

ChiTransformer: Towards Reliable Stereo from Cues

Qing Su
 Georgia State University
 qsu3@gsu.edu

Shihao Ji
 Georgia State University
 sji@gsu.edu

Abstract

Current stereo matching techniques are challenged by restricted searching space, occluded regions and sheer size. While single image depth estimation is spared from these challenges and can achieve satisfactory results with the extracted monocular cues, the lack of stereoscopic relationship renders the monocular prediction less reliable on its own especially in highly dynamic or cluttered environments. To address these issues in both scenarios, we present an optic-chiasm-inspired self-supervised binocular depth estimation method, wherein vision transformer (ViT) with a gated positional cross-attention (GPCA) layer is designed to enable feature-sensitive pattern retrieval between views, while retaining the extensive context information aggregated through self-attentions. Monocular cues from a single view are thereafter conditionally rectified by a blending layer with the retrieved pattern pairs. This crossover design is biologically analogous to the optic-chiasma structure in human visual system and hence the name, *ChiTransformer*. Our experiments show that this architecture yields substantial improvements over state-of-the-art self-supervised stereo approaches by 11%, and can be used on both rectilinear and non-rectilinear (e.g., fisheye) images.

1. Introduction

In the context of computer vision, nowadays almost all mainstream depth estimation methods are deep learning based and can be roughly categorized into two prevalent methodologies, namely, stereo matching and monocular depth estimation. Stereo matching has traditionally been the most investigated area due to its strong connection to the human visual system. The task is to find or estimate the correspondences of all the pixels in two rectified images [3, 5, 54]. Virtually, all the current works resort to convolutional neural network (CNN) based methods to calculate the matching cost since it was first introduced to the task by [16, 71] in 2015. Following the seminal work of FlowNet [16], more than 150 papers have been published

¹<https://github.com/ISL-CV/ChiTransformer.git>

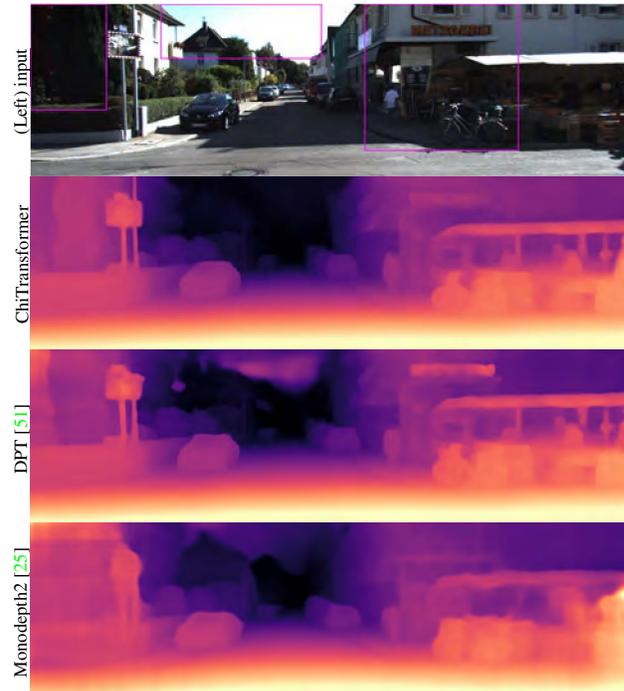


Figure 1. Depth estimation from cues. Monocular depth estimation relies on depth cues to make prediction. While our self-supervised stereo method, *ChiTransformer*, leverages retrieved and rectified depth cues from a stereo to make superior predictions with context consistency.

using CNN-related methods [39], pushing the performance forward by more than 50%. Some deep-seated issues such as thin structures, large texture-less areas, and occlusions have been mitigated or addressed over time [32, 73]. So far, stereo matching is the most adopted technique in majority of passive stereo applications.

However, the applications that entail depth estimation grow increasingly demanding as visual systems are greatly downsized and installed on platforms with higher mobility (e.g., UAV, commercial robots). This indicates a more congested, cluttered and dynamic operating environment where the once side issues become major ones, i.e., large disparity,

severe occlusion and non-rectilinear images that might be involved. Therefore, most existing stereo matching methods are not set up for this new trend and fail to address these issues properly.

On the other hand, monocular depth estimation (MDE) is spared from these challenges as depth is estimated from a single view. Following [19], current works leverage deep models to derive more descriptive cues to achieve superior predictions. More recent works focus on fusing multi-scale information to further improve the pixel level depth estimation [40, 44]. Lately, vision transformer is exploited in the task and yields globally organized and coherent predictions with finer granularity [6, 51]. State-of-the-art MDE methods can achieve impressive results with relative accuracy $\delta^3 > 0.99$ with supervised training [21, 26, 43, 69]. However, the reliability of MDE estimation is essentially based on the assumption that scenes in the real world are mostly regular. Therefore, due to the lack of stereo relation, the MDE is more delimited to its training dataset and susceptible to “unfamiliar” scenes. This renders the MDE alone not reliable in safety-critical applications, such as autonomous driving and visual-aided UAV.

From the discussion above, we can see the limitations and advantages of stereo matching and MDE are complementary. Therefore, in this paper we propose a novel method that jointly addresses their limitations by crossing over the stereo and MDE approaches such that stereo information can be injected into the MDE process to rectify and improve the estimation made from depth cues.

We therefore introduce ChiTransformer, an optic-chiasm-inspired self-supervised binocular depth estimation network. ChiTransformer adopts the recent vision transformer (ViT) [15] as backbone and extends the encoder-only transformer to an encoder-decoder structure similar to the ones for natural language processing [14, 59]. Rather than end-to-all connections, interleaved connections are employed for cross-attention to enable progressive instillation of the encoded depth cues from a nearby view to the master view in a self-regressive process. Our main contribution is the design of a retrieval cross-attention layer. Instead of attending over multi-level contextual relation of the encodings in the regular multi-head attention (MHA), the cross-attention mechanism of ChiTransformer aims to retrieve depth cues with strong contextual and feature coincidence from the other view. To achieve this, we condition the initial state (query) with a self-adjoint operator without breaking the convergence rule of modern Hopfield network [50]. The positive-definite operator is spectrally decomposed to enable polarized attention within the encoded feature space to emphasize on certain cues while preserving as much of the original information as possible. We show that this design facilitates reliable retrieval and leads to finer feature-consistent details on top of the globally coherent es-

imation. Moreover, the model can be further extended to non-rectilinear images such as fisheye by using a gated positional embedding [12]. We model the epipolar geometry with learnable quadratic polynomials of relative positions. Considering the per-pixel labeled data is challenging to acquire at scale and let alone for the non-rectilinear images, we choose to train the model with self-supervised learning strategy tailored from the work [25].

In contrast to traditional stereo methods, our approach foregoes pixel-level matching optimization but leverages the context-infused depth cues of both views to improve the overall depth prediction. With a global receptive field, ChiTransformer is not only not restricted to certain epipolar geometry, e.g., the horizontal collinear epipolar lines of rectified regular stereo pairs, but also able to treat large disparity. Furthermore, with the inherent capability of depth estimation within a single image, estimation at large occluded area can be properly handled rather than being masked out, interpolated, or left untreated. Enhanced from current MDE methods, our approach provides reliable prediction with guided cues in stereo pair which makes ChiTransformer more suitable for complex and dynamic environments.

Experiments are conducted on depth estimation tasks that provide stereo pairs. Our result shows that ChiTransformer delivers an improvement of more than 11% compared to the top-performing self-supervised stereo methods. The architecture is also tested on stereo tasks to evaluate the gain brought by stereo cues and the underlying reliability yielded by the instilled stereo information. To show the potential of ChiTransformer in non-rectilinear images, we train our model to predict the distances on the translated synthetic fisheye sequences from [18] and achieve visually satisfactory results.

2. Related Work

Since the publications of [19, 20], the end-to-end trainable CNN-based models have been the prototype architecture for dense depth [24, 25, 52] or disparity estimation [29, 31, 56, 67]. The principal idea is to leverage learned representation to improve matching cost [36, 54] or depth cues [7] with contextual information of appropriately large local area. The prevalent encoder-decoder structure enables progressive down- and up-sampling of representation at different scales [10, 17, 44, 66, 74], and the intermediate results from previous layers are often reused to recover fine-grained estimation while ensuring sufficiently large context.

After showing exemplary performance on a broad range of NLP tasks, attention or in particular transformer has demonstrated competitive or superior capability in vision tasks such as image recognition [15, 57], object detection [9, 76], semantic segmentation [68], super-resolution [65], image inpainting [72], image generation [53], text-image synthesis [1], etc. The successes also sparked interest in the

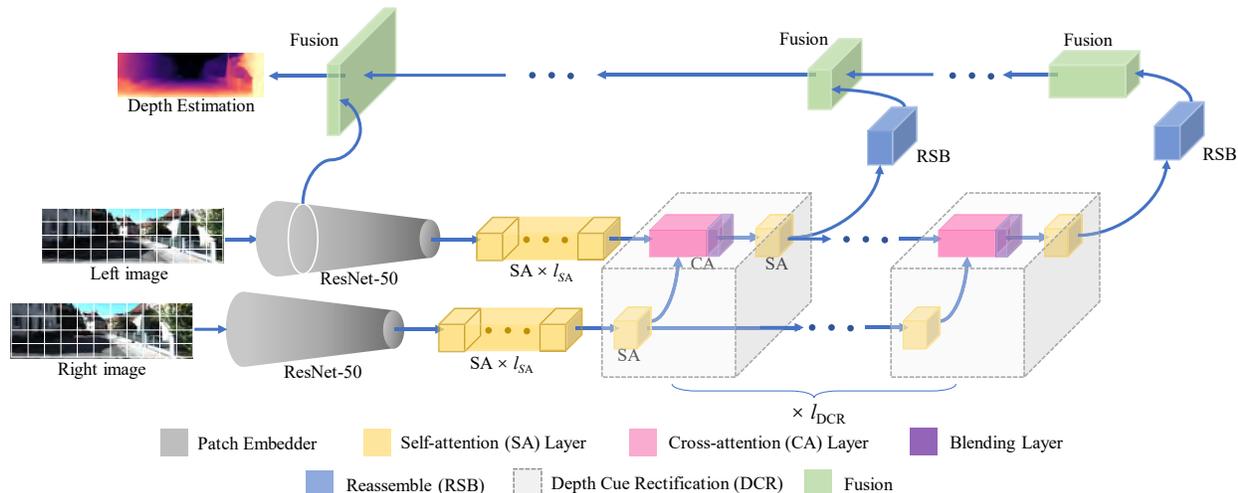


Figure 2. Architecture of ChiTransformer. A stereo pair (left: master, right: reference) is initially embedded into tokens through a Siamese ResNet-50 tower. The 2D-organized tokens from the two images are flattened and then augmented with learnable positional embeddings and an extra class token, respectively. Then tokens are fed into two self-attention (SA) stacks of size l_{SA} in parallel. After that, tokens are fed into a series ($\times l_{DCR}$) of depth rectification blocks (DCR) in each of which tokens of reference image go through an SA layer while tokens of master go through a polarized cross-attention (CA) layer, followed by an SA layer. In the polarized CA layer, relevant tokens from the output of reference SA are fetched to rectify the master’s depth cues. Tokens from different stages are afterwards reassembled into an image-like arrangement at multiple resolution (blue) and progressively fused and up-sampled through fusion block to generate a fine-grained depth estimation.

community of stereo and depth estimation. [63] leveraged cascaded attentions to calculate the matching cost along the epipolar lines and achieved competitive results among self-supervised stereo matching methods [2, 34, 42, 75]. More recently, vision transformer was leveraged in place of convolution network as backbone for dense depth prediction in [51] and achieved a significant improvement by 28% compared to the state-of-the-art convolutional counterparts. A mini-ViT block [6] is employed in the refinement stage to facilitate the adaptive depth bin calculation, and the work tops KITTI [23] and NYUv2 [55] leaderboards. Inspired by [51], our method leverages the capability of ViT in learning long range complex context information to rectify depth cues instead of performing stereo matching.

Most works discussed above are fully supervised, which necessitate pixel-wise labeled ground truth for training. However, it is challenging to acquire dense annotations at scale in many real-world settings. One of the workarounds is to adopt self-supervised learning. For stereo self-supervised training, usually pixel disparities of synchronized stereo pairs are predicted [2, 34, 42, 66, 75], while for monocular self-supervised training, not only depth but also camera pose has to be estimated to help reconstruct the image and constrain the estimation network [8, 24, 60, 70, 75]. Considering the versatility and the potential application environment of our method, we choose self-supervised training for ChiTransformer.

3. Method

This section introduces the overall architecture of the ChiTransformer with elaboration on the key building

blocks. We follow the configuration of the vision transformer [15] as the backbone and maintain the prevalent overall encoder-decoder structure because of their repeatedly verified success in various dense prediction tasks. We show the interplay of the encoded representations or cues between a stereo pair in ChiTransformer and how they can be effectively converted into dense depth prediction. The intuition for the elicitation and success of this method is discussed.

3.1. Architecture

Overview: The complete architecture of ChiTransformer is shown in Figure 2. ChiTransformer employs a pair of hybrid vision transformers as backbone with ResNet-50 [27] for patch embedding. The parameters of two ResNet-50 are shared to ensure consistency in representations. The image patch embeddings are first projected to 768 dimensions, then flattened and summed with positional embeddings before fed into attention blocks. For an image of size $H \times W$, with the patch size of $P \times P$, the result is a set $T = \{t_0, t_1, \dots, t_{N_p}\}$, where $N_p = \frac{H \cdot W}{P^2}$ and t_0 is the class token. Here, patches are in the role of “words” for transformer, but we will refer to patches as “words” or “tokens” interchangeably hereafter. The attention block for the reference view closely follows the design in [15] with class token included, whereas master tokens are self-attended in the first multiple SA layers, followed by cross-attention (CA) and self-attention (SA) layers in an interleaved fashion. The output tokens of the master ViT (and reference ViT in training) are then reassembled into an image-like arrangement.

Feature representations I^s at different scales $s \in S$ are progressively aggregated and fused into the final depth estimation in the fusion block, which is modified from RefineNet [45]. Fusion block is shared for both views in training phase but dedicated to the master view in inference.

Attention layers: Self-attention layer is the crucial part for transformers and other attention-based methods to achieve superior performance over their non-attention competitors. The key advantage is that complex context information can be gathered in a global scope. With multiple layers of SA, encodings get progressively tempered with the context information as it goes deeper into the attention layers. This mechanism begets the globally coherent predictions. Therefore, instead of putting immediate connection to the CA layer, we place multiple ($l_{SA} = 4$) SA layers at the output of the ResNet-50. Cues with appropriate amount of context information result in more reliable pattern retrieval in the subsequent CA layers. This design improves both training convergence and prediction performance.

The *cross-attention* layer is our key contribution in ChiTransformer. It is the enabler of the stereopsis through the fusion of high-level depth cue expressions from two views. We argue that the effectiveness of the traditional 4-step strategy would be largely weakened as the sources of ill-posedness, such as occlusion, wider and closer range of depth, depth discontinuity and nonlinearity, become increasingly frequent or prominent. Current deep learning-based methods rely on the learned rich representations to construct cost volume which is then regularized to make estimation. The output quality, in this case, largely depends on both the quality of the representations and the conformity of the scene to the matching regularizing assumptions [54]. While good representations can be learned with many approaches, there are few ways to fix up an impaired cost volume when scenes are far away from being appropriate for stereo matching. Therefore, instead of clinging to the matching strategy, we propose a novel pattern retrieval mechanism inspired by associative memory to retrieve the correspondent pattern from the other view. We assume that a set of patterns can be learned to separate well such that each pattern can be retrieved at least in meta-stable state, i.e., fixed average of similar patterns [50]. Modeled by modern Hopfield network [13, 38], the retrieval rule of the associative memory elegantly coincides with the attention mechanism of the transformer. Naturally, we leverage cross-attention layer to retrieve patterns (tokens) from the reference to the master view. To facilitate reliable effective retrieval, we devise a new attention mechanism – polarized attention, which enables feature-sensitive retrieval while preserving the context information contained in the pattern without breaching the convergence rule. From [63], we observe that direct attention over representations at the

output of CNN reduces to cosine similarity-based matching. Without position-dependent context information over extensive scope, patterns are liable to ill-posedness and low separability.

Given a token pair $(m_{t_i}, m_{t'_i}), \forall i \in \{1, \dots, N_p\}$, from the preceding CA layer and the class token pair (m_{t_0}, r_{t_0}) , where m indicates the master view, r the reference view, and t' denotes the retrieved tokens, depth cues are then rectified through the following blending process:

$$f_{proj}(m_{t_i}, m_{t'_i}) = \text{MLP}([m_{t_i}^\top, m_{t'_i}^\top]), \quad (1)$$

$$m_{t_i} = m_{t_i} + \text{Heat}(p_{a_i}) \cdot \text{LN}(f_{proj}(m_{t_i}, m_{t'_i})). \quad (2)$$

We set $m_{t_0} = r_{t_0}$ to unify the expression. GELU [28] is used for MLP nonlinearity, LN, layer normalization [4], p_{a_i} , the vector of attention scores of m_{t_i} , and Heat , the confidence score calculated with stabilized attention entropy as

$$\text{Heat}(p_{a_i}) = 1 - g(\text{H}(p_{a_i}), \tau, c), \quad (3)$$

where $\text{H}(p_{a_i}) = -\sum_{k=1}^{N_p} p_{a_i,k} \log(p_{a_i,k} + \epsilon)$, and $g(\cdot)$ is a clamping function (e.g., sigmoid or smoothstep) with temperature τ and offset c . Heat is set to 1 for class token. By doing so, tokens retrieved back in fixed state, i.e., with a very low entropy, would be securely rectified whereas those with high entropy are inhibited from being updated as they are very likely to reside in occluded areas. Thus, the depth in occluded or “uncertain” areas are left to the power of SA layers to speculate its value with context information and rectified cues from neighboring non-occluded areas.

Fusion block: Our convolutional decoder follows the refinement block in [45, 51]. The output of attention layers $\mathbf{t} \in \mathbb{R}^{(N_p+1) \times D}$ is reassembled into an image-like arrangement $\mathbb{R}^{H' \times W' \times D'}$ through a four-step operation:

$$\text{RSB} = (\text{rescale} \circ \text{reshape} \circ \text{MLP} \circ \text{cat}). \quad (4)$$

The class token is concatenated with all other tokens (by broadcasting) before being projected to dimension D' to get \mathbf{t}_{λ_0} . Then it is reshaped into a 2D shape per the original arrangement of the image embedding. Finally, \mathbf{t}_{λ_0} is re-sampled to size $\frac{H}{P} s_l \times \frac{W}{P} s_l \times D_l$ for different scales at level l . Re-sampling method is 2D transposed convolution for $s_l > 1$ (up-sampling), and strided 2D convolution for $s_l < 1$ (down-sampling). For our model, features from level $l_{attn} = \{11, 7, 3\}$ in attention blocks (12 in total for ChiTransformer-8) and level $l_{res} = \{1, 0\}$ (first 2 blocks) in ResNet-50 are reassembled. The reassembled feature maps from those levels are consecutively fused through customized feature fusion block from RefineNet [45]. At each level, feature map is up-sampled by a factor of 2 and finally

the depth estimation map reaches the original resolution of the input images.

The architecture of ChiTransformer is structurally similar and biologically analogous to the *optic-chiasma* structure in our visual system, where visual field covered by both eyes is fused to enable the processing of binocular depth perception by stereopsis [54], hence the name of our model.

3.2. Polarized Attention

We propose a new attention mechanism to highlight or suppress features, which is much like signal polarization but in feature domain. Ideally, for a set of tokens represented in tensor $\mathbf{t} = (t_1, \dots, t_N)$ that is well separated, highlighting or suppressing can be potentially achieved in token-wise granularity. However, ideal separability is hard to achieve in practice because the attention tensor \mathbf{A} for regular attention mechanism is calculated as

$$\mathbf{A} = \text{softmax} \left(\beta \mathbf{t}^\top \mathbf{W}_t^\top \mathbf{W}_\xi \xi \right), \quad (5)$$

which is prone to be noisy with joint activation over all channels and hard to learn directly for \mathbf{W} 's. While the prevalent MHA seeks for multi-level context instead of retrieval since tokens are mapped to different (sub-)spaces for each head that generates its own attention weights and output with the projected tokens. To achieve retrieval behavior, without loss of generality, we stick to the Hopfield network update rule to ensure retrieval behavior and condition the query pattern with a self-adjoint operator $\mathbf{G} \in \mathbb{R}^{d \times d}$,

$$\xi' = \mathbf{t} \text{softmax} \left(\beta \mathbf{t}^\top \mathbf{G} \xi \right), \quad (6)$$

where β is the scale factor set to be $1/\sqrt{d}$. We assume the constraint that the query and memory should stay in the same sub-space which is satisfied by a positive-definite \mathbf{G} decomposed as $\mathbf{G} = \mathbf{M}^\top \mathbf{M}$. It can further be spectrally decomposed to get:

$$\xi' = \mathbf{t} \text{softmax} \left(\beta \mathbf{t}^\top \mathbf{U}^\top \mathbf{\Lambda} \mathbf{U} \xi \right), \quad (7)$$

where \mathbf{U} is an orthogonal matrix and $\mathbf{\Lambda}$ is a positive diagonal matrix.

To achieve feature-sensitive retrieval while factoring in all the information in the embeddings, we desire $\text{diag}(\mathbf{\Lambda})$ not to be zero abounded, i.e., feature selection. To achieve that and also enable multi-modal retrieval, multiple $\mathbf{\Lambda}$ s are learned and we desire $\prod_{i=1}^s \mathbf{\Lambda}_i$ close to \mathbf{I} such that if one feature is highlighted in one mode it should be suppressed in other modes. As such, the new attention mechanism becomes

$$\xi' = \mathbf{W}_{\text{cat}}^s \left[\mathbf{t} \cdot \text{softmax} \left(\beta \mathbf{t}^\top \mathbf{U}^\top \mathbf{\Lambda}_i \mathbf{U} \xi \right) \right]. \quad (8)$$

For our model, ξ is the tokens from the master view, \mathbf{t} is the tokens from reference view, $s = 2$, \mathbf{W} projects the concatenated tokens back to its original dimension and ξ' is the retrieved tokens from the reference view. The lowest entropy

of the s heads is used in Heat calculation for rectification. The effectiveness of the retrieval mechanism is reflected in Figure 3, where up-scaled attention map of different tokens are overlapped with the reference view.

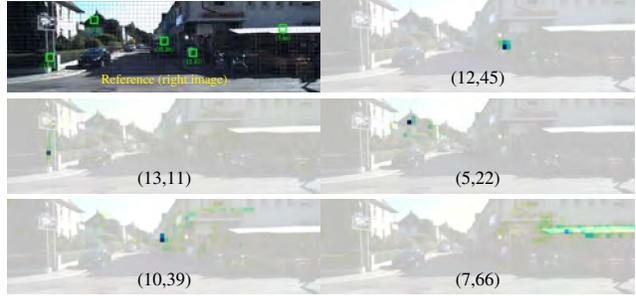


Figure 3. Attention map of $\mathbf{A}_{ij}^{s=1}$ of tokens retrieved in metastable state in the CA layer of the 2^{nd} DCR. Darker green corresponds to higher attention score. As expected, tokens of distinctive features have more centralized attention, while tokens in texture-less area result in more dispersed attention.

3.3. Learnable Epipolar Geometry

Token separability may be limited by the memory size and image content (e.g., existence of repetitive or uniform texture in the image). To further ensure secure retrieval without corrupting the encoded information, we constrain the attention mechanism with epipolar geometry through a *gated positional cross attention* (GPCA) following [12]. In GPCA, positional embedding is modeled as trainable quadratic polynomial of relative positional encoding $\mathbf{v}_{pos}^\top \mathbf{r}_{ij}$ [11]. For regular rectified stereo, candidate retrievals reside within collinear horizontal lines. Therefore, we set $\mathbf{v}_{pos} = -\alpha(0, 0, 0, 0, 0, 1, \dots, 0)$, $\mathbf{r} = (1, \delta_1, \delta_2, \delta_1 \delta_2, \delta_1^2, \delta_2^2, 0, \dots, 0)$. For non-rectilinear images, e.g., fisheye, \mathbf{v}_{pos} is a vector of trainable curve coefficients, which will be discussed when we present results on fisheye images.

In the equations above, \mathbf{r} is the position vector of (δ_1, δ_2) which are the relative coordinates with respect to the query. The locality strength $\alpha > 0$ determines how focused attention is along the horizontal line (i.e., when $\delta_2 = 0$). The positional attention scores are calculated as softmax normalized L_2 distance between the attended tokens and the query:

$$\mathbf{A}_{pos,ij} = \text{softmax} \left(\mathbf{v}_{pos}^\top \mathbf{r}_{ij} \right). \quad (9)$$

With the learnable gating parameter λ , the GPCA attention scores are calculated as:

$$\mathbf{A}_{ij}^s = \text{norm} \left[(1 - \sigma(\lambda)) \mathbf{A}_{cnt,ij}^s + \sigma(\lambda) \mathbf{A}_{pos,ij} \right], \quad (10)$$

where $\text{norm}[\mathbf{x}] = \frac{x_{ij}}{\sum_k x_{ik}}$, σ is the sigmoid function, and $\mathbf{A}_{cnt,ij}^s$ is the content attention score calculated by polarized attention of head s . To avoid GPCA from being stuck at $\lambda \gg 1$, we initialize $\lambda = 1$ for all layers.

3.4. Regularization

Matrix \mathbf{U} has to be orthogonal to guarantee that query and memory are attended in the same space. However, \mathbf{U} in each layer is trainable parameter; even though it can be initialized with orthogonal matrices, during training process the orthogonality may not hold. Therefore, we introduce an orthogonality regularization loss to \mathbf{U} as:

$$L_o(\mathbf{U}) = \frac{1}{d^2} \|\mathbf{U}^\top \mathbf{U} - \mathbf{I}\|_{\mathcal{F}}, \quad (11)$$

where d is the size of \mathbf{U} and $\|\cdot\|_{\mathcal{F}}$ is the Frobenius norm of matrix. Although \mathbf{U} can be orthogonalized through Cayley’s parameterization, it is computational expensive for large matrix as inversion is involved and we found it is more difficult to converge and unstable in our case.

To induce the diagonal matrix $\mathbf{\Lambda}$ to be trained into the desired form, we modified Hoyer regularizer [30] to mitigate the proportional scaling issue and at the same time to pull $\mathbf{\Lambda}$ away from being identity matrix. We introduce the following regularization:

$$L_{\Lambda}(\mathbf{\Lambda}) = \frac{\prod_{i=1}^s |\mathbf{\Lambda}_i|_e - \mathbf{I}_1}{\prod_{i=1}^s \|\mathbf{\Lambda}_i\|_{\mathcal{F}}}, \quad (12)$$

where $|\cdot|_e$ is the element-wise absolute function. Identity matrix is only one of the possible but unpreferred solution for $\prod_{i=1}^s |\mathbf{\Lambda}_i|_e - \mathbf{I} = 0$. Take 2D matrix for example, diagonal matrices $\Lambda_1 = (5, \frac{1}{5})$, $\Lambda_2 = (\frac{1}{5}, 5)$ is also a solution. However, when Λ_i s are approaching to this optimal, the denominator of Eq. 12 is much larger (~ 25) than that of Λ_i s close to the identity matrix (2.0). In addition, the final loss is a combination of reconstruction loss and regularization loss; with a proper hyperparameter μ_{λ} , the learned $\mathbf{\Lambda}$ s will be pulled away from \mathbf{I} .

3.5. Training

In this section, we provide details of the training method we used. We closely followed the self-supervised stereo training method provided in [25]. The model is trained to predict the target image from the other viewpoint in a stereo pair. Unlike classical binocular and multi-view stereo methods, the image synthesis process in our case is constrained by predicted depth instead of disparity as an intermediary variable. Specifically, given a target image I_t , a source image I'_t , and the predicted depths D_t , through the relative pose between two views $T_{t \rightarrow t'}$ calculated with the provided stereo base width (0.54m for KITTI) and calibration information, the correspondent coordinates between two images can be calculated. Following [33], the target image can be reconstructed from source image using bilinear sampling, which is sub-differentiable.

The depth prediction should minimize the photometric reprojection error constructed for both master and reference view as follows:

$$L_p = \omega \cdot \text{pe}(I_t, I_{t \rightarrow t}) + (1 - \omega) \cdot \text{pe}(I_{t'}, I_{t \rightarrow t'}), \quad (13)$$

where ω is the weight for master view, and $\text{pe}(\cdot)$ is the photometric reconstruction error [64]:

$$\text{pe}(X, Y) = \frac{\kappa}{2} (1 - \text{SSIM}(X, Y)) + (1 - \kappa) \|X - Y\|_1 \quad (14)$$

$\kappa = 0.85$ and $I_{t' \rightarrow t}$ is the reprojected image:

$$I_{t' \rightarrow t} = \text{bi-sample} \langle \text{proj}(D_t, T_{t \rightarrow t'}, K) \rangle, \quad (15)$$

where K is the pre-computed intrinsic matrix, proj is the resulting image coordinates projected from source view through

$$p'_t := K T_{t \rightarrow t'} D_t [p_t] K^{-1} p_t \quad (16)$$

and $\text{bi-sample} \langle \cdot \rangle$ is the bilinear sampler.

We also enforce edge-aware smoothness in the depths to improve depth-feature consistency defined as

$$L_s = |\partial_x d_t^*| e^{-|\partial_x I_t|} + |\partial_y d_t^*| e^{-|\partial_y I_t|}, \quad (17)$$

where $d_t^* = d/\bar{d}_t$ is the mean-normalized inverse depth in [61].

Unlike existing self-supervised stereo-matching methods that rely on predicted values to generate confidence map to detect occlusions, e.g. left-right consistency check, ChiTransformer detects occluded area on the fly in the form of heat map in the rectification stage. During training, heat map from the last GPCA layer is up-sampled to the output resolution and used as a mask m_h in loss computation. For stereo training, static camera and synchronous movement between objects and camera are not issues, hence we do not apply the binary auto-masking to block out the static area in the image.

During inference, only the master ViT output is up-scaled and refined to make the prediction. While in the training stage, both ViT towers in ChiTransformer are trained in tandem to predict depth and calculate losses L_p and L_s .

Final Training Loss By combining the reconstruction loss, per-pixel smoothness from two views and the regularizations for the matrices \mathbf{U} and $\mathbf{\Lambda}$, the final training loss is:

$$\mathcal{L} = \text{mean}(m_h \odot L_p) + \mu_s L_s + \mu_o L_o + \mu_{\lambda} L_{\lambda}, \quad (18)$$

where μ_* are the hyperparameters that balance the contributions from different loss terms.

Our models are implemented in PyTorch. With pre-trained ResNet-50 patch feature extractor and partial refinement layers from [51], the model is trained for 30 epochs using Adam [37] with a batch size of 12 and input resolution of 1216×352 . We use learning rate $1e-5$ for the ResNet-50 and $1e-4$ for the rest part of the network in the first 20 epoch, and then is decayed to $1e-5$ for the remaining epochs. We set $\omega = 0.6$, $\mu_s = 1e-4$, $\mu_o = 1e-7$ and $\mu_{\lambda} = 1e-3$ in our experiments.

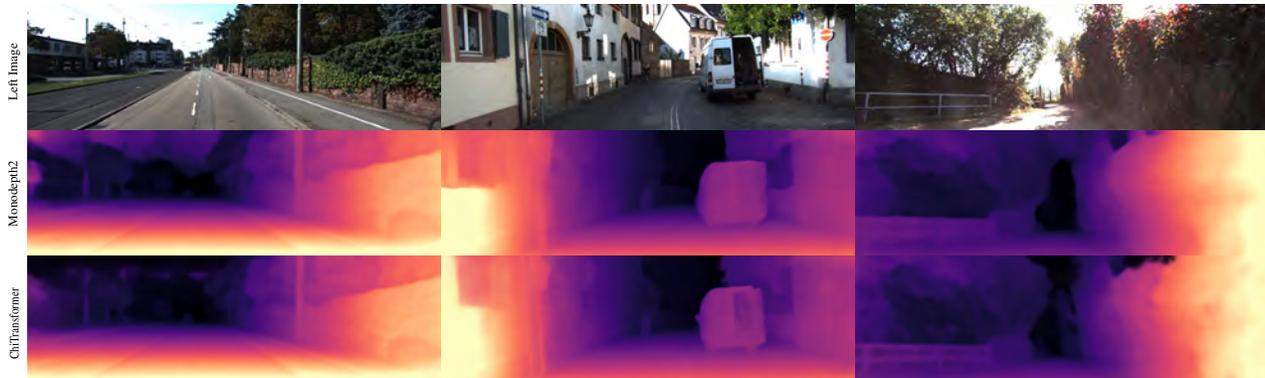


Figure 4. Sample results compared with self-stereo-supervised fully-convolutional network Monodepth2. ChiTransformer shows better global coherence (e.g., sky region, sides of image) and provides feature consistent details.

4. Experiments

The model is trained on KITTI 2015 [23]. We show that our model significantly improves accuracy compared to its top CNN-based counterparts. Side-by-side comparisons are given in this section with the state-of-the-art self-supervised stereo methods [42, 62, 63]. Ablation study is conducted to validate that several features in ChiTransformer contribute to the improved prediction. Finally, we extend our model to fisheye images and yield visually satisfactory result.

4.1. KITTI 2015 Eigen Split

We divide the KITTI dataset following the method of Eigen et al. [19]. Same intrinsic parameters are applied to all images by setting the camera principal point at the image center and the focal length as the average focal length of KITTI. For stereo training, the relative pose of a stereo pair is set to be pure horizontal translation of a fixed length (0.54m) according to the KITTI sensor setup. For a fair comparison, depth is truncated to 80m according to stan-

Table 1. QUANTITATIVE RESULTS

	Method	NOC					
		D1 (bg)	D1 (fg)	D1 (all)	D1 (bg)	D1 (fg)	D1 (all)
Supervised	DispNet [47]	4.11	3.72	4.05	4.32	4.41	4.43
	GC-Net [35]	2.02	5.58	2.61	2.21	6.16	2.87
	iResNet [44]	2.07	2.76	2.19	2.25	3.40	2.44
	PSMNet [10]	1.71	4.31	2.14	1.86	4.62	2.32
Self-supervised	Yu et al. [34]	-	-	8.35	-	-	19.14
	Zhou et al. [75]	-	-	8.61	-	-	9.91
	SegStereo [66]	-	-	7.70	-	-	8.79
	OASM [42]	5.44	17.30	7.39	6.89	19.42	8.98
	PASMnet'192 [63]	5.02	15.16	6.69	5.41	16.36	7.23
	Flow2Stereo [46]	4.77	14.03	6.29	5.01	14.62	6.61
	pSGM [41]	4.20	10.08	5.17	4.84	11.64	5.97
	MC-CNN-WS [58]	3.06	9.42	4.11	3.78	10.93	4.97
	SsSMnet [74]	2.46	6.13	3.06	2.70	6.92	3.40
	PVSstereo [62]	2.09	5.73	2.69	2.29	6.50	2.99
	ChiT-8 (ours)	2.24	4.33	2.56	2.50	5.49	3.03
	ChiT-12 (ours)	2.11	3.79	2.38	2.34	4.05	2.60

Comparison of our model to the state-of-the-art self-supervised binocular stereo methods. The lower the better for all metrics.

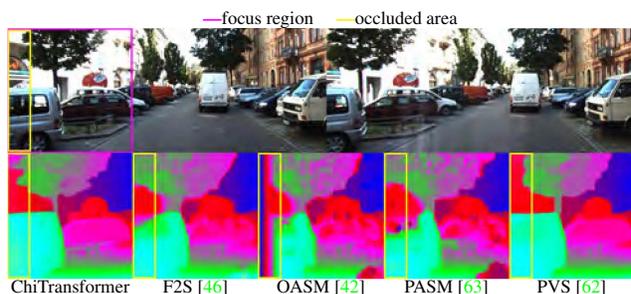


Figure 5. Comparison of predictions in left side occluded area.

dard practice [24].

4.2. Quantitative Results

We compare the results of the two different configurations of our model with state-of-the-art self-supervised stereo approaches. ChiT-8 has 4 SA layers followed by 4 rectification blocks, while ChiT-12 has 6 SA layers and 6 rectification blocks. The results in Table 1 show that ChiTransformer outperforms most of the existing methods, particularly in the prediction of the foreground regions. This trait is as expected that foreground regions are more likely to be abounded by distinctive features that benefit depth cue rectification. Qualitative results in occluded region comparing with existing self-supervised stereo methods are shown in Figure 5. Their complete predictions are given in supplementary material. With depth cues from both views, ChiTransformer provides more details consistent to the image features compared to existing self-supervised methods.

We also compare our method with top self-stereo-supervised MDE methods to show the reliability gain in terms of accuracy improvement. For a fair comparison, we choose the models that are trained on KITTI with stereo supervision. Methods trained over multiple datasets are not considered. Quantitative results are shown in Table 2. Side-by-side prediction comparison is shown in Figure 4.

Table 2. COMPARISON WITH SELF-STEREO-SUPERVISED MONOCULAR METHODS

Method	AbsRel	SqRel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Garg et al. [22]	0.152	1.226	5.849	0.246	0.784	0.921	0.967
3Net(R50) [49]	0.129	0.996	5.281	0.223	0.831	0.939	0.974
3Net(VGG) [49]	0.119	1.201	5.888	0.208	0.844	0.941	0.978
SuperDepth+pp (1024×382) [48]	0.112	0.875	4.968	0.207	0.852	0.947	0.977
Monodepth2 [25]	0.109	0.873	4.960	0.209	0.864	0.948	0.975
ChiT-12 (ours)	0.073	0.634	3.105	0.118	0.924	0.989	0.997

All models listed in the table above are trained with self-supervised methods using stereo pair. Same as the monocular methods, ChiTransformer relies on depth cues to estimate depth, only with extra information from a second image.

Table 3. ABLATION STUDY

	AbsRel	RMSE	RMSElog	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
ChiT+P	0.106	4.845	0.204	0.878	0.960	0.981
ChiT+G+LEG	0.101	4.783	0.203	0.895	0.966	0.983
ChiT+P+Linear	0.092	4.535	0.201	0.889	0.964	0.987
ChiT+P+LEG	0.085	3.924	0.181	0.906	0.979	0.991

Evaluations for different settings of ChiTransformer (ChiT) trained on KITTI 2015 with Eigen split. "P" denotes the polarized attention. "G" stands for the direct learning of matrix G . "LEG" represents the feature of learnable epipolar geometry. "Linear" is the single line attention zone. ChiT with only P enabled has the lowest score due to inferior token separability. With the addition of LEG, the model: ChiT+P+LEG, becomes the top performer and show the advantage of P over G compared with ChiT+G+LEG. ChiT+P+Linear has the 2nd best performance. The serration effect due to "Linear" is largely mitigated by the long range context information and SA layers.

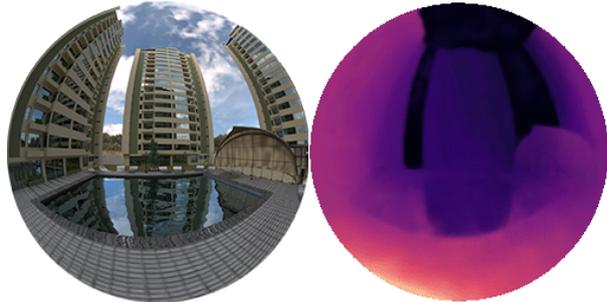
4.3. Ablation Study

To understand how each major feature contributes to the overall performance of ChiTransformer, ablation study is conducted by suppressing or activating specific components of the model. We observe that each component in our model is designed to push the performance a bit forward which aggregates into a sizable improvement. Here we provide some insights on the major features based on observation.

Self-attention layer largely improves the separability of each token with long range complex contextual information. Without SA layer, the retrieval process would take up a hopping behavior and result in erroneous predictions.

Polarized attention We learn the matrix G through its spectral decomposition to gain more control over its behavior. Direct learning of G tends to result in feature negligence as the major features contained within a token dominate or take all the reward. With the complementary feature highlighting-suppressing strategy as we desire $\prod \Lambda$ to be close to \mathbf{I} , the features from both parties can be attended. Meanwhile, since Λ is not porous with zeros, i.e., no Lasso regularization involved, all information contained in the token is more or less attended.

Learnable Epipolar Geometry Intuitively, for rectified stereo a pixel pair is guaranteed to reside in the same horizontal line and hence the attending space should be that line. However, the slotted attention region hurts the inter-line connection and cause serrated effects on vertical features even that feature is distinctive, e.g., an edge. Whilst the learnable epipolar geometry in GPCA solves the problem by allowing global but focused view over the lines and at the same time further improves the cross-line separability. Quantitative results are given in Table 3. Qualitative results are given in supplementary material.



Synthetic fisheye image

Planar depth estimation

Figure 6. Example result of ChiTransformer for fisheye depth estimation. With learnable epipolar curve $\mathbf{v}_{pos,ij} = (1, a, b, c, d, e)_{ij}$ (constant term is set to 1 to avoid proportional scaling) and circular masks, ChiTransformer can directly work on circular image without warping.

4.4. Example Results on Fisheye Images

Finally, due to its versatility, ChiTransformer can be applied to non-rectilinear (e.g., fisheye) images without warping. An example result is provided in Figure 6.

5. Conclusion

By investigating the limitations of the two prevalent methodologies of depth estimation, we present ChiTransformer, a novel and versatile stereo model that generates reliable depth estimation with rectified depth cues instead of stereo matching. With the three major contributions: (1) polarized attention mechanism, (2) learnable epipolar geometry, and (3) the depth cue rectification method, our model outperforms the existing self-supervised stereo methods and achieves state-of-the-art accuracy. In addition, due to its versatility, ChiTransformer can be applied to fisheye images without warping, yielding visually satisfactory results.

6. Acknowledgment

This research was sponsored in part by VMware Inc. for its university research fund and the Army Research Laboratory under Cooperative Agreement #W911NF-22-2-0025. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

- [1] Ramesh Aditya, Pavlov Mikhail, Goh Gabriel, and Gray Scott. Dalle: Creating images from text. In *OpenAI*, 2021. 2
- [2] Aria Ahmadi and Ioannis Patras. Unsupervised convolutional neural networks for motion estimation. In *2016 IEEE international conference on image processing (ICIP)*, pages 1629–1633. IEEE, 2016. 3
- [3] Alex M Andrew. Multiple view geometry in computer vision. *Kybernetes*, 2001. 1
- [4] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. *arXiv preprint arXiv:1606.08415*, 2016. 4
- [5] Stephen T Barnard and Martin A Fischler. Computational stereo. *ACM Computing Surveys (CSUR)*, 14(4):553–572, 1982. 1
- [6] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4009–4018, 2021. 2, 3
- [7] Amlaan Bhoi. Monocular depth estimation: A survey. *arXiv preprint arXiv:1901.09402*, 2019. 2
- [8] Arunkumar Byravan and Dieter Fox. Se3-nets: Learning rigid body motion using deep neural networks. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 173–180. IEEE, 2017. 3
- [9] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020. 2
- [10] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5410–5418, 2018. 2, 7
- [11] Jean-Baptiste Cordonnier, Andreas Loukas, and Martin Jaggi. On the relationship between self-attention and convolutional layers. *arXiv preprint arXiv:1911.03584*, 2019. 5
- [12] Stéphane d’Ascoli, Hugo Touvron, Matthew Leavitt, Ari Morcos, Giulio Biroli, and Levent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. *arXiv preprint arXiv:2103.10697*, 2021. 2, 5
- [13] Mete Demircigil, Judith Heusel, Matthias Löwe, Sven Upgang, and Franck Vermet. On a model of associative memory with huge storage capacity. *Journal of Statistical Physics*, 168(2):288–299, 2017. 4
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 3
- [16] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Häusser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2758–2766, 2015. 1
- [17] Shivam Duggal, Shenlong Wang, Wei-Chiu Ma, Rui Hu, and Raquel Urtasun. Deeppruner: Learning efficient stereo matching via differentiable patchmatch. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4384–4393, 2019. 2
- [18] Andrea Eichenseer and André Kaup. A data set providing synthetic and real-world fisheye video sequences. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1541–1545. IEEE, 2016. 2
- [19] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *arXiv preprint arXiv:1406.2283*, 2014. 2, 7
- [20] Philipp Fischer, Alexey Dosovitskiy, Eddy Ilg, Philip Häusser, Caner Hazirbas, Vladimir Golkov, Patrick Van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. *arXiv preprint arXiv:1504.06852*, 2015. 2
- [21] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2002–2011, 2018. 2
- [22] Ravi Garg, Vijay Kumar Bg, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *European conference on computer vision*, pages 740–756. Springer, 2016. 8
- [23] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. 3, 7
- [24] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 270–279, 2017. 2, 3, 7
- [25] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *ICCV*, pages 3828–3838, 2019. 1, 2, 6, 8
- [26] Xiaoyang Guo, Hongsheng Li, Shuai Yi, Jimmy Ren, and Xiaogang Wang. Learning monocular depth by distilling cross-domain stereo networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 484–500, 2018. 2
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3
- [28] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1607.06450*, 2016. 4

- [29] Yuxin Hou, Juho Kannala, and Arno Solin. Multi-view stereo by temporal nonparametric fusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2651–2660, 2019. [2](#)
- [30] Patrik O Hoyer. Non-negative matrix factorization with sparseness constraints. *Journal of machine learning research*, 5(9), 2004. [6](#)
- [31] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. Deepmvs: Learning multi-view stereopsis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2821–2830, 2018. [2](#)
- [32] Eddy Ilg, Tonmoy Saikia, Margret Keuper, and Thomas Brox. Occlusions, motion and depth boundaries with a generic network for disparity, optical flow or scene flow estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 614–630, 2018. [1](#)
- [33] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. *Advances in neural information processing systems*, 28:2017–2025, 2015. [6](#)
- [34] J Yu Jason, Adam W Harley, and Konstantinos G Derpanis. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In *European Conference on Computer Vision*, pages 3–10. Springer, 2016. [3](#), [7](#)
- [35] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 66–75, 2017. [7](#)
- [36] Salman Khan, Hossein Rahmani, Syed Afaq Ali Shah, and Mohammed Bennamoun. A guide to convolutional neural networks for computer vision. *Synthesis Lectures on Computer Vision*, 8(1):1–207, 2018. [2](#)
- [37] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [38] Dmitry Krotov and John J Hopfield. Dense associative memory for pattern recognition. *Advances in neural information processing systems*, 29:1172–1180, 2016. [4](#)
- [39] Hamid Laga, Laurent Valentin Jospin, Farid Boussaid, and Mohammed Bennamoun. A survey on deep learning techniques for stereo-based depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. [1](#)
- [40] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326*, 2019. [2](#)
- [41] Yeongmin Lee, Min-Gyu Park, Youngbae Hwang, Youngsoo Shin, and Chong-Min Kyung. Memory-efficient parametric semiglobal matching. *IEEE Signal Processing Letters*, 25(2):194–198, 2017. [7](#)
- [42] Ang Li and Zejian Yuan. Occlusion aware stereo matching via cooperative unsupervised learning. In *Asian Conference on Computer Vision*, pages 197–213. Springer, 2018. [3](#), [7](#)
- [43] Ruibo Li, Ke Xian, Chunhua Shen, Zhiguo Cao, Hao Lu, and Lingxiao Hang. Deep attention-based classification network for robust depth prediction. In *Asian Conference on Computer Vision*, pages 663–678. Springer, 2018. [2](#)
- [44] Zhengfa Liang, Yiliu Feng, Yulan Guo, Hengzhu Liu, Wei Chen, Linbo Qiao, Li Zhou, and Jianfeng Zhang. Learning for disparity estimation through feature constancy. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2811–2820, 2018. [2](#), [7](#)
- [45] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1925–1934, 2017. [4](#)
- [46] Pengpeng Liu, Irwin King, Michael R Lyu, and Jia Xu. Flow2stereo: Effective self-supervised learning of optical flow and stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6648–6657, 2020. [7](#)
- [47] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4040–4048, 2016. [7](#)
- [48] Sudeep Pillai, Rareş Ambrus, and Adrien Gaidon. Superdepth: Self-supervised, super-resolved monocular depth estimation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 9250–9256. IEEE, 2019. [8](#)
- [49] Matteo Poggi, Fabio Tosi, and Stefano Mattoccia. Learning monocular depth estimation with unsupervised trinocular assumptions. In *2018 International conference on 3d vision (3DV)*, pages 324–333. IEEE, 2018. [8](#)
- [50] Hubert Ramsauer, Bernhard Schäfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Thomas Adler, Lukas Gruber, Markus Holzleitner, Milena Pavlović, Geir Kjetil Sandve, et al. Hopfield networks is all you need. *arXiv preprint arXiv:2008.02217*, 2020. [2](#), [4](#)
- [51] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *ICCV*, 2021. [1](#), [2](#), [3](#), [4](#), [6](#)
- [52] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *arXiv preprint arXiv:1907.01341*, 2019. [2](#)
- [53] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Advances in neural information processing systems*, pages 14866–14876, 2019. [2](#)
- [54] Daniel Scharstein and Richard Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International journal of computer vision*, 47(1):7–42, 2002. [1](#), [2](#), [4](#), [5](#)
- [55] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *European conference on computer vision*, pages 746–760. Springer, 2012. [3](#)
- [56] Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. A benchmark for the evaluation of rgb-d slam systems. In *2012 IEEE/RSJ international*

- conference on intelligent robots and systems*, pages 573–580. IEEE, 2012. [2](#)
- [57] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021. [2](#)
- [58] Stepan Tulyakov, Anton Ivanov, and Francois Fleuret. Weakly supervised learning of deep metrics for stereo reconstruction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1339–1348, 2017. [7](#)
- [59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [2](#)
- [60] Sudheendra Vijayanarasimhan, Susanna Ricco, Cordelia Schmid, Rahul Sukthankar, and Katerina Fragkiadaki. Sfmnet: Learning of structure and motion from video. *arXiv preprint arXiv:1704.07804*, 2017. [3](#)
- [61] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2022–2030, 2018. [6](#)
- [62] Hengli Wang, Rui Fan, Peide Cai, and Ming Liu. Pvsstereo: Pyramid voting module for end-to-end self-supervised stereo matching. *IEEE Robotics and Automation Letters*, 6(3):4353–4360, 2021. [7](#)
- [63] Longguang Wang, Yulan Guo, Yingqian Wang, Zhengfa Liang, Zaiping Lin, Jungang Yang, and Wei An. Parallax attention for unsupervised stereo correspondence learning. *IEEE transactions on pattern analysis and machine intelligence*, 2020. [3](#), [4](#), [7](#)
- [64] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. [6](#)
- [65] Fuzhi Yang, Huan Yang, Jianlong Fu, Hongtao Lu, and Bainig Guo. Learning texture transformer network for image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5791–5800, 2020. [2](#)
- [66] Guorun Yang, Hengshuang Zhao, Jianping Shi, Zhidong Deng, and Jiaya Jia. Segstereo: Exploiting semantic information for disparity estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 636–651, 2018. [2](#), [3](#), [7](#)
- [67] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 767–783, 2018. [2](#)
- [68] Linwei Ye, Mrigank Rochan, Zhi Liu, and Yang Wang. Cross-modal self-attention network for referring image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10502–10511, 2019. [2](#)
- [69] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5684–5693, 2019. [2](#)
- [70] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1983–1992, 2018. [3](#)
- [71] Jure Zbontar, Yann LeCun, et al. Stereo matching by training a convolutional neural network to compare image patches. *J. Mach. Learn. Res.*, 17(1):2287–2318, 2016. [1](#)
- [72] Yanhong Zeng, Jianlong Fu, and Hongyang Chao. Learning joint spatial-temporal transformations for video inpainting. In *European Conference on Computer Vision*, pages 528–543. Springer, 2020. [2](#)
- [73] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. Ga-net: Guided aggregation net for end-to-end stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 185–194, 2019. [1](#)
- [74] Yiran Zhong, Yuchao Dai, and Hongdong Li. Self-supervised learning for stereo matching with self-improving ability. *arXiv preprint arXiv:1709.00930*, 2017. [2](#), [7](#)
- [75] Chao Zhou, Hong Zhang, Xiaoyong Shen, and Jiaya Jia. Unsupervised learning of stereo matching. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1567–1575, 2017. [3](#), [7](#)
- [76] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. [2](#)