

AFT-VO: Asynchronous Fusion Transformers for Multi-View Visual Odometry Estimation

Nimet Kaygusuz, Oscar Mendez, Richard Bowden¹

Abstract—Motion estimation approaches typically employ sensor fusion techniques, such as the Kalman Filter, to handle individual sensor failures. More recently, deep learning-based fusion approaches have been proposed, increasing the performance and requiring less model-specific implementations. However, current deep fusion approaches often assume that sensors are synchronised, which is not always practical, especially for low-cost hardware. To address this limitation, in this work, we propose AFT-VO, a novel transformer-based sensor fusion architecture to estimate VO from multiple sensors. Our framework combines predictions from asynchronous multi-view cameras and accounts for the time discrepancies of measurements coming from different sources.

Our approach first employs a Mixture Density Network (MDN) to estimate the probability distributions of the 6-DoF poses for every camera in the system. Then a novel transformer-based fusion module, AFT-VO, is introduced, which combines these asynchronous pose estimations, along with their confidences. More specifically, we introduce *Discretiser* and *Source Encoding* techniques which enable the fusion of multi-source asynchronous signals.

We evaluate our approach on the popular *nuScenes* and *KITTI* datasets. Our experiments demonstrate that multi-view fusion for VO estimation provides robust and accurate trajectories, outperforming the state of the art in both challenging weather and lighting conditions.

I. INTRODUCTION

Visual Odometry (VO) can be described as the process of estimating the relative pose and motion of a moving camera given a series of images. Traditional approaches have been studied for decades and are used to extract motion in applications ranging from robotics to autonomous driving. They generally follow the standard procedure of feature detection followed by robust matching and use projective geometry to estimate camera motion & therefore odometry. However, these approaches are not robust against poor lighting conditions and low textured environments, where the feature detection and matching tend to fail. In recent years, learning-based VO approaches have become popular and achieved competitive performance [29], [16], [31] due to their ability to extract high-level representations without relying on hand-crafted features.

Most of the current VO approaches employ a single monocular or stereo camera. However, VO is prone to failure when the quality of the captured image is degraded by occlusion, lighting change, etc. Various approaches have been proposed to address this issue by combining cameras with non-visual sensors such as IMUs [5], [4]. They demonstrate improved performance when camera motion estimation fails.

However, employing multiple cameras is a clear way to provide robustness to individual camera failures. This has led to employing multiple cameras with different views as an alternative way to provide robustness to individual camera failures [27], [23].

Multi-camera systems [17], [27], follow an optimisation-based approach, however, they assume that the sensors are time-synchronised. Unfortunately, synchronous multi-source signals are hard to acquire. This is because obtaining images from multiple cameras at the same time requires ‘gen-locked’ or synchronised hardware and sufficient bandwidth to capture and process those images with the associated cost implication. Thus new advancements are needed that can enable the use of low-cost asynchronous sensors.

Motivated by the limitations of the current multi-camera approaches, in this work, we aim to estimate robust VO using images obtained from multiple unsynchronised cameras. To achieve this, we introduce a novel transformer-based deep fusion framework. Our approach can fuse signals from any number of sources and, more importantly, it does not require those sources to be synchronised or to be at the same frequency.

As inputs, our approach takes variable-length video sequences from any number of cameras attached to a vehicle. We first process every video stream separately and estimate independent motion and uncertainties using a deep MDN-based VO module [14]. Due to the asynchronous nature of these signals, naive fusion approaches, such as concatenating representations corresponding to the same time step, are not applicable. To address this, we introduce a novel *Discretiser* module which enables our system to positionally encode continuous timestamps. Thus, we also introduce *Source Encoding* which enables the fusion module to differentiate signals coming from different cameras as predictions from different cameras do not have any indication of their origin.

The contributions of this paper can be summarised as:

- We propose AFT-VO, a transformer-based deep fusion framework which can fuse asynchronous signals from multiple sources. We apply our fusion approach to the multi-view VO estimation problem.
- We introduce two new types of encoding approach into transformers: (1) *Discretiser*: to represent the temporal relationship between asynchronous time series (2) *Source Encoding*: to inform the model about the source of the information.
- We evaluate our approach on two popular public datasets, namely, KITTI and the nuScenes datasets, and achieve state-of-the-art performance.

¹All authors are with the University of Surrey, UK {n.kaygusuz, o.mendez, r.bowden}@surrey.ac.uk

II. RELATED WORK

Traditional VO approaches can be split into two categories which are direct and feature-based methods. Direct methods estimate the motion by minimising the photometric error [8], [21]. Feature-based methods estimate the motion by extracting and matching a set of feature points on consecutive frames [20], [22], [32]. While direct methods perform well in low texture settings, feature-based methods are more successful at higher velocities.

In recent years, deep learning-based VO methods achieved promising results. These methods learn to extract high-level features from the data and therefore do not need hand-crafted features. Current learning-based VO approaches can be classified as supervised, which we employ in this work, or unsupervised. One of the earliest supervised approaches was proposed by Mohanty *et al.* [18] who estimate VO using a Convolutional Neural Network (CNN). Wang *et al.* [29] combine a CNN with a Recurrent Neural Network (RNN) to enhance temporal dependencies. Kaygusuz *et al.* [14] utilise an MDN to estimate VO and its uncertainty. Saputra *et al.* [25] introduce geometric loss constraints to increase generalisation capability. Xue *et al.* [31] propose a memory module to store global information which is then used to refine the estimated poses. To reduce the dependency of labelled data, unsupervised methods have also been studied. Zhou *et al.* [35] proposed learning depth and motion jointly. Li *et al.* [16] and Zhan *et al.* [33] use stereo images to recover the scaled VO. While unsupervised methods are more easily scaled, to date they have not performed as well as supervised methods. Additionally, both supervised and unsupervised methods rely on single-camera models, which are liable to fail in difficult environmental conditions (*e.g.* glare, low texture, etc).

To overcome motion estimation failure, researchers have studied sensor fusion techniques. Geneva *et al.* [10] propose a multi-sensor fusion system based on a factor graph-based optimisation framework. Sola *et al.* [26] employed a Kalman Filter to fuse measurements from a multi-camera system. Liu *et al.* [17] used multiple cameras to build a classical VO approach. Zhang *et al.* [34] studied a multi-view visual-inertial odometry system.

In recent years, learning-based fusion for motion estimation has been studied. Kaygusuz *et al.* [15] introduced a learning-based multi-view fusion approach in which they assume that the cameras are synchronised. Clark *et al.* [5] proposed a learning-based approach to visual-inertial odometry by simply concatenating features from a camera and an inertial sensor. Chen *et al.* [4] also studied visual-inertial odometry and introduced a selective fusion approach. However, they relied on only inertial measurements when the camera images are degraded by occlusion, sunlight etc. Contrary to previous techniques, in this work, we propose the first learning-based asynchronous multi-view VO system.

Historically, deep learning approaches have used recurrent architectures to model time series. More recently, transformer models [28] have been introduced for time series data, in

the domain of machine translation. They have drastically increased the performance of sequence-to-sequence tasks such as, speech recognition [13], language understanding [6] and sign language translation [3].

Inspired by the recent success of transformers, in this work we propose a novel architecture to fuse measurements for the estimation of VO. We formulate our fusion approach as a sequence-to-sequence problem where the inputs are asynchronous and multi-source and do not necessarily have the same cardinality as the output.

III. MULTI-VIEW VISUAL ODOMETRY

We study multi-view VO as a fusion problem with observations captured by different sensors at different timestamps. Our task is to estimate vehicle trajectory, $Y_{1:U}$, given variable-length asynchronous video sequences, $\mathcal{V} = \{\mathcal{V}^0, \dots, \mathcal{V}^K\}$ from $(K+1)$ cameras all rigidly mounted to the same vehicle. A video stream coming from the k^{th} camera, C^k , can be written as:

$$\mathcal{V}^k = \{(I_0^k, t_0^k), \dots, (I_n^k, t_n^k), \dots, (I_{N^k}^k, t_{N^k}^k)\} \quad (1)$$

where I_n^k represent the n^{th} video frame captured at the timestamp t_n^k , and N^k represents the total number of captured frames by the k^{th} camera. Note that due to asynchronicity, frames coming from different cameras, I_n^i and I_n^j , that share the same order index, n , do not necessarily share the same timestamp, *i.e.* $t_n^i \neq t_n^j$.

Our network can be broken into two sections: a Mixture Density Network (MDN) and our novel Asynchronous Fusion Transformer (AFT). For each camera, we first estimate a mixture of Gaussian distributions of the 6-DoF relative poses using an MDN. We then fuse all pose distributions estimated through different arms of the network in the AFT module.

A. Mixture Density Network (MDN)

Learning-based approaches to VO generally regress a 6-DoF pose directly. In this work, we employ an MDN similar to [14] and estimate the probability distribution of the 6-DoF pose, conditioned on the input images. Estimating the pose as a mixture of distributions provides our fusion module with an indicator of the estimation uncertainty.

We first extract latent representations from consecutive image pairs in each camera stream using a CNN. We employ FlowNet [7] as our CNN backbone which has proven to be successful for learning features that are geometrically meaningful. We formalise this operation for two consecutive images, (I_{t-1}^k, I_t^k) , from camera C^k at time t as:

$$w_t^k = \text{CNN}([I_{t-1}^k, I_t^k]). \quad (2)$$

where $[]$ operation represents the channel-wise concatenation of consecutive images.

Motion estimation is highly dependant on temporal information. Thus, we employ an RNN, which takes in the latent representation, w_t^k , at time t to produce temporally enriched representations, r_t^k , as:

$$r_t^k = \text{RNN}(w_t^k, h_{t-1}^k). \quad (3)$$

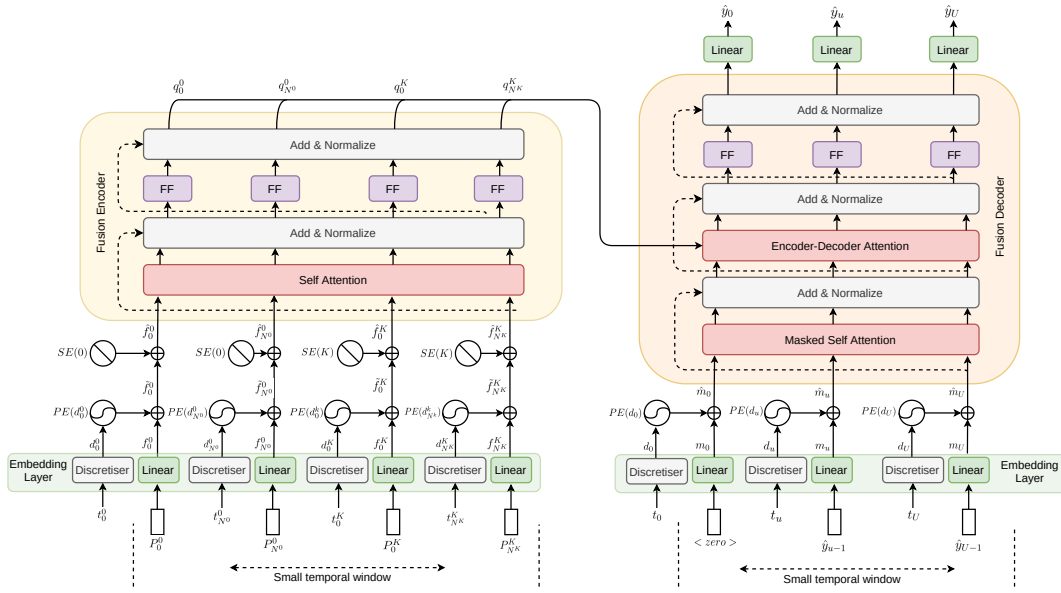


Fig. 1: Overview of the proposed asynchronous fusion module. To fuse asynchronous information and differentiate signals that are coming from different cameras, we utilise two novel encoding approaches, *Discretiser* and *Source Encoding (SE)*.

where h_{t-1}^k represents the RNN's previous hidden state.

We then pass the temporally enhanced representations through the MDN module. We build our mixture model as a multivariate Gaussian distribution and formalise a mixture model, P_t^k as:

$$P_t^k = \{(\alpha_t^1 \mathcal{N}_t(\mu^1, \sigma^1))^k, \dots, (\alpha_t^X \mathcal{N}_t(\mu^X, \sigma^X))^k\} \quad (4)$$

where X denotes the number of mixture components, while μ , σ and α represent the mean, standard deviation and mixture coefficient, respectively. A linear combination of mixture components produces the probability density of the target pose, $y_{(t-1,t)}^k$, and can be written as:

$$p(y_{(t-1,t)}^k | r_t^k) = \sum_{i=1}^X (\alpha_i(r_t^k))^k \phi_i(y_{(t-1,t)}^k | r_t^k) \quad (5)$$

where ϕ_i is the conditional density function for the i^{th} component. The mixture coefficient, $(\alpha_i(r_t^k))^k$, represents the probability of the target pose, $y_{(t-1,t)}^k$, being generated by the i^{th} component. We train our MDN module by minimising the negative log likelihood of the ground truth poses coming from our estimated mixture model.

B. Asynchronous Fusion Transformer (AFT)

We consider the task of multi-view VO estimation as a sequence-to-sequence learning problem and propose a novel transformer-based approach to fuse complementary information coming from multiple asynchronous cameras.

The input to the asynchronous fusion module is the pose and uncertainty estimates from the MDN module with their corresponding timestamps, $\{(P_0^k, t_0^k), \dots, (P_n^k, t_n^k), \dots, (P_{N^k}^k, t_{N^k}^k)\}$. One thing to note is that due to asynchronicity, estimates coming from different cameras within a selected time interval, *e.g.* C^i and C^j , will not necessarily share the same cardinality,

i.e. $\diamond(N^i \neq N^j)$, or timestamps of the same time indices, n , will not necessarily be equal, *i.e.* $\diamond(t_n^i \neq t_n^j)$. We first project the predictions from individual cameras into a latent representation using a linear layer as:

$$f_n^k = \text{Linear}(P_n^k); \quad (6)$$

where P_n^k is the pose estimate and its uncertainty (See Section III-A) from the k^{th} camera at the n^{th} time step.

Transformers do not inherently contain any temporal modelling, and function by utilising attention mechanisms. To this end, positional encoding techniques are used to inject the order of the inputs by creating a unique vector for each discrete step. Although positional encoding has been applied to sequence-to-sequence learning problems, such as text-to-text [24], [28] and video-to-text tasks [1], they assume the consecutive items to be equidistant in the time domain *e.g.* the distance between every consecutive input is equal. This is a valid assumption for tasks where input comes from a single source with a fixed frequency. However, this is not applicable to our problem as the information comes from multiple asynchronous sources where the inputs are not equally distant from each other. To represent continuous time information and to differentiate signals coming from different sources, we propose two novel encoding approaches, namely *Discretiser* and *Source Encoding*.

Discretiser: In this module, our goal is to represent the continuous timing information with a set of discrete equidistant time step bins. Although the timings of captured images are generally reported in integer format, *i.e.* microseconds, the cardinality of this space is much larger than what transformers are designed to work with, *i.e.* number of words in a sentence. Also, the information change in this time domain is much slower than the text domain. For example, information gained from 1 microsecond of a 5 seconds video

is much smaller than what is gained from a word in a sentence with a cardinality of 10. Furthermore, the frequency of sensors can be varied throughout a capture, resulting in non-equidistant frames.

To address this issue we propose to discretise the continuous time domain into bins. Given a sequence of timestamps from an arbitrary time window, we first normalise them by finding the earliest stamp and subtracting it from all others. We then divide the time axis into smaller chunks, Z , and group the measurements into bins that have been captured close in time. Chunk size, Z , affects the discretisation sensitivity. In our experiments we choose 20ms as the chunk size, however reducing it does not affect the performance significantly. We formalise this operation as:

$$d_n^k = \text{round}\left(\frac{t_n^k - \min(\{t_0^0, \dots, t_{N^0}^0\}, \dots, \{t_0^K, \dots, t_{N^K}^K\})}{Z}\right) \quad (7)$$

where d_n^k represents the discretised form, *i.e.* the bin index, of the timestamp t_n^k , and Z is the quantisation step size. Note that the d_n^k is rounded down to the nearest integer. This allows us to convert the time axis into fewer discrete units, where each integer represents a bin of timestamps.

Next, we use these discretised representations to obtain relative position embeddings for inputs in a sequence. We employ the positional encoding approach proposed in [28], which uses a dense sequential set of integers, *i.e.* the position of the word in a sentence, to index the positional encoding matrix. However, our discretisation based use of positional encoding works with a sparse set of integers, *i.e.* the bin indexes of timestamps. This allows our model to learn the temporal relationship between asynchronous multi-source signals that lie on the continuous time axis. We achieve this by using d_n^k as an index to select from the positional encoding matrix. We formalise positional encoding of the inputs as:

$$\tilde{f}_n^k = f_n^k + \text{PositionalEncoding}(\text{Discretiser}(t_n^k)) \quad (8)$$

Source Encoding: Transformer architectures do not have any built-in way to identify the source of signals. However, our model needs to be able to differentiate between sources so that it can exploit the complementary information coming from them. To achieve this we propose *Source Encoding*. For each source in our system, we generate a unique one-hot vector. We then pass these one-hot representations through a linear layer that is learnt during training to obtain unique source encoding for each camera. We then combine this additional source information into the positionally encoded inputs as:

$$\hat{f}_n^k = \tilde{f}_n^k + \text{SourceEncoding}(k) \quad (9)$$

Fusion Encoder: After embedding the relative temporal position and the camera source information into the input representations, we fuse them using a transformer encoder model. The inputs to the encoder are first modelled by a self-attention layer which learns the correlations between different timestamps. The outputs of the self-attention layer

are then passed through a position-wise feed forward layer. Residual connections and layer normalisation are added after every operation. We formulate the fusion encoding process as:

$$q_n^k = \text{FusionEncoder}\left(\hat{f}_n^k | \{\hat{f}_{1:N^0}^0 \dots \hat{f}_{1:N^K}^K\}\right) \quad (10)$$

where q_n^k denotes the fused representation corresponding to the n^{th} frame coming from the k^{th} camera. Note that to produce q_n^k , the encoder attends to the representations of all predictions from all available sources in an arbitrarily chosen time window.

Fusion Decoder: The encoder produces a set of fused representations, $Q = \{q_n^k \forall n, k\}$, within a time window. We use these latent representations, Q , to decode the vehicle's odometry in an autoregressive manner.

During training, we concatenate a zero vector, $\langle 0 \rangle$, to the beginning of the target pose sequence, $\{\langle 0 \rangle, t_0, \dots, (\hat{y}_{u-1}, t_u), \dots, (\hat{y}_{U-1}, t_U)\}$, which is then used as the input of the decoder. This operation is analogous to the addition of the special beginning of sentence token in symbolic sequence-to-sequence tasks, which allow teacher forcing to be applied during training. We inject positional encoding using the *Discretiser* similar to the fusion encoder module. These embedded representations, \hat{m}_u , are then passed to a masked self-attention layer. Unlike the self-attention utilised in the encoder module, masking prevents the previous time steps from gaining information from the future while estimating a pose. This limitation is key, as it allows the system to function at inference time when the information is not available to the model. We then combine the output of the fusion encoder and decoder self-attention. These representations are passed to an encoder-decoder attention module. This module learns the relationship between source and target pose sequences. As with the fusion encoder, all layers are followed by residual connections and layer normalisation. We formulate this fusion decoder process as:

$$\hat{y}_{(t-1,t)} = \text{FusionDecoder}\left(\hat{m}_u | \hat{m}_{1:u-1}, \{q_{1:N^0}^0 \dots q_{1:N^K}^K\}\right) \quad (11)$$

During inference, we query the decoder with timestamps for which we want a pose estimate. In our experiments, we used the ground truth timestamps that are available to decode the vehicle's odometry, but the model can handle arbitrary timestamps. We train our asynchronous fusion transformer model using a Mean Squared Error (MSE) loss function as:

$$\mathcal{E} = \frac{1}{U} \sum_{t=1}^U \|y_{(t-1,t)}^\tau - \hat{y}_{(t-1,t)}^\tau\|_2^2 + \Omega \|y_{(t-1,t)}^\psi - \hat{y}_{(t-1,t)}^\psi\|_2^2 \quad (12)$$

where $y_{(t-1,t)}^\tau$ and $y_{(t-1,t)}^\psi$ represents the translation and rotation of the 6-DoF ground truth poses, respectively. Ω is a parameter to increase the importance of rotation errors (100 in our experiments) and U is the number of time steps within a given arbitrary time interval.

IV. EXPERIMENTS

In this section, we evaluate the proposed AFT-VO approach and compare it against the single and multi-sensor state-of-the-art VO techniques.

A. Dataset & Implementation Details

We conduct our experiments using two popular autonomous driving datasets, namely KITTI [9] and nuScenes [2]. KITTI has become the benchmark dataset for motion estimation. However, it is not an ideal dataset for asynchronous multi-camera fusion systems, considering it only has a stereo camera pair. For completeness, we evaluate our model on KITTI to demonstrate that our system can also work with synchronous cameras. In addition, the nuScenes dataset has 6 asynchronous cameras giving a 360 degree view, which makes it an ideal dataset to evaluate our approach.

We implement our model using the PyTorch deep learning framework. Pre-trained FlowNet weights are used for our CNN initialisation [7]. We build our transformers with 4 layers for both encoder and decoder, each layer containing 512 hidden units and 4 heads. The network parameters are initialised with the Xavier approach. We train our model using the Adam optimiser ($\beta_1 = 0.9, \beta_2 = 0.999$) with a batch size of 32 and a learning rate of 0.0005. To assess the performance of our approach we use the standard Relative Pose Error (RPE) metric from the *evo* evaluation toolkit [11].

B. Ablation Studies

In our first set of experiments, we ablate input sources and different components of the AFT-VO. We conduct these experiments on the nuScenes dataset. We perform two sets of ablation studies to measure (1) the effect of fusing different sets of cameras, (2) the benefits of the proposed modules.

Camera Ablation: In our first ablation study, we compare the effect of fusing different sets of cameras with complementary views.

Table I shows fusion results with different sets of cameras. As can be seen, the model that utilises a single view, namely front camera (F), has the largest error. Fusing other view combinations such as front and back cameras (F+B), and three front cameras (F+FL+FR), improves the performance. One interesting finding is that the two-camera system (F+B), outperforms the three-camera system, (F+FR+FL). We believe this is due to the significant increase in view coverage between these two setups. For example, there is natural overlap between cameras F, FR and FL resulting in considerable redundancy whereas (F+B) provides more complementary information.

Ultimately, we see that fusing all the cameras yields the best results, further suggesting that the proposed approach benefits from the complementary information coming from different views.

Module Ablation: In our second set of ablation studies, we evaluate the effects of individual modules on the performance, namely *Discretiser* and *Source Encoding*.

TABLE I: Average RPE (m) results from the camera ablation experiments. (F: Front, FR: Front Right, FL: Front Left, B: Back, BR: Back Right, BL: Back Left)

Fused Camera	RMSE	Max	Mean \pm std
F	0.05482	0.17238	0.04393 \pm 0.03190
F+FR+FL	0.05226	0.17684	0.04157 \pm 0.03102
F+BR+BL	0.04870	0.16094	0.03889 \pm 0.02899
F+B	0.04853	0.16183	0.03864 \pm 0.02867
F+FR+FL+B	0.04608	0.15116	0.03696 \pm 0.02682
All	0.04529	0.14983	0.03637 \pm 0.02644

To evaluate the importance of the *Discretiser*, we devise two experiments, namely (-D-Equi) and (-D-None). In (-D-Equi), we use standard positional encoding without the *Discretiser*, which means the model assumes that the consecutive frames are equidistant, *i.e.* all the cameras share a fixed frequency. In the second setup, (-D-None), we do not provide any timing information to our AFT-VO, removing both the *Discretiser* and the positional encoding. To ablate *Source Encoding*, (-SE), the unique one-hot vector that represents the camera sources is removed, so that the model does not know the source of its inputs.

We first share the overall performance of the proposed approach where the camera views and modules are used in Table II (See AFT-VO (ours)). As can be seen, the proposed approach benefits from both the *Discretiser* and the *Source Encoding* modules. Removing the *Discretiser*, (-D-Equi), degrades the model performance drastically (resulting in nearly five times the error rate). An interesting finding is that removing the temporal information completely (-D-None) outperforms using it sequentially (-D-Equi), in the case of asynchronous data. We argue that when the data is asynchronous, using the positional encoding matrix sequentially sends the wrong time information to the fusion model. This fundamentally corrupts the temporal information, thus ultimately lowering the performance of fusion, even when compared to the case where no time information is given. Overall, this experiment shows the importance of providing accurate time information when fusing asynchronous measurements, thus validating the importance of our *Discretiser* module.

TABLE II: Average RPE (m) results from the module ablation experiments. The minus sign (-) represents that the module is ablated. SE: Source Encoding, D: Discretiser.

Ablation	RMSE	Max	Mean \pm std
AFT-VO (ours)	0.04529	0.14983	0.03637 \pm 0.02644
-D-Equi	0.24042	0.82568	0.20731 \pm 0.11833
-D-None	0.16027	0.56651	0.13559 \pm 0.08446
-SE	0.04685	0.14979	0.03731 \pm 0.02778

C. Comparison Against the State-of-the-Art

In our next set of experiments, we compare the performance of our approach against state-of-the-art motion estimation techniques, namely monocular ORB-SLAM [20], DeepVO [29], MDN-VO [14] and an Extended Kalman Filter (EKF) based fusion.

We run ORB-SLAM with its global loop closure disabled in order to achieve a comparison in the context of VO,

TABLE III: Average RPE (m) results on the nuScenes test sequences, categorised by weather/lighting conditions. NR represents where an approach failed to produce results.

	Daylight (64 seq.)			Rain (12 seq.)			Night (24 seq.)		
	RMSE	Max	Mean \pm std	RMSE	Max	Mean \pm std	RMSE	Max	Mean \pm std
ORB-SLAM [20]	0.398	3.281	0.120 \pm 0.381	0.762	10.19	0.143 \pm 0.736	NR	NR	NR
DeepVO [29]	0.086	0.537	0.057 \pm 0.066	0.070	0.390	0.049 \pm 0.049	0.122	0.742	0.085 \pm 0.087
MDN-VO [14]	0.067	0.529	0.039 \pm 0.053	0.060	0.434	0.037 \pm 0.046	0.108	0.734	0.071 \pm 0.080
EKF [19]	0.070	0.352	0.046 \pm 0.052	0.065	0.368	0.044 \pm 0.047	0.160	0.895	0.089 \pm 0.129
AFT-VO (ours)	0.039	0.129	0.031 \pm 0.023	0.038	0.133	0.029 \pm 0.024	0.066	0.213	0.053 \pm 0.037

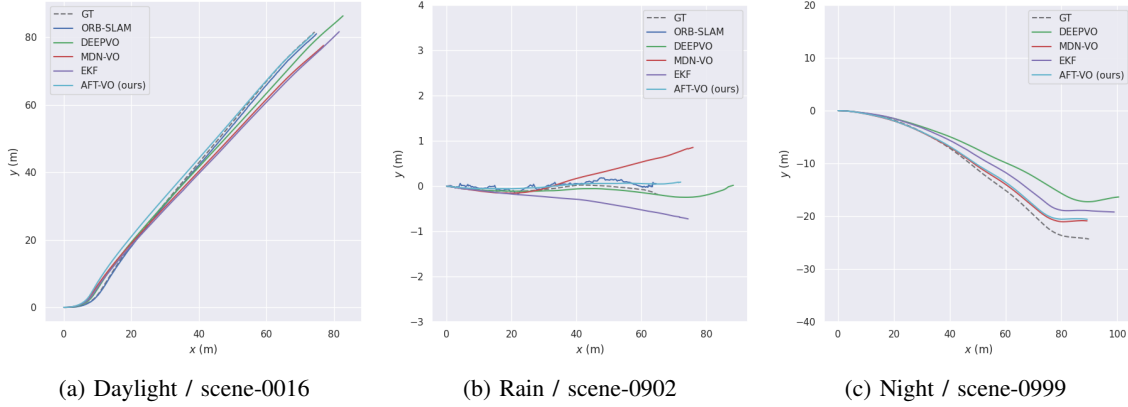


Fig. 2: Estimated VO trajectories on nuScenes - an example from every category (Daylight, Rain, Night) is shown.

as has been done previously in the literature [16], [30]. We use the PyTorch implementation [12] of DeepVO. In order to compare our results with a classic asynchronous fusion algorithm, we employ an EKF and use its ROS implementation [19].

We use the front camera images for DeepVO, MDN-VO and ORB-SLAM which is a standard in VO approaches. We feed the EKF with the VO and uncertainty estimations from the MDN module (See Section III-A). The reason behind this choice is to be able to compare the performance of our fusion module independently from its inputs. Note that although AFT-VO performance is comparable to the work of Kaygusuz *et al.* [15], their approach does not account for time in an explicit manner while fusing information coming from multiple cameras. Instead, they remove asynchronicity from the data by interpolating the missing time points for predictions coming from all cameras based on a common time axis, which has the advantage of looking into the future. Unlike their approach, our AFT-VO model does not need such an interpolation step due to the proposed time encoding and hence can fuse asynchronous camera information in real-time.

In our nuScenes experiments, we split the test set into three categories defined as: daylight, rain and night. The reason behind this choice is to provide a fair comparison for classical VO approaches. For example, as a classical, feature-based approach, ORB-SLAM could not initialise on the night-time sequences, due to insufficient illumination and texture. Thus, we could not report any night-time results for ORB-SLAM.

Average RPEs on the testing sequences with respect to their categories can be seen in Table III. As a classical approach, ORB-SLAM needs to track hand-crafted features

in consecutive frames. In our experiments, even though it achieves good performance in some scenarios, *i.e.* in good weather and illumination, it loses track in others which degrades its overall performance. By contrast, DeepVO and MDN-VO achieve better performance than ORB-SLAM in all categories. The EKF performs the second-best in daylight and rain. However, it performs worse than DeepVO and MDN-VO in night driving conditions. We believe this is due to the initial parameter selection for the Kalman Filter. We choose its initial parameters *e.g.* process noise, initial state covariance, that achieves the best performance across all test sequences. It is worth noting that, although EKF and our fusion approach use the MDN outputs (see Section III-A), our fusion module performs better than the EKF. As can be seen, our proposed method, AFT-VO, achieves the best performance in all three categories which validates our transformer-based fusion approach to estimate VO.

We demonstrate an example of estimated VO trajectories from three different categories in Figure 2. It can be seen that the VO trajectories estimated by our approach, AFT-VO, are generally more accurate than the other approaches. Although ORB-SLAM’s performance in daylight is successful, it tends to lose track in the rainy scenario which decreases its performance. As a learning-based approach, DeepVO and MDN-VO are able to produce results even for night scenes.

Results on the KITTI dataset can be seen in Table IV. As with the nuScenes dataset, AFT-VO achieves the best overall mean performance across all sequences. Even though ORB-SLAM achieves the best performance on Sequence 03, it has larger error on other sequences. In general, DeepVO and MDN-VO results are better than ORB-SLAM. However, AFT-VO achieves the best results on most of the sequences. It is worth noting that, even though KITTI has two syn-

TABLE IV: RPE (m) results on the KITTI Dataset

Sequence	ORB-SLAM [20]			DeepVO [16]			MDN-VO [14]			AFT-VO (ours)		
	RMSE	Max	Mean \pm std	RMSE	Max	Mean \pm std	RMSE	Max	Mean \pm std	RMSE	Max	Mean \pm std
03	0.03	0.20	0.03 \pm 0.02	0.08	0.23	0.08 \pm 0.04	0.12	0.41	0.11 \pm 0.06	0.09	0.29	0.08 \pm 0.05
05	0.25	0.67	0.20 \pm 0.15	0.24	0.57	0.21 \pm 0.12	0.16	0.36	0.15 \pm 0.08	0.10	0.26	0.08 \pm 0.04
06	0.34	0.65	0.30 \pm 0.17	0.16	0.31	0.14 \pm 0.08	0.20	0.45	0.18 \pm 0.09	0.17	0.33	0.15 \pm 0.08
07	0.17	0.35	0.13 \pm 0.09	0.14	0.35	0.12 \pm 0.07	0.08	0.51	0.07 \pm 0.05	0.10	0.29	0.09 \pm 0.04
10	0.30	0.95	0.23 \pm 0.20	0.21	0.47	0.19 \pm 0.08	0.14	0.32	0.13 \pm 0.06	0.13	0.32	0.12 \pm 0.06
mean	0.22	0.56	0.18 \pm 0.13	0.17	0.39	0.15 \pm 0.08	0.14	0.41	0.13 \pm 0.07	0.12	0.30	0.10 \pm 0.06

chronised cameras, our fusion approach achieves the best results on most sequences further validating the proposed transformer-based fusion approach.

V. CONCLUSION

In this paper, we presented a novel transformer-based fusion architecture to estimate VO. We formalised VO fusion as a sequence-to-sequence problem and employ transformers to fuse measurements from multiple asynchronous cameras. We proposed a novel *Discretiser* method which enables the network to positionally encode continuous timestamps, allowing transformer architectures to seamlessly handle asynchronous data. We also proposed *Source Encoding* to enable our fusion module to differentiate the source of measurements that it is getting. We evaluated AFT-VO on two popular datasets, namely nuScenes and KITTI, and reported state-of-the-art performance. As future work, we are planning to expand our asynchronous fusion approach by including other sensor modalities such as IMU.

VI. ACKNOWLEDGEMENTS

This paper was supported by the EPSRC project ROSSINI (EP/S016317/1).

REFERENCES

- [1] Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text. *NeurIPS*, 2021.
- [2] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, 2020.
- [3] Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. Sign language transformers: Joint end-to-end sign language recognition and translation. In *CVPR*, 2020.
- [4] Changhao Chen, Stefano Rosa, Yishu Miao, Chris Xiaoxuan Lu, Wei Wu, Andrew Markham, and Niki Trigoni. Selective sensor fusion for neural visual-inertial odometry. In *CVPR*, 2019.
- [5] Ronald Clark, Sen Wang, Hongkai Wen, Andrew Markham, and Niki Trigoni. Vinet: Visual-inertial odometry as a sequence-to-sequence learning problem. In *AAAI Conference on Artificial Intelligence*, 2017.
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019.
- [7] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *ICCV*, 2015.
- [8] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *PAMI*, 2017.
- [9] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *Int J Rob Res*, 2013.
- [10] Patrick Geneva, Kevin Eickenhoff, and Guoquan Huang. Asynchronous multi-sensor fusion for 3D mapping and localization. In *ICRA*, 2018.
- [11] Michael Grupp. evo: Python package for the evaluation of odometry and slam. Note: <https://github.com/MichaelGrupp/evo>, 2017.

- [12] Chi-Wei Hsiao. Note: <https://github.com/ChiWeiHsiao/DeepVO-pytorch>, 2017.
- [13] Kazuki Irie, Albert Zeyer, Ralf Schlüter, and Hermann Ney. Language modeling with deep transformers. *Proc. Interspeech*, 2019.
- [14] Nimet Kaygusuz, Oscar Mendez, and Richard Bowden. MDN-VO: Estimating visual odometry with confidence. In *IROS*, 2021.
- [15] Nimet Kaygusuz, Oscar Mendez, and Richard Bowden. Multi-camera sensor fusion for visual odometry using deep uncertainty estimation. In *ITSC*, 2021.
- [16] Ruihao Li, Sen Wang, Zhiqiang Long, and Dongbing Gu. Undeepvo: Monocular visual odometry through unsupervised deep learning. In *ICRA*, 2018.
- [17] Peidong Liu, Marcel Geppert, Lionel Heng, Torsten Sattler, Andreas Geiger, and Marc Pollefeys. Towards robust visual odometry with a multi-camera system. In *IROS*, 2018.
- [18] Vikram Mohanty and et al. Deepvo: A deep learning approach for monocular visual odometry. *arXiv preprint arXiv:1611.06069*, 2016.
- [19] Thomas Moore and Daniel Stouch. A generalized extended Kalman Filter implementation for the robot operating system. In *AIS 13*. 2016.
- [20] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. ORB-SLAM: a versatile and accurate monocular slam system. *IEEE transactions on robotics*, 2015.
- [21] Richard A Newcombe, Steven J Lovegrove, and Andrew J Davison. DTAM: Dense tracking and mapping in real-time. In *ICCV*, 2011.
- [22] David Nistér, Oleg Naroditsky, and James Bergen. Visual odometry. In *CVPR*, 2004.
- [23] Danda Pani Paudel, Cédric Demonceaux, Adlane Habed, and Pascal Vasseur. 2D–3D synchronous/asynchronous camera fusion for visual odometry. *Autonomous Robots*, 2019.
- [24] Jan Rosendahl, Viet Anh Khoa Tran, Weyue Wang, and Hermann Ney. Analysis of positional encodings for neural machine translation. *IWSLT, Hong Kong, China*, 2019.
- [25] Muhamad Risqi U Saputra, Pedro PB de Gusmao, Sen Wang, Andrew Markham, and Niki Trigoni. Learning monocular visual odometry through geometry-aware curriculum learning. In *ICRA*, 2019.
- [26] Joan Sola, André Monin, Michel Devy, and Teresa Vidal-Calleja. Fusing monocular information in multicamera SLAM. *IEEE transactions on robotics*, 2008.
- [27] Michael J Tribou, Adam Harmat, David WL Wang, Inna Sharf, and Steven L Waslander. Multi-camera parallel tracking and mapping with non-overlapping fields of view. *Int J Rob Res*, 2015.
- [28] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [29] Sen Wang, Ronald Clark, Hongkai Wen, and Niki Trigoni. Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks. In *ICRA*, 2017.
- [30] Sen Wang, Ronald Clark, Hongkai Wen, and Niki Trigoni. End-to-end, sequence-to-sequence probabilistic visual odometry through deep neural networks. *Int J Rob Res*, 2018.
- [31] Fei Xue, Xin Wang, Shunkai Li, Qiuyuan Wang, Junqiu Wang, and Hongbin Zha. Beyond tracking: Selecting memory and refining poses for deep visual odometry. In *CVPR*, 2019.
- [32] Anqi Joyce Yang, Can Cui, Ioan Andrei Bărsan, Raquel Urtasun, and Shenlong Wang. Asynchronous multi-view SLAM. In *ICRA*, 2021.
- [33] Huangying Zhan, Ravi Garg, Chamara Saroj Weerasekera, Kejie Li, Harsh Agarwal, and Ian Reid. Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In *CVPR*, 2018.
- [34] Chaofan Zhang, Yong Liu, Fan Wang, Yingwei Xia, and Wen Zhang. Vins-mkf: A tightly-coupled multi-keyframe visual-inertial odometry for accurate and robust state estimation. *Sensors*, 2018.
- [35] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017.