_B$ are supervised by the cycle-consistency loss.

Smoothing loss. The flow fields themselves should be smooth except at depth discontinuities. By enforcing uniformity in homogenous areas, the accuracy and visual quality of the warped images improves, thereby allowing a closer match for the cycle-consistency L_C . We introduce a term to penalize flow-field gradients which do not coincide with image gradients,

$$\begin{aligned} L_{S_1} &= \left| h({}^A f_B) * (1 - h(I_A)) \right| \\ L_{S_2} &= \left| h({}^B f_A) * (1 - h(I_B)) \right| \\ L_{S_3} &= \left| h({}^{AB} f_A) * (1 - h({}^A I_B)) \right| \\ L_{S_4} &= \left| h({}^{BA} f_B) * (1 - h({}^B I_A)) \right|, \end{aligned} \quad (7)$$

where $h(x)$ retrieves the Sobel gradient magnitude of x , normalised between 0 and 1.

Combined loss function. We define the overall loss as a balanced combination,

$$L_{total} = \alpha L_C + \beta L_B + \gamma L_F + \delta L_R + \epsilon L_S, \quad (8)$$

where $\alpha, \beta, \gamma, \delta, \epsilon$ are training weights for each loss type. For simplicity of notation, we refer to each loss type as the sum of its components, e.g. $L_C = L_{C_1} + L_{C_2}$.

IV. EXPERIMENTAL RESULTS

In this section we evaluate our approach under different combinations of spectra. We show that our approach is both capable of producing competitive state-of-the-art results, as

TABLE I: RGB-FIR evaluation results for precision (Pr), recall (Re) and F1 score (F_1) on the VAP dataset [32]. Scene 2 is omitted, as reported by other authors. As reference, the top three approaches (the highest performers for this dataset) estimate segmentation masks on registered images. The bottom three register ground truth masks using image data. The repeated baseline values represent the same evaluation. PWC-Net in parentheses was designed for fully-supervised RGB-RGB only.

Method	Metric	Scene 1		Scene 3		Overall		
		RGB	FIR	RGB	FIR	RGB	FIR	Mean
St-Charles [45]	Pr	0.820	0.755	0.716	0.514	0.768	0.635	0.701
	Re	0.810	0.975	0.688	0.969	0.749	0.972	0.861
	F_1	0.815	0.851	0.702	0.672	0.758	0.762	0.760
GrabCut [46]	Pr	0.685	0.808	0.653	0.847	0.669	0.828	0.748
	Re	0.759	0.896	0.929	0.916	0.844	0.906	0.875
	F_1	0.721	0.850	0.737	0.880	0.744	0.865	0.804
St-Charles [47]	Pr	0.894	0.860	0.788	0.749	0.841	0.804	0.821
	Re	0.902	0.901	0.918	0.937	0.910	0.919	0.914
	F_1	0.898	0.880	0.848	0.833	0.873	0.857	0.866
Baseline	Pr	0.536	0.536	0.559	0.559	0.548	0.548	0.548
	Re	0.525	0.525	0.535	0.535	0.530	0.530	0.530
	F_1	0.529	0.529	0.545	0.545	0.537	0.537	0.537
(PWC-Net) [10]	Pr	0.448	0.337	0.693	0.129	0.571	0.233	0.402
	Re	0.497	0.279	0.729	0.091	0.613	0.185	0.399
	F_1	0.467	0.294	0.702	0.095	0.585	0.195	0.390
Ours	Pr	0.700	0.799	0.720	0.694	0.750	0.707	0.728
	Re	0.569	0.768	0.518	0.779	0.669	0.649	0.659
	F_1	0.622	0.781	0.595	0.734	0.702	0.665	0.683

well as generalising to more than a single fixed pair of spectra. In order to demonstrate this, we evaluate three different cases:

- 1) **RGB-FIR**: We show that in addition to the previous two scenarios, our approach can recover flow fields and solve the correspondence problem between thermal and visible images. We evaluate the task of transferring annotations to the other spectrum.
- 2) **RGB-NIR**: We show that without modifying the architecture, our approach can solve correspondence between the visible and NIR spectra without ground truth supervision. We demonstrate its effectiveness with competitive evaluation scores on an automotive stereo dataset.
- 3) **RGB-RGB**: We show that our self-supervised training approach is not limited to non-visible spectra, and it can still achieve competitive results when compared to other supervised state-of-the-art approaches.

Each of these instances do not require a network modification, demonstrating the flexibility of our method. For each, we train with a batch size of 8 and SGD optimizer, and optimize our hyperparameters with optuna [48]. For RGB-RGB, the learning rate was set to $4.3e-05$, with losses weighted as $\alpha = 3.4e-1$, $\beta = 3.6e-4$, $\gamma = 6.7e-1$, $\delta = 6.9e-2$, and $\epsilon = 2.7e-1$. Our approach is able to recover both the vertical and horizontal disparity, whereas many benchmarks are restricted to horizontal disparity.

A. RGB-FIR evaluation

We first evaluate on the VAP dataset [32], using the RGB and FIR modalities only. Ground truth correspondence is not provided. We instead use human body segmentation as a proxy, by warping the ground truth segmentation mask from one spectrum to the other, and comparing against the

TABLE II: PittsStereo RGB-NIR evaluation results. Score for each material category is the RMSE of disparity in pixels. PWC-Net in parentheses was designed for fully-supervised RGB-RGB only.

Method	Bag	Cloth.	Com.	Glass	Glossy	Light	Skin	Veg.	Mean
CMA [49]	4.63	6.42	1.60	2.55	3.86	5.17	3.39	4.42	4.00
ANCC [50]	2.57	2.85	1.36	2.27	2.41	2.43	2.32	4.82	2.63
DASC [23]	1.33	0.80	0.82	1.50	1.82	1.24	1.59	1.09	1.28
Liang [42]	0.80	0.98	0.68	0.67	1.05	0.80	1.04	0.68	0.84
Zhi [31]	0.90	1.22	0.51	1.05	1.57	1.08	1.01	0.69	1.00
(PWC-Net) [10]	4.88	5.67	11.20	6.06	3.33	3.92	7.00	13.5	6.95
Ours	0.91	0.90	0.64	1.18	1.49	1.00	1.47	1.10	1.08




Fig. 3: Qualitative results. **Column 1**: Example RGB and FIR input images I_A and I_B . **Column 2**: Original masks from $[I_A, I_B]$, and used for baseline scores. **Column 3**: Warped RGB with FIR masks $[^A I_B, I_B]$. **Column 4**: Warped FIR with RGB masks $[^B I_A, I_A]$.

other mask. This process is significant because it provides an easier way to obtain annotations on challenging modalities. A state-of-the-art RGB segmentation algorithm may be used to generate accurate masks which can then be densely registered to another viewpoint and spectrum.

Training. We break the left \rightarrow right RGB \rightarrow FIR stereo assumption by horizontally flipping both images in each pair with a probability of 0.5 during training. With data augmentation consistent between sample pairs, the static backgrounds may lead to overfitting to the camera distortion. To prevent this we randomly crop pairs independently, i.e. in a single pair, the RGB may be cropped to a slightly different region as the FIR. This introduces unpaired regions in each image, which are masked from the loss. For each training sample, the network is forced to warp a region of the image, in any direction, to align and minimize the training losses. This can only be achieved by learning how to match features across different spectra.

Results. The three approaches at the top of table I estimate the masks on registered images, whereas we use the provided masks with unregistered images. This distinction is important when comparing scores. When information is only present in one modality, a correct flow field will not give a perfect score, whereas segmentation on those images will reflect what is present. In the qualitative comparison in fig. 3, both RGB \rightarrow FIR and FIR \rightarrow RGB flow fields successfully align the masks (note that the masks are not seen by our network). However, poor FIR image quality around the borders results the bottom of its mask being cut off which harms our evaluation scores.

The ‘Baseline’ row in table I represents the scores on unwarped, unregistered images (our evaluation’s input). The majority of the masks already overlap but are not aligned. The decreased scores of PWC-Net [10] shows the difficulty of the task, whereas ours improves the stereo registration.




Fig. 4: Qualitative results on the RGB-NIR evaluation dataset. Brightness and contrast increased for clarity. **Column 1:** Input RGB images I_A . **Column 2:** Target NIR images I_B . **Column 3:** Flow fields ${}^A f_B$ & ${}^B f_A$ applied to input images. **Column 4:** Flow fields ${}^{AB} f_A$ & ${}^{BA} f_B$ to return warped images to their input state.

B. RGB-NIR evaluation

We also evaluate on the PittsStereo RGB-NIR dataset (automotive) [31]. The test set is annotated with a small number of point correspondences for each image pair, labelled with a material category including vegetation, glass, and lights.

Training. The dataset only exhibits horizontal disparity. No 2D NIR datasets could be found. We do not restrict our network to 1D; it learns to ignore the vertical component. We follow a similar training regime as for the RGB-FIR (section IV-A). Since spectral bands are closer, we are able to augment the data by jittering the RGB values by 3%.

Results. The RMSE error for each material type can be found in Table II. Our approach clearly outperforms existing feature based methods. The state-of-the-art deep learning approaches are specifically targeted at RGB-NIR. Our performance is comparable in spite of this, noting that errors are of the magnitude of a single pixel, accounting for $< 0.2\%$ of the image width. Figure 4 shows the flow fields estimated from both spectra and viewpoints agree, with clear object boundaries.

Both neural network approaches, from Zhi *et al.* [31] and Liang *et al.* [42] leverage image-to-image translation. This is possible because the majority of materials in the NIR spectrum closely resemble their grayscale counterparts in the visible image. However, this is an inherent restriction preventing their transfer to other spectra.

C. RGB-RGB evaluation

We evaluate on the KITTI 2015 scene-flow dataset [35]. Despite the fact that all images are RGB, our design choices made to avoid direct left-right photometric image comparison do not cause an issue for same-spectrum correspondence

TABLE III: KITTI optical flow evaluation 2015 results. ‘AEE’ is the Average flow field End-Point Error in pixels. ‘Error’ is the percentage of erroneous pixels, classed as having an end-point error of ≥ 3 px and $\geq 5\%$.

Method	Self-supervised	AEE	Error
FlowNet2 [9]	✗	2.30	10.41%
UnFlow-CSS-ft [11]	✗	1.86	11.11%
PWC-Net [10]	✗	2.16	9.60%
LiteFlowNet [51]	✗	5.58	9.38%
LiteFlowNet2 [52]	✗	4.32	7.62%
UnFlow-CSS [11]	✓	8.10	23.30%
UnOS (FlowNet only) [36], [53]	✓	7.88	23.75%
Wang <i>et al.</i> [54]	✓	8.88	31.20%
Janai <i>et al.</i> [55]	✓	6.59	22.94%
DF-Net [56]	✓	8.98	25.70%
Ours	✓	4.53	21.12%




Fig. 5: KITTI 2015 evaluation (test set). **Row 1:** Input images I_A . **Row 2:** Estimated flow fields ${}^A f_B$. **Row 3:** Error heatmaps.

estimation. It should be noted that we do not make use of techniques specific to RGB-RGB correspondence estimation, nor do we estimate additional constraints such as vehicle egomotion. We follow the same training regime as for the RGB-NIR in section IV-B.

Results. We compare against popular recent approaches and present quantitative results in table III, including supervised approaches for context. Compared with other recent self-supervised optical flow approaches, we achieve the lowest error on the test set. This is even without using photometric loss or explicit occlusion reasoning. None of the compared approaches are capable of cross-spectral matching. Example qualitative flow fields and error heatmaps are shown in fig. 5.

V. CONCLUSION

In summary, we have presented a correspondence estimation method agnostic to the spectra that it is operating in. We have shown that our approach tackles RGB-FIR, RGB-NIR, and RGB-RGB. This is all while achieving results comparable to state-of-the-art methods, many of which that have been fine-tuned to their specific modality. More generally we have demonstrated that, at a fundamental level, cross-spectral correspondence estimation is a problem that can be solved in a generic way. Additionally, our approach demonstrates that this can all be done in a self-supervised manner. This enables correspondence estimation algorithms to be trained in domains that have very little training data and virtually no annotations. With this, we hope this that the interest in the cross-spectral domain continues to grow, to expand the currently limited variety of datasets and algorithms.

REFERENCES

- [1] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, ISBN: 0521540518, 2004.
- [2] J. M. Brady, "ASSET-2: Real-Time Motion Segmentation and Shape Tracking," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, no. 8, pp. 814–820, 1995.
- [3] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [4] B. D. Lucas and T. Kanade, "An Iterative Image Registration Technique with an Application to Stereo Vision," *DARPA Image Understanding Workshop*, pp. 121–130, 1981.
- [5] P. Besl and N. McKay, "A Method for Registration of 3-D Shapes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 2, pp. 239–256, 1992.
- [6] C. Liu, J. Yuen, and A. Torralba, "Sift flow: Dense correspondence across scenes and its applications," *Dense Image Correspondences for Computer Vision*, pp. 15–49, 2015.
- [7] E. Tola, V. Lepetit, and P. Fua, "DAISY: An efficient dense descriptor applied to wide-baseline stereo," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 5, pp. 815–830, 2010.
- [8] A. Dosovitskiy, P. Fischery, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. V. D. Smagt, D. Cremers, and T. Brox, "FlowNet: Learning optical flow with convolutional networks," *Proceedings of the IEEE International Conference on Computer Vision*, vol. 2015 Inter, pp. 2758–2766, 2015.
- [9] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox, "FlowNet 2.0: Evolution of optical flow estimation with deep networks," *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, vol. 2017-Janua, pp. 1647–1655, 2017.
- [10] D. Sun, X. Yang, M. Y. Liu, and J. Kautz, "PWC-Net: CNNs for Optical Flow Using Pyramid, Warping, and Cost Volume," *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, vol. D, pp. 8934–8943, 2018.
- [11] S. Meister, J. Hur, and S. Roth, "UnFlow: Unsupervised learning of optical flow with a bidirectional census loss," *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pp. 7251–7259, 2018.
- [12] X. Liu, Y. Ai, B. Tian, and D. Cao, "Robust and Fast Registration of Infrared and Visible Images for Electro-Optical Pod," *IEEE Transactions on Industrial Electronics*, vol. 66, no. 2, pp. 1335–1344, 2019.
- [13] S. Hwang, J. Park, N. Kim, Y. Choi, and I. S. Kweon, "Multispectral pedestrian detection: Benchmark dataset and baseline," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 07-12-June, pp. 1037–1045, 2015.
- [14] A. Okazawa, T. Takahata, and T. Harada, "Simultaneous transparent and non-transparent object segmentation with multispectral scenes," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2019)*, 2019, pp. 4977–4984.
- [15] E. Gebhardt and M. Wolf, "CAMEL Dataset for Visual and Thermal Infrared Multiple Object Detection and Tracking," *Proceedings of AVSS 2018 - 2018 15th IEEE International Conference on Advanced Video and Signal-Based Surveillance*, pp. 1–6, 2019.
- [16] J. W. Davis and M. A. Keck, "A Two-Stage Template Approach to Person Detection in Thermal Imagery," in *Proc. Workshop on Applications of Computer Vision*, 2005.
- [17] M. M. Zhang, J. Choi, K. Daniilidis, M. T. Wolf, and C. Kanan, "VAIS: A dataset for recognizing maritime imagery in the visible and infrared spectrums," *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, vol. 2015-October, pp. 10–16, 2015.
- [18] C. Li, H. Cheng, S. Hu, X. Liu, J. Tang, and L. Lin, "Learning Collaborative Sparse Representation for Grayscale-Thermal Tracking," *IEEE Transactions on Image Processing*, vol. 25, no. 12, pp. 5743–5756, 2016.
- [19] C. Li, X. Liang, Y. Lu, N. Zhao, and J. Tang, "RGB-T Object Tracking: Benchmark and Baseline," pp. 1–13, 2018.
- [20] P. Ricaurte, C. Chilán, C. A. Aguilera-Carrasco, B. X. Vintimilla, and A. D. Sappa, "Feature point descriptors: Infrared and visible spectra," *Sensors (Switzerland)*, vol. 14, no. 2, pp. 3690–3701, 2014.
- [21] M. Brown and S. Süssstrunk, "Multispectral SIFT for scene category recognition," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2011, pp. 177–184.
- [22] S. Saleem and R. Sablatnig, "A robust SIFT descriptor for multispectral images," *IEEE Signal Processing Letters*, vol. 21, no. 4, pp. 400–403, 2014.
- [23] S. Kim, D. Min, B. Ham, S. Ryu, M. N. Do, and K. Sohn, "DASC: Dense adaptive self-correlation descriptor for multi-modal and multi-spectral correspondence," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 07-12-June, pp. 2103–2112, 2015.
- [24] K. Hajebi and J. S. Zelek, "Structure from infrared stereo images," *Proceedings of the 5th Canadian Conference on Computer and Robot Vision, CRV 2008*, pp. 105–112, 2008.
- [25] T. Mouats, N. Aouf, L. Chermak, and M. A. Richardson, "Thermal stereo odometry for UAVs," *IEEE Sensors Journal*, vol. 15, no. 11, pp. 6335–6347, 2015.
- [26] Y. Li and R. L. Stevenson, "Multimodal Image Registration With Line Segments by Selective Search," *IEEE Transactions on Cybernetics*, vol. 47, no. 5, pp. 1285–1298, 2017.
- [27] P. Viola and W. M. Wells, "Alignment by Maximization of Mutual Information," *International Journal of Computer Vision*, vol. 24, no. 2, pp. 137–154, 1997.
- [28] M. I. Belghazi, A. Baratin, S. Rajeswar, S. Ozair, Y. Bengio, A. Courville, and R. D. Hjelm, "Mutual Information Neural Estimation," *35th International Conference on Machine Learning, ICML 2018*, vol. 2, pp. 864–873, 2018.
- [29] G. Choe, S. H. Kim, S. Im, J. Y. Lee, S. G. Narasimhan, and I. S. Kweon, "RANUS: RGB and NIR urban scene dataset for deep scene parsing," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1808–1815, 2018.
- [30] G. A. Bilodeau, A. Torabi, P. L. St-Charles, and D. Riahi, "Thermal-visible registration of human silhouettes: A similarity measure performance evaluation," *Infrared Physics and Technology*, vol. 64, pp. 79–86, 2014.
- [31] T. Zhi, B. R. Pires, M. Hebert, and S. G. Narasimhan, "Deep Material-Aware Cross-Spectral Stereo Matching," in *IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018.
- [32] C. Palmero, A. Clapés, C. Bahnsen, A. Møgelmoose, T. B. Moeslund, and S. Escalera, "Multi-modal RGB–Depth–Thermal Human Body Segmentation," *Int. J. Comput. Vis.*, vol. 118, pp. 217–239, 2016.
- [33] W. Treible, P. Saponaro, S. Sorensen, A. Kolagunda, M. O'Neal, B. Phelan, K. Sherbondy, and C. Kambhamettu, "CATS: A color and thermal stereo benchmark," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2961–2969.
- [34] C. Li, W. Xia, Y. Yan, B. Luo, and J. Tang, "Segmenting Objects in Day and Night: Edge-Conditioned CNN for Thermal Image Semantic Segmentation," vol. 1, pp. 1–12, 2019.
- [35] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [36] Y. Wang, Z. Yang, P. Wang, Y. Yang, C. Luo, and W. Xu, "Joint Unsupervised Learning of Optical Flow and Depth by Watching Stereo Videos," oct 2018.
- [37] C. Godard, O. Mac Aodha, and G. J. Brostow, "Unsupervised monocular depth estimation with left-right consistency," *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, vol. 2017-Janua, pp. 6602–6611, 2017.
- [38] C. Godard, O. Mac Aodha, M. Firman, and G. Brostow, "Digging Into Self-Supervised Monocular Depth Estimation," no. 1, 2018.
- [39] Y.-C. Chen, Y.-Y. Lin, M.-H. Yang, and J.-B. Huang, "CrDoCo: Pixel-Level Domain Transfer With Cross-Domain Consistency," pp. 1791–1800, 2020.
- [40] C. A. Aguilera, F. J. Aguilera, A. D. Sappa, and R. Toledo, "Learning Cross-Spectral Similarity Measures with Deep Convolutional Neural Networks," *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pp. 267–275, 2016.
- [41] W. Treible, P. Saponaro, Y. Liu, A. Das Gupta, V. Veerendraveer, S. Sorensen, and C. Kambhamettu, "CATS 2: Color And Thermal Stereo Scenes with Semantic Labels," in *CVPR Workshops*, 2019, pp. 0–4.
- [42] M. Liang, X. Guo, H. Li, X. Wang, and Y. Song, "Unsupervised Cross-spectral Stereo Matching by Learning to Synthesize," 2019.
- [43] S. Jeong, S. Kim, K. Park, and K. Sohn, "Learning to Find Unpaired Cross-Spectral Correspondences," *IEEE Transactions on Image Processing*, vol. 28, no. 11, pp. 5394–5406, 2019.
- [44] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 9906 LNCS, pp. 694–711, 2016.
- [45] P. L. St-Charles, G. A. Bilodeau, and R. Bergevin, "Universal Background Subtraction Using Word Consensus Models," *IEEE Trans. Image Process.*, vol. 25, no. 10, pp. 4768–4781, 2016.

- [46] C. Rother, V. Kolmogorov, and A. Blake, "'GrabCut'- Interactive foreground extraction using iterated graph cuts," *ACM Trans. Graph. – Proc. ACM SIGGRAPH 2004*, vol. 23, no. 3, p. 309, 2004.
- [47] P. L. St-Charles, G. A. Bilodeau, and R. Bergevin, "Online Mutual Foreground Segmentation for Multispectral Stereo Videos," *International Journal of Computer Vision*, vol. 127, no. 8, pp. 1044–1062, 2019.
- [48] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A next-generation hyperparameter optimization framework," *Proceedings of the 25rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 2623–2631, 2019.
- [49] W. C. Chiu, U. Blanke, and M. Fritz, "Improving the kinect by cross-modal stereo," *BMVC 2011 - Proc. Br. Mach. Vis. Conf. 2011*, pp. 1–10, 2011.
- [50] Yong Seok Heo, Kyong Mu Lee, and Sang Uk Lee, "Robust Stereo Matching Using Adaptive Normalized Cross-Correlation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 4, pp. 807–822, apr 2011.
- [51] T. W. Hui, X. Tang, and C. C. Loy, "LiteFlowNet: A Lightweight Convolutional Neural Network for Optical Flow Estimation," *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, pp. 8981–8989, 2018.
- [52] T.-W. Hui, X. Tang, and C. C. Loy, "A Lightweight Optical Flow CNN - Revisiting Data Fidelity and Regularization," pp. 1–15, 2019.
- [53] Y. Wang, P. Wang, Z. Yang, C. Luo, Y. Yang, and W. Xu, "UnOS: Unified unsupervised optical-flow and stereo-depth estimation by watching videos," *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, vol. 2019-June, pp. 8063–8073, 2019.
- [54] Y. Wang, Y. Yang, Z. Yang, L. Zhao, P. Wang, and W. Xu, "Occlusion Aware Unsupervised Learning of Optical Flow," nov 2017.
- [55] J. Janai, F. Güney, A. Ranjan, M. Black, and A. Geiger, "Unsupervised Learning of Multi-Frame Optical Flow with Occlusions," *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 11220 LNCS, pp. 713–731, 2018.
- [56] Y. Zou, Z. Luo, J.-b. Huang, and I. Separate, "DF-Net: Unsupervised Joint Learning of Depth and Flow using Cross-Task Consistency," *Eur. Conf. Comput. Vis.*, 2018.