

AnyCam: Learning to Recover Camera Poses and Intrinsic from Casual Videos

Felix Wimbauer^{1,2,3}

Weirong Chen^{1,2,3}

Dominik Muhle^{1,2}

Christian Rupprecht³

Daniel Cremers^{1,2}

¹Technical University of Munich

²MCML

³University of Oxford

felix.wimbauer@tum.de

cremers@tum.de

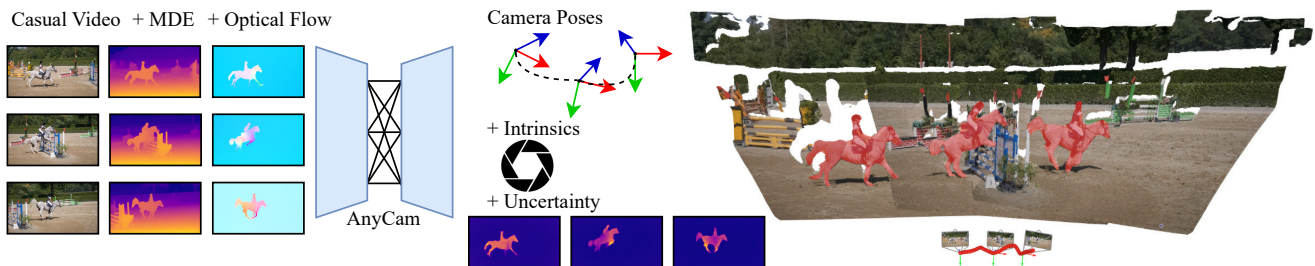


Figure 1. **AnyCam**. Given a casual video and pretrained monocular depth estimation (MDE) and optical flow networks, AnyCam outputs camera poses, camera intrinsics, and uncertainty maps in a single forward pass. The uncertainty maps represent probable movement in the scene. By using a novel loss formulation, AnyCam can be trained on a large corpus of unlabelled videos mostly obtained from YouTube.

Abstract

Estimating camera motion and intrinsics from casual videos is a core challenge in computer vision. Traditional bundle-adjustment based methods, such as SfM and SLAM, struggle to perform reliably on arbitrary data. Although specialized SfM approaches have been developed for handling dynamic scenes, they either require intrinsics or computationally expensive test-time optimization and often fall short in performance. Recently, methods like Dust3r have reformulated the SfM problem in a more data-driven way. While such techniques show promising results, they are still 1) not robust towards dynamic objects and 2) require labeled data for supervised training. As an alternative, we propose **AnyCam**, a fast transformer model that directly estimates camera poses and intrinsics from a dynamic video sequence in feed-forward fashion. Our intuition is that such a network can learn strong priors over realistic camera poses. To scale up our training, we rely on an uncertainty-based loss formulation and pre-trained depth and flow networks instead of motion or trajectory supervision. This allows us to use diverse, unlabelled video datasets obtained mostly from YouTube. Additionally, we ensure that the predicted trajectory does not accumulate drift over time through a lightweight trajectory refinement step. We test AnyCam on established datasets, where it delivers accurate camera poses and intrinsics both qualitatively and quantitatively. Furthermore, even with trajectory refinement, AnyCam is significantly faster than existing works for SfM in

dynamic settings. Finally, by combining camera information, uncertainty, and depth, our model can produce high-quality 4D pointclouds. For more details and code, please check out our project page: fwmb.github.io/anycam

1. Introduction

Estimating camera motion and intrinsics from a casual, dynamic video is a long-standing problem in 3D computer vision. Such scene parameters serve as the basis for a plethora of more complex applications, ranging from novel view synthesis to reconstruction. In particular, an efficient and robust solution to this problem would unlock the huge corpus of online video data from sources like YouTube to train 3D foundation models. To this day, the availability of 3D data is one of the main limitations when creating general-purpose 3D models. [33]

Traditional SfM (and SLAM) systems like COLMAP [32] excel at reconstruction in settings where a static scene is captured from many well-distributed camera angles. However, they generally fail when there are dynamic parts in the scene, the camera follows a suboptimal trajectory, or the image quality suffers from artifacts. This makes them unsuitable for casual videos. Recently, a number of works have proposed to incorporate certain deep learning-based components into or instead of the classic SfM pipeline to make it more robust. For example, Dust3r [44] proposes a network to predict dense 3D pointmaps from multiple

views. Camera pose and intrinsics can then be recovered by aligning the point maps from every view. FlowMap [34] reformulates the problem by learning per-frame depth and aligning them based on flow and point tracks. While these proposed approaches show promising results, they are still restricted to static scenes and fail when dynamic objects are too prominent. There are methods specifically designed to deal with dynamic scenes, *e.g.* LeapVO [5] and ParticleSfM [56]. However, they are generally trained in a supervised way and require calibrated camera intrinsics. Given the general lack of intrinsics information in casual videos, the applicability of such methods in this domain is limited.

This work presents **AnyCam**, a novel method targeted explicitly for robust camera motion and intrinsics estimation in dynamic casual videos. Unlike most existing approaches, we utilize a transformer-based model to *directly* predict relative camera poses and intrinsics for a sequence of video frames. Crucially, this end-to-end formulation allows our network to learn strong priors over plausible camera poses and to become robust towards imperfect inputs. Additionally, we design a training scheme that can ingest raw videos and does not require any labels. Our formulation effectively models data uncertainty and automatically filters out dynamic objects. This enables training on a collection of datasets obtained from YouTube and other casual video sources, resulting in strong generalization capabilities. Furthermore, we present a lightweight test-time refinement step to avoid long-term drift.

Our experiments thoroughly test AnyCam’s zero-shot capabilities in camera pose estimation and intrinsics recovery. We achieve state-of-the-art results on several challenging dynamic benchmarks and perform on par with methods trained fully supervised with labeled data. Additionally, by combining camera information, uncertainties, and depth maps, we can create high-fidelity 4D pointclouds.

2. Related Work

2.1. Foundational Models for Depth and Flow

In the tasks of monocular depth (MDE) and flow estimation, well-generalizable foundation models have replaced early deep learning approaches [8, 12, 13] in the last years. For MDE, DepthAnything [50] uses a data engine to construct a large corpus of automatically annotated data to learn relative depth estimation. Additional fine-tuning allows for metric depth estimates. DepthAnythingV2 [51] finetunes the previous model using synthetic data for better performance. Metric3D [53] and Metric3Dv2 [16] transform images to canonical camera intrinsics with a fixed focal length. DepthPro [1] proposes a two-stage training curriculum with a second stage solely on synthetic data to sharpen boundary predictions. DepthCrafter [17] leverages a conditional diffusion model to predict temporally consistent depth maps

for videos. In this work, we utilize UniDepth [28] for metric MDE, which uses a geometric invariance loss on different image augmentation to enforce consistency.

RAFT [38] presented the state of the art for optical flow estimation for a long time. It improved previous methods by introducing a recurrent look-up operator on correlation volumes to iteratively refine flow predictions without needing coarse-to-fine flow pyramids. GMFlow [46] avoids correlation volumes and instead leverages the properties of transformers for global matching on feature maps. This removes the need for iterative steps to improve runtime performance. UniMatch [47] extends GMFlow network by tasks of disparity and depth prediction to enable cross-task transfer learning of a single transformer network.

We rely on both off-the-shelf MDE and Optical Flow networks to benefit from strong geometric priors during training and inference.

2.2. SfM and SLAM

For many decades, the problem of recovering camera parameters and geometry from images has been formulated as the Structure-from-Motion (SfM) pipeline [14, 24, 26]. While many different implementations of the SfM pipeline exist, COLMAP [32] has emerged as the standard due to its robustness and flexibility. One of the drawbacks of SfM methods is their high computational cost. Simultaneous Location and Mapping (SLAM) [9, 10, 22, 23] approaches employ a similar pipeline to SfM but focus on the efficient processing of consecutive video frames. In recent years, these classical optimization-based approaches were enhanced by learned components [2, 7, 19, 21, 29–31, 52]. However, relying on epipolar geometry [15] or photometric consistency [10] makes them susceptible to high error on highly dynamic scenes. The strong focus on self driving data provided datasets with mostly static environments [4, 11, 36], an assumption that does not hold for casual videos.

2.3. Learning Based SfM and SLAM

Largely learning-based methods started to replace classical SLAM and SfM systems due to improved robustness [39]. DROID-SLAM extends the framework of RAFT [38] by an update operator on both depth and pose estimates. A final differentiable bundle adjustment (BA) layer produces the final pose estimates. ParticleSfM [56] utilizes dense correspondences inside a BA framework to optimize poses. The dense correspondences are initialized from optical flow, and dynamic points are filtered using trajectory-based motion segmentation. CasualSAM [55] predicts both depth and movement from images to get frame-to-frame motion. A global optimization aligns the scale of the prediction and refines the poses. Dust3R [44] is a dense multi-view stereo method that regresses point coordinates between an image pair. This allows it to be extended to either SfM or

SLAM. FlowMap [34] proposes to reconstruct a scene by overfitting a depth network to it and aligning depth maps via correspondences from flow or point tracking. LEAP-VO [5] combines visual and temporal information of video sequences to improve the tracking accuracy of points and identify occluded and dynamic points. A sliding window bundle adjustment then optimizes the poses. The concurrent work of MonST3R [54] finetunes Dust3r on mostly synthetic data to generalize it to dynamic scenes. While these works achieve impressive progress, they generally obtain poses from aligning depth and point maps or by optimizing them per-scene. This makes it hard to inject prior information about camera motion. In contrast, our method uses a neural network to predict a trajectory, which can effectively learn priors over realistic camera motions.

3. Method

In the following, we first describe our transformer-based model **AnyCam**, which directly predicts camera pose and intrinsics for video sequences. Further, we introduce an uncertainty-based training scheme, which enables training on large-scale dynamic video datasets without any ground truth labels, and a lightweight test-time refinement strategy.

3.1. Preliminaries

As input, our pipeline receives a sequence of n video frames $\mathbf{I}^i \in ([0, 1]^3)^\Omega, i \in N$, where $N = \{1, \dots, n\}$ denotes the set of frame indices and $\Omega = \{1, \dots, H\} \times \{1, \dots, W\}$ denotes the pixel lattice. Using off-the-shelf depth and optical flow predictors, we further obtain dense depth maps $\mathbf{D}^i \in \mathbb{R}_+^{1 \times H \times W}$, as well as dense optical flows $\mathbf{F}^{i \rightarrow j} \in \mathbb{R}^{2 \times H \times W}$ from any frame i to j . We generally assume a simplified pinhole model that is constant for the duration of a sequence. Therefore, camera intrinsics for a video are modeled via a single focal length $f \in \mathbb{R}_+$. Let $\pi_f(\mathbf{x}) : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ be the corresponding projection function that maps a point $\mathbf{x} \in \mathbb{R}^3$ in the camera’s coordinate system to the respective pixel $\mathbf{p} \in \mathbb{R}^2$ on the image plane parametrized by focal length f . Accordingly, $\pi_f^{-1}(\mathbf{p}, d) : (\mathbb{R}^2, \mathbb{R}_+) \rightarrow \mathbb{R}^3$ denotes the unprojection function where d is the depth value for \mathbf{p} . Let $\mathbf{P}^{i \rightarrow j} \in \text{SE3}$ describe the *relative* camera pose, consisting of a rotation and translation, between frames i and j . To map a pixel location \mathbf{p} from frame i to frame j given the pixel’s depth value d_p in frame i , we write:

$$\mathbf{p}' = \pi_f \left(\mathbf{P}^{i \rightarrow j} \pi_f^{-1}(\mathbf{p}, d_p) \right) \quad (1)$$

3.2. Transformer for Camera Prediction

Camera pose and intrinsics are crucial for many 3D tasks. While SfM and SLAM can be highly accurate in well-posed, static environments, they are very sensitive to outliers. Therefore, dynamic objects pose a significant challenge to them. Not only do they make convergence to the

correct solution more difficult, but they also often lead to catastrophic failures like loss of tracking or degenerate solutions. Different techniques were introduced to make the underlying optimizations more robust. These range from robust cost functions and RANSAC, to filtering out objects that could possibly move using off-the-shelf segmentation models. Nevertheless, these techniques are either hand-crafted, struggle to generalize across many domains, or present a significant computational overhead.

This paper aims to build a robust camera pose estimation system that can readily be applied to any casual video. We base our design on three key insights:

1. The space of plausible camera poses in videos is much more constrained than in general multi-view settings.
2. Strong, data-driven priors can help resolve ambiguities introduced by dynamic objects.
3. Incorporating information from all pixels instead of selected keypoints makes detecting outliers easier.

Pose from a feed-forward network. For a natural video captured by a single camera, the space of plausible camera movements is well-constrained and continuous. This is especially true when also considering the temporal context of videos. However, this prior information is hard to inject into classical SfM and SLAM approaches [45]. Hand-crafted priors often fail to generalize to different domains and struggle to capture the full variety of plausible motions. On the other hand, when trained on large and diverse datasets, neural networks can learn strong priors over the distribution of plausible camera poses. This knowledge makes them robust to small outliers or noisy measurements and provides reliable generalization abilities. Therefore, we choose a neural network for the prediction of frame-to-frame poses.

We rely on a transformer architecture, which as input receives frames \mathbf{I}^i , depth predictions \mathbf{D}^i , and flow predictions $\mathbf{F}^{i \rightarrow i+1}$ for a short sequence. First, a backbone extracts features for each timestep separately. Stacked self-attention layers allow the model to then exchange information between the frames, producing $n - 1$ pose tokens $\phi^{i \rightarrow i+1}, i \in \{1, \dots, n - 1\}$ and updated feature maps. Additionally, a single sequence token ϕ^{seq} is predicted.

To decode the features, a frame prediction head $\mathcal{H} = (\mathcal{H}^P, \mathcal{H}^\sigma)$ outputs the 6DOF camera poses $\mathbf{P}^{i \rightarrow i+1}$ and pixel-wise uncertainty maps $\sigma^i \in \mathbb{R}_+^\Omega$ for every timestep. A sequence head \mathcal{H}^{seq} additionally determines a single set of camera intrinsics for the video. In the following, we explain how to robustly design these decoder heads.

Robust intrinsics hypotheses. Camera intrinsics, which in our case are parametrized by focal length f , are notoriously difficult to disentangle from the camera movement. For example, the same motion pattern in the optical flow can be described by mostly rotation if the camera has a

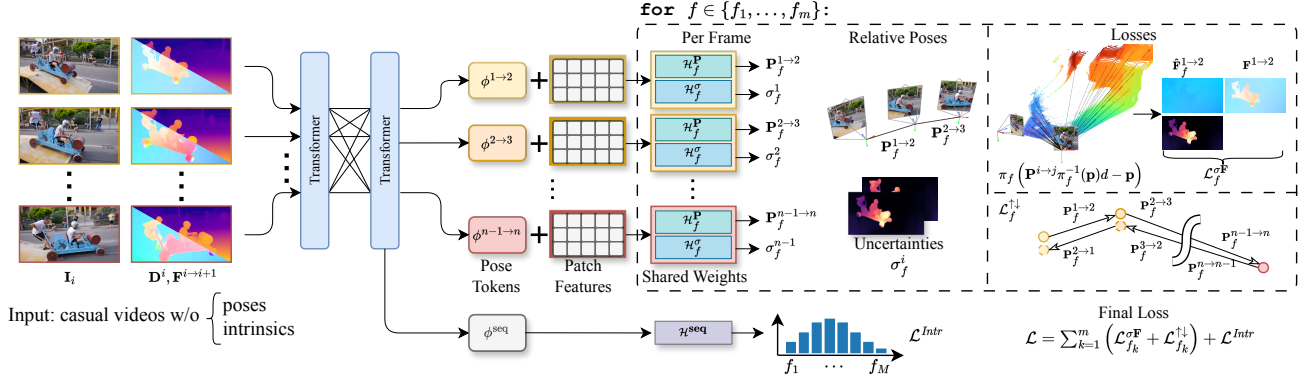


Figure 2. **Architecture.** AnyCam processes a sequence of frames from a casual video with corresponding depth maps and optical flow. A backbone extracts feature maps per image. Information sharing between frames is enabled by multiple attention layers that process the features of all sequence images. The transformer architecture outputs one pose token $\phi^{i \rightarrow j}$ per timestep and an additional sequence token ϕ^{seq} . The pose tokens are processed using multiple intrinsic hypotheses $f \in \{f_1, \dots, f_m\}$, parametrized by frame prediction heads $(\mathcal{H}_f^P, \mathcal{H}_f^\sigma)$. The sequence head \mathcal{H}^{seq} predicts the likelihood scores of the different hypotheses. The model is trained end-to-end via a reprojection loss, a pose consistency loss between forward and backward pose predictions, and a KL-divergence loss.

wide field of view or by mostly translation, if the field of view is narrow. Because camera intrinsics have such a major effect, designing a robust system to recover them is vital. While focal length could, in principle, also be predicted by the network directly, we find that it makes training unstable, and the network does not converge to a meaningful result

Instead of considering the focal length as a free variable during training, we reformulate it as a property of the model itself. Considering a fixed set of m candidate focal lengths $\{f_1, \dots, f_m\}$, we train m individual frame prediction heads $\{\mathcal{H}_{f_1}, \dots, \mathcal{H}_{f_m}\}$. Every head \mathcal{H}_f predicts pose $\mathbf{P}^{i \rightarrow i+1}$ and uncertainty map σ_f^i under the assumption that the given sequence was captured with a camera of focal length f . For every candidate prediction, we later compute an individual loss. The sequence head \mathcal{H}^{seq} learns likelihood scores $\mathcal{P} = (\rho_{f_1}, \dots, \rho_{f_m})$ for the different candidates. The focal length and corresponding poses with the highest likelihood are the final output of the model.

$$f_{final} = \arg \max_{f \in \{f_1, \dots, f_m\}} \mathcal{P} \quad (2)$$

This idea is similar to [34], but we use candidates as part of the *model* itself rather than only at the *loss* level.

3.3. Dynamics-aware Pose Training

We aim to train our model on large, unlabelled datasets. Here, our loss formulation follows two main objectives:

1. Leveraging multi-view information to recover local camera motion between adjacent frames.
2. Using the context of longer sequences to learn realistic long-range camera motion patterns.

Uncertainty-aware flow loss. Using the predicted relative pose $\mathbf{P}_f^{i \rightarrow j}$ between two frames i and j , and the corresponding depth map \mathbf{D}^i , we can project all pixel locations

\mathbf{p}_{uv} from frame i into frame j . Subtracting the original pixel locations from the projected pixels locations \mathbf{p}'_f yields the induced optical flow $\hat{\mathbf{F}}_f^{i \rightarrow j}$:

$$\begin{aligned} \hat{\mathbf{F}}_{f,uv}^{i \rightarrow j} &= \mathbf{p}'_{f,uv} - \mathbf{p}_{uv} \\ &= \pi_f \left(\mathbf{P}^{i \rightarrow j} \pi_f^{-1}(\mathbf{p}_{uv}, d_{uv}) \right) - \mathbf{p}_{uv} \end{aligned} \quad (3)$$

Assuming a static world, the induced flow $\hat{\mathbf{F}}_f^{i \rightarrow j}$ will match the reference optical flow $\mathbf{F}^{i \rightarrow j}$ optimally if the predicted pose $\mathbf{P}_f^{i \rightarrow j}$ and depth \mathbf{D}^i is as close to the real pose as possible. However, dynamic objects lead to inconsistencies in the reference optical flow, which would deteriorate gradients during optimization. From a statistical point of view, these inconsistencies cannot be captured via our flow induction formulation. Thus, we choose to model them via so-called aleatoric uncertainty and use the predicted uncertainty map \mathbf{U}_i . For a pair of frames i and j , this leads us to the following loss function:

$$\ell_{f,uv}^{\mathbf{F}^{i \rightarrow j}} = \left\| \hat{\mathbf{F}}_{f,uv}^{i \rightarrow j} - \mathbf{F}_{uv}^{i \rightarrow j} \right\|_1 \quad (4)$$

$$\mathcal{L}_f^{\sigma \mathbf{F}^{i \rightarrow j}} = -\frac{1}{|\Omega|} \sum_{uv \in \Omega} \ln \frac{1}{\sqrt{2}\sigma_{f,uv}^i} \exp -\frac{\sqrt{2}\ell_{f,uv}^{\mathbf{F}^{i \rightarrow j}}}{\sigma_{f,uv}^i} \quad (5)$$

which is then summed up over every pair of neighboring frames within the sequence:

$$\mathcal{L}_f^{\sigma \mathbf{F}} = \sum_{i=1}^{n-1} \mathcal{L}_f^{\sigma \mathbf{F}^{i \rightarrow i+1}} \quad (6)$$

Intuitively, the model learns to downweight areas in the input frame where there will likely be a high loss, *i.e.* dynamic objects. In turn, the pose supervision signal mostly comes from areas that can likely be captured via the induced

flow, *i.e.* static parts. In practice, the uncertainty map does not only make pose training more stable, but it can also be used in downstream tasks as a motion segmentation map. Finally, due to its dense nature, the loss is robust to small outliers, such as local inaccuracies of the reference flow.

Learning camera motion patterns. While the flow loss $\mathcal{L}_f^{\sigma F}$ provides a strong training signal *per frame*, the model could, in theory, minimize the loss without considering adjacent frames. However, it is important for the model to also rely on the context of the *entire sequence* when predicting camera poses. For example, optical flow predictions for some frames could be inaccurate, or dynamic objects could occlude the relevant static part. We introduce an additional loss term and a dropout training strategy to make the model consider the entire sequence.

We first leverage the fact that when reversing the sequence, the relative poses $\mathbf{P}_f^{j \rightarrow i}$ between frames are the same as the inverted poses of the original sequence $(\mathbf{P}_f^{i \rightarrow j})^{-1}$. The model becomes more robust towards inaccurate inputs by enforcing consistency between the forward and backward pose predictions during training. We rely on a simple L1 loss, where \mathbf{I}_4 is the 4×4 identity matrix:

$$\mathcal{L}_f^{\uparrow \downarrow} = \sum_{i=1}^{n-1} \left\| \left(\mathbf{P}_f^{i \rightarrow i+1} \right)^{-1} \mathbf{P}_f^{i+1 \rightarrow i} - \mathbf{I}_4 \right\|_{1,1} \quad (7)$$

Furthermore, we utilize the temporal dependency of poses within a sequence. Given the overall trajectory, the model should learn strong motion priors. To encourage this, we apply a dropout scheme during training by setting values in the pose tokens to zero with probability p_{drop} . This encourages the network to exchange information between frames in the self-attention layers. The network needs to learn realistic motion patterns to overcome the noise injected by dropout. Note that applying dropout on the pose tokens does not affect the uncertainty map.

Finding the best candidate. While the flow loss $\mathcal{L}_f^{\sigma F}$ is very effective for learning pose and uncertainty, it also serves as a helpful proxy. It indicates how well the model can reconstruct the observed optical flow using the predicted poses $\mathbf{P}_f^{i \rightarrow j}$ and corresponding focal length f . After the model has converged during training, it will predict (close to) the optimal poses assuming focal length f . $\mathcal{L}_f^{\sigma F}$ now is mainly dependent on f . Thus, the magnitude of $\mathcal{L}_f^{\sigma F}$ during training reflects how close f is to the true but unknown actual focal length. We leverage this fact and train the sequence head \mathcal{H}^{seq} to predict which of the focal candidates will yield the lowest flow loss.

In practice, we find that, while correct on average, arg min on the different loss terms is too noisy to enable stable training. Therefore, we first convert the (inverted)

flow loss terms to probability scores using softmax and then optimize the Kullback-Leibler divergence \mathbf{KL}_{div} . Let $\mathcal{P} = (\rho_{f_i}, \dots, \rho_{f_m})$ be the output of \mathcal{H}^{seq} :

$$\mathcal{L}^{Intr} = \mathbf{KL}_{div}(\mathcal{P}, \text{softmax}(-\mathcal{L}_{f_1}^{\sigma F}, \dots, -\mathcal{L}_{f_m}^{\sigma F})) \quad (8)$$

Final loss. By combining all loss terms, we obtain the final training objective, through which our model learns poses, uncertainties, and intrinsics from unlabelled video sequences:

$$\mathcal{L} = \sum_{k=1}^m \left(\lambda_{\sigma F} \mathcal{L}_{f_k}^{\sigma F} + \lambda_{\uparrow \downarrow} \mathcal{L}_{f_k}^{\uparrow \downarrow} \right) + \lambda_{Intr} \mathcal{L}^{Intr} \quad (9)$$

3.4. Test-time Refinement

Our model predicts relative camera poses for a sequence of frames. This allows our model to be directly used as a visual odometry system by chaining the relative poses together:

$$\mathbf{P}^j = \prod_{i=1}^j \mathbf{P}^{i \rightarrow i+1} \quad (10)$$

While this works well for short to medium-length sequences, it is prone to accumulating drift over a longer time. Small inaccuracies in earlier relative poses will propagate throughout the trajectory.

We perform lightweight test-time refinement using bundle adjustment (BA) [41] to overcome this issue. By chaining multiple optical flow maps $\mathbf{F}^{1 \rightarrow i+1}$, $\mathbf{F}^{i+1 \rightarrow i+2}$, \dots , pixels can be tracked through multiple frames. We initialize the BA system with our predicted trajectory and then optimize over coarse pixel tracks in a sliding-window fashion. Crucially, the BA optimization requires the predicted uncertainty maps to downweigh pixel tracks on dynamic objects. For more details on the BA optimization, please refer to the supplementary material.

4. Experiments

In the following, we thoroughly evaluate the overall performance of our model regarding camera pose estimation and intrinsics recovery. We put a special focus on testing the zero-shot generalization capabilities of our network. Further, we also validate and demonstrate the effectiveness of our design choices.

4.1. Setup

Data. Our training formulation requires no labeled data and is robust towards dynamic objects and suboptimal image quality. This is a key benefit of our method, as dynamic videos with 3D labels are scarce. To highlight this strength, we rely on a diverse mix of datasets (YouTube-VOS [48], RealEstate10K [57], WalkingTours [42], OpenDV [49],

| Category | Method | No Supervis. | Approx. Runtime | Sintel | | | TUM-RGBD (dynamics) | | |
|----------------|---------------------------|--------------|-----------------|--------------|------------------------|----------------------|---------------------|------------------------|----------------------|
| | | | | ATE↓ | RPE _{trans} ↓ | RPE _{rot} ↓ | ATE↓ | RPE _{trans} ↓ | RPE _{rot} ↓ |
| w/ intrinsics* | DROID-SLAM [39] | | <2min | 0.175 | 0.084 | 1.912 | - | - | - |
| | DPVO [40] | | <2min | 0.115 | 0.072 | 1.975 | - | - | - |
| | LeapVO [5] | | <2min | 0.089 | 0.066 | 1.250 | 0.068 | 0.008 | 1.686 |
| w/o intrinsics | ParticleSfM [56] | | <20min | 0.129 | 0.031 | <u>0.525</u> | - | - | - |
| | MonST3R [†] [54] | | <2min | 0.108 | 0.042 | 0.732 | <u>0.063</u> | <u>0.009</u> | 1.217 |
| | Robust-CVD [18] | ✓ | ? | 0.360 | 0.154 | 3.443 | 0.153 | 0.026 | 3.528 |
| | CasualSAM [55] | ✓ | >1h | 0.141 | <u>0.035</u> | 0.615 | 0.071 | 0.010 | 1.712 |
| | Ours w/o ref. | ✓ | <20sec | <u>0.099</u> | 0.045 | 0.567 | 0.095 | 0.025 | <u>1.050</u> |
| | Ours | ✓ | <2min | 0.078 | 0.031 | 0.427 | 0.056 | 0.005 | 0.927 |

Table 1. **Pose estimation in dynamic environments.** Absolute trajectory error (ATE) and relative pose error for translation (RPE_{trans}) and rotation (RPE_{rot}) on the Sintel and TUM-RGBD (dynamics) datasets. We compare against other learning-based VO/SLAM systems. AnyCam achieves competitive performance against systems trained in a supervised manner as well as methods that have ground truth camera intrinsics available. We consistently outperform Robust-CVD and are on par with CasualSAM, which has a significantly higher runtime than AnyCam. *excluded from the evaluation. [†]concurrent work. Many measurements are taken over from [54].

| | Dataset | Source | Domain | Dyn. | Seq. | Frames |
|--------------|---------------|------------|-----------|------|------|--------|
| <i>Train</i> | YouTube-VOS | YouTube | Diverse | ✓ | 3471 | 87K |
| | RealEstate10K | YouTube | Indoor | ✗ | 6929 | 90K |
| | WalkingTours | YouTube | Outdoor | ✓ | 278 | 53K |
| | OpenDV | YouTube | Driving | ✓ | 532 | 105K |
| | EpicKitchens | GoPro | Ego | ✓ | 167 | 100K |
| <i>Test</i> | Sintel | Blender | Synthetic | ✓ | 14 | 629 |
| | TUM-RGBD | MoCap | Indoor | ✓ | 8 | 720 |
| | Davis | YouTube | Diverse | ✓ | 90 | 6118 |
| | Aria EA | Glasses | Ego | ✓ | 30 | 19200 |
| | Waymo | Sensor Rig | Driving | ✓ | 64 | 3067 |

Table 2. **Dataset Mix:** We train our method on a diverse dataset mix obtained mostly from YouTube. Evaluation happens on unseen datasets. For many datasets, custom splits were used.

EpicKitchens [6]) based on YouTube or individual GoPro capture, as shown in Tab. 2. None of the datasets has ground truth 3D labels (note that for some, Colmap was later applied to obtain proxy labels, but we do not use them). Following existing works on pose estimation in dynamic environments [5, 54, 56], we perform evaluation on the Sintel [3] and the dynamic subset of TUM-RGBD [35]. Furthermore, we test AnyCam qualitatively on three other datasets: Davis [27] (diverse videos from YouTube), Waymo [37] (autonomous driving), and Aria Everyday Activities [20].

Implementation details. We rely on the recent UniDepth [28] and UniMatch [47] methods to obtain depth and flow maps. Our model is implemented in PyTorch and is initialized with a pretrained DinoV2 [25] ViT backbone. We train AnyCam in two stages with different sequence lengths, first 2 and then 8 frames. All frames are sampled at a resolution of 336×336 from the videos in the datasets. During inference, we crop the input video to squares of 336 and pass sequences of up to 100 frames to the model at once. The model is configured to use 32 focal length candidates

ranging from $0.1H$ to $3.5H$ where H is the image height. Training converges after 250k iterations at a batch size of 16 (seq. len 2) and 4 (seq. len 8) per GPU, and takes around two days on two NVIDIA A100 40GB GPUs per stage. For a video of 50 frames, it takes around 1) 15 seconds to obtain flows and depths, 2) 5 seconds for AnyCam to predict a trajectory, and 3) 90 seconds for test-time refinement. For more implementation details and hyperparameters, please refer to the supplementary material.

4.2. Camera Pose Estimation

We first test AnyCam’s ability to recover camera trajectory in challenging environments and compare it to state-of-the-art methods in that domain. We group methods by the amount of data they require both during training and test time. SLAM and SfM systems for dynamic environments, like DPVO [40] and LeapVO [5], are trained with ground truth motion or trajectory data and mostly require camera intrinsics during test time. Methods like CasualSAM [55] require neither special training data nor intrinsics but are slow because they rely on much more costly test-time optimization. In comparison, AnyCam works in a feed-forward way, and optional test-time refinement is very lightweight.

Our model achieves strong quantitative results throughout all benchmarks as shown in Tab. 1. The low relative pose errors for translation (RPE_{trans}) and rotation (RPE_{rot}), even without test-time refinement, confirm that our model learns meaningful motion priors during training and generalizes to other datasets. In particular, this is noteworthy compared to specialized SLAM systems like LeapVO, which result in significantly higher local pose errors.

However, when only relying on feed-forward predictions, our model can suffer from drift, denoted by a slightly higher absolute trajectory error (ATE). This is a result of chaining the predicted relative poses $\mathbf{P}^{i \rightarrow j}$, where small errors in earlier frames accumulate over time. Our test-time

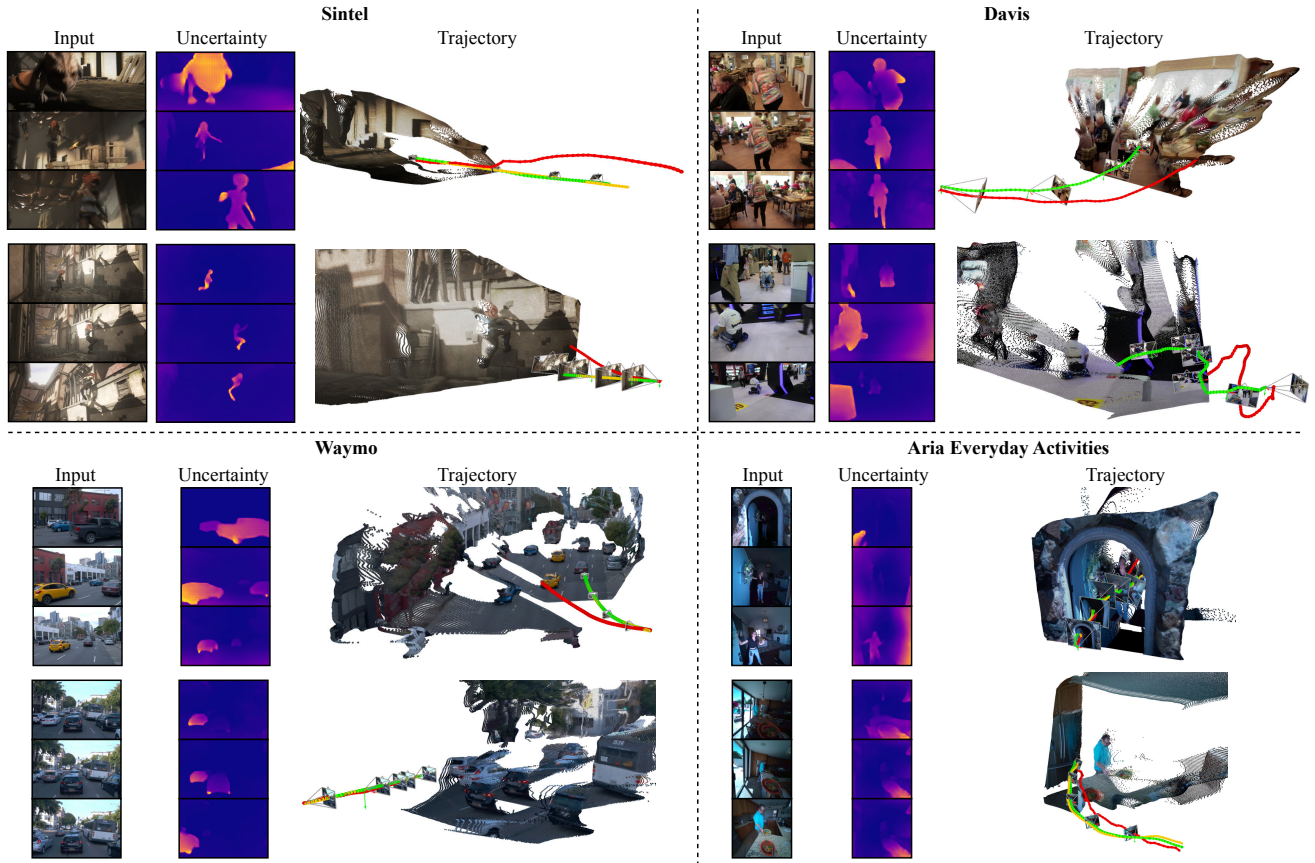


Figure 3. **Qualitative results on various datasets.** Red: forward-pass prediction, Green: refined trajectory, Yellow: GT (if available). AnyCam is able to predict high-quality pose estimates on challenging scenes in dynamic environments. The uncertainty maps show objects with a high likelihood of movement, such as persons or cars, that would produce inconsistencies in the induced optical flow. Pose refinement with bundle adjustment further aligns the trajectory towards reducing the error compared to the ground truth poses.

refinement effectively corrects for this by leveraging longer-range dependencies between frames. As a result, the drift is reduced, and AnyCam matches the ATE for some SLAM systems like LeapVO, which requires camera intrinsics.

As can be seen in Tab. 1, AnyCam also generalizes to different domains like autonomous driving or egocentric vision. In particular, it works well in settings where classical SfM or SLAM systems are difficult to operate in. For example, in the first sequence of Aria Everyday Activities, a person enters a house via an opening door. Not only does the lighting change drastically, but also the 3D environment transforms as the door opens. Still, the camera motion trajectory is recovered correctly. During the second Waymo sequence, a vehicle moves in stop-and-go traffic. As can be observed in the uncertainty maps, AnyCam correctly identifies the moving vehicles on the left side as moving, while the uncertainty for the bus on the right is low. This indicates that the model actually relies on motion cues and not only on semantics, which would be the case when filtering out (potentially) moving objects with segmentation models.

| Method | AFE _(px) ↓ | RFE(%) ↓ |
|----------|-----------------------|--------------|
| UniDepth | 447.4 | 0.357 |
| Dust3r | 434.0 | 0.364 |
| Ours | 252.2 | 0.181 |

Table 3. **Camera intrinsics estimation.** We measure the mean absolute focal error (AFE) and mean relative focal error (RFE) across all sequences of the Sintel dataset.

4.3. Intrinsic Recovery

For most casual videos, particularly when they originate from the internet, camera intrinsics are not available. Nevertheless, they are crucial for high-quality pose estimation and reconstruction. Therefore, we also evaluate the accuracy of the recovered focal length and report the results in Tab. 3. UniDepth predicts depth maps and focal length from a single image. Our superior performance against UniDepth (which provides the depth maps for AnyCam) suggests that our video-based training and inference offer benefits for fo-

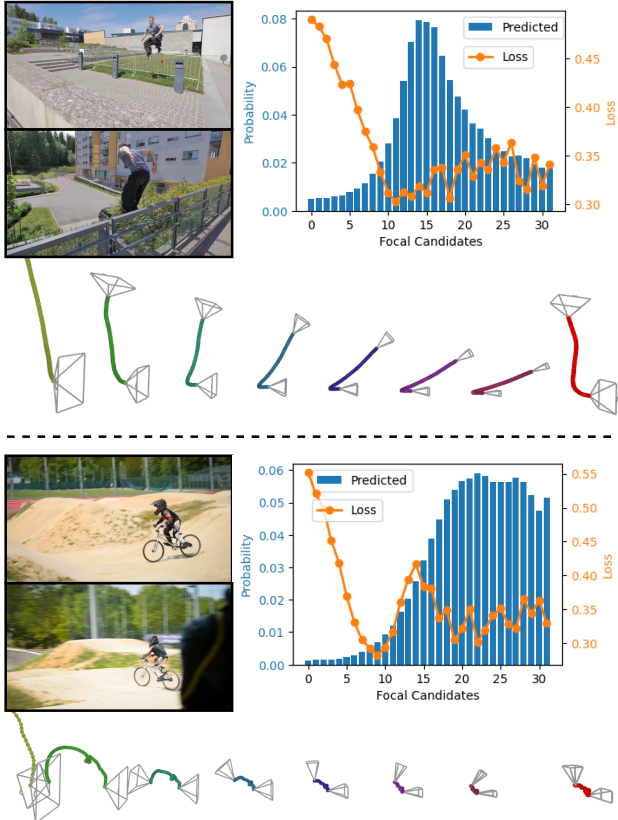


Figure 4. **Focal length candidates.** Predicted likelihood and computed flow loss for different focal candidate (FL) hypotheses for two sequences of the Davis dataset. **Below:** Estimated trajectories for selected FL hypotheses increasing from left to right. **Red** trajectory shows the trajectory for the FL with the highest likelihood. The predicted likelihood tends to be more stable than the loss when estimating the best candidate.

cal length recovery. While Dust3r also processes sequential frames, it is trained on classical multi-view settings and fails to generalize well to challenging dynamic scenarios.

Additionally, we test our approach of training several prediction heads assuming different focal length candidates. As seen in Fig. 4, the individual heads predict significantly varying pose trajectories. For higher focal lengths on the right side, motion is modeled more through translation and only small rotations. The lower the focal length, the more the network predicts strong rotation. This highlights the importance of selecting the correct focal length candidate.

In both examples, the predicted likelihood scores seem to pick a candidate that is very close to the actual focal length, and the selected trajectory is plausible. While the calculated loss is a very good indicator of which focal length is optimal, it can be noisy and is not guaranteed to always be correct. In the second example, the field-of-view is narrow, indicating a high focal length. However, the minimal loss is achieved for a fairly low focal length. We hypothesize that

| Seq. len. | | Refinement | | Sintel | | |
|-----------|----------|------------|--------|------------------|---------------------------------|-------------------------------|
| Train | Infer | Simple | Uncert | ATE \downarrow | RPE $_{\text{trans}}\downarrow$ | RPE $_{\text{rot}}\downarrow$ |
| 2 | 2 | | | 0.134 | 0.045 | 0.393 |
| 2 | ≥ 8 | | | 0.179 | 0.052 | 0.537 |
| 8 | 2 | | | 0.182 | 0.053 | 0.617 |
| 8 | ≥ 8 | | | <u>0.099</u> | 0.045 | 0.567 |
| 8 | ≥ 8 | ✓ | | 0.136 | <u>0.036</u> | 0.440 |
| 8 | ≥ 8 | | ✓ | 0.078 | 0.031 | <u>0.427</u> |

Table 4. **Model ablation on Sintel:** We test the effect of 1) sequence length during training and inference, and 2) enabling refinement with or without uncertainty.

the model learned to effectively filter out this noise, leading to a more stable and more accurate predictor compared to selecting the candidate with the lowest loss.

4.4. Model Ablations

Finally, we evaluate the design choices that went into our model and report the results in Tab. 4. First, we test whether AnyCam can learn meaningful priors from sequences. For this, a second model *AnyCam₂* trained only on sequences of length 2 (*i.e.* pairs of frames) is used as comparison. This model only predicts a single relative pose between the two frames and does not learn to use the context of longer sequences. The RPE $_{\text{trans}}$ and RPE $_{\text{rot}}$ metrics indicate that *AnyCam₂* can still perform accurate prediction between pairs of frames. However, the poor ATE results show that it leads to a noticeably higher drift. When passing a longer sequence to *AnyCam₂* at test time, the model cannot use this additional information and even gets confused, leading to overall poor results. Similarly, the final trajectory will suffer from drift when training with sequences of 8 frames but only passing longer sequences to the model at inference. Only provided longer sequences, the model can utilize the learned priors, leading to improved overall scores.

Additionally, we verify the effectiveness of our refinement strategy. Naive bundle adjustment struggles to provide an improvement, even when initialized with the predicted trajectory and focal length. Only when also weighing the pixel tracks through AnyCam’s uncertainty, bundle adjustment can improve the overall trajectory quality.

5. Conclusion

AnyCam demonstrates that pose estimation and intrinsics recovery can effectively be performed by a feed-forward network with lightweight test-time refinement. It is possible to achieve competitive performance with supervised learning methods in the tasks of SfM and SLAM. Through a novel training pipeline, AnyCam can be trained on a large corpus of casual video data.

Acknowledgements: This work was funded by the ERC Advanced Grant “SIMULACRON” (agreement #884679), the GNI Project “AI4Twinning”, and the DFG project CR 250/26-1 “4D YouTube”.

References

- [1] Aleksei Bochkovskii, Amaël Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan R Richter, and Vladlen Koltun. Depth pro: Sharp monocular metric depth in less than a second. *arXiv preprint arXiv:2410.02073*, 2024. [2](#)
- [2] Eric Brachmann, Alexander Krull, Sebastian Nowozin, Jamie Shotton, Frank Michel, Stefan Gumhold, and Carsten Rother. Dsac-differentiable ransac for camera localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6684–6692, 2017. [2](#)
- [3] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *European Conf. on Computer Vision (ECCV)*, pages 611–625. Springer-Verlag, 2012. [6](#)
- [4] 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 multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. [2](#)
- [5] Weirong Chen, Le Chen, Rui Wang, and Marc Pollefeys. Leap-vo: Long-term effective any point tracking for visual odometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19844–19853, 2024. [2](#), [3](#), [6](#)
- [6] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *Proceedings of the European conference on computer vision (ECCV)*, pages 720–736, 2018. [6](#)
- [7] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236, 2018. [2](#)
- [8] 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 *Proceedings of the IEEE international conference on computer vision*, pages 2758–2766, 2015. [2](#)
- [9] Jakob Engel, Thomas Schöps, and Daniel Cremers. Lsd-slam: Large-scale direct monocular slam. In *European conference on computer vision*, pages 834–849. Springer, 2014. [2](#)
- [10] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence*, 40(3):611–625, 2017. [2](#)
- [11] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. [2](#)
- [12] 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](#)
- [13] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3828–3838, 2019. [2](#)
- [14] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. [2](#)
- [15] Richard I Hartley. In defense of the eight-point algorithm. *IEEE Transactions on pattern analysis and machine intelligence*, 19(6):580–593, 1997. [2](#)
- [16] Mu Hu, Wei Yin, Chi Zhang, Zhipeng Cai, Xiaoxiao Long, Hao Chen, Kaixuan Wang, Gang Yu, Chunhua Shen, and Shaojie Shen. Metric3d v2: A versatile monocular geometric foundation model for zero-shot metric depth and surface normal estimation. *arXiv preprint arXiv:2404.15506*, 2024. [2](#)
- [17] Wenbo Hu, Xiangjun Gao, Xiaoyu Li, Sijie Zhao, Xiaodong Cun, Yong Zhang, Long Quan, and Ying Shan. DepthCrafter: Generating consistent long depth sequences for open-world videos. *arXiv preprint arXiv:2409.02095*, 2024. [2](#)
- [18] Johannes Kopf, Xuejian Rong, and Jia-Bin Huang. Robust consistent video depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1611–1621, 2021. [6](#)
- [19] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. *arXiv preprint arXiv:2406.09756*, 2024. [2](#)
- [20] Zhaoyang Lv, Nicholas Charron, Pierre Moulon, Alexander Gamino, Cheng Peng, Chris Sweeney, Edward Miller, Huixuan Tang, Jeff Meissner, Jing Dong, et al. Aria everyday activities dataset. *arXiv preprint arXiv:2402.13349*, 2024. [6](#)
- [21] Dominik Muhle, Lukas Koestler, Krishna Murthy Jatavallabhula, and Daniel Cremers. Learning correspondence uncertainty via differentiable nonlinear least squares. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13102–13112, 2023. [2](#)
- [22] Raul Mur-Artal and Juan D Tardós. Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE transactions on robotics*, 33(5):1255–1262, 2017. [2](#)
- [23] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics*, 31(5):1147–1163, 2015. [2](#)
- [24] John Oliensis. A critique of structure-from-motion algorithms. *Computer Vision and Image Understanding*, 80(2): 172–214, 2000. [2](#)
- [25] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. [6](#)
- [26] Onur Özyeşil, Vladislav Voroninski, Ronen Basri, and Amit Singer. A survey of structure from motion*. *Acta Numerica*, 26:305–364, 2017. [2](#)

- [27] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 724–732, 2016. 6
- [28] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. Unidepth: Universal monocular metric depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10106–10116, 2024. 2, 6
- [29] René Ranftl and Vladlen Koltun. Deep fundamental matrix estimation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 284–299, 2018. 2
- [30] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4938–4947, 2020.
- [31] Paul-Edouard Sarlin, Ajaykumar Unagar, Mans Larsson, Hugo Germain, Carl Toft, Viktor Larsson, Marc Pollefeys, Vincent Lepetit, Lars Hammarstrand, Fredrik Kahl, et al. Back to the feature: Learning robust camera localization from pixels to pose. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3247–3257, 2021. 2
- [32] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. 1, 2
- [33] Yawar Siddiqui, Tom Monnier, Filippos Kokkinos, Mahendra Kariya, Yanir Kleiman, Emilien Garreau, Oran Gafni, Natalia Neverova, Andrea Vedaldi, Roman Shapovalov, et al. Meta 3d assetgen: Text-to-mesh generation with high-quality geometry, texture, and pbr materials. *arXiv preprint arXiv:2407.02445*, 2024. 1
- [34] Cameron Smith, David Charatan, Ayush Tewari, and Vincent Sitzmann. Flowmap: High-quality camera poses, intrinsics, and depth via gradient descent. *arXiv preprint arXiv:2404.15259*, 2024. 2, 3, 4
- [35] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of rgb-d slam systems. In *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, 2012. 6
- [36] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020. 2
- [37] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020. 6
- [38] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 402–419. Springer, 2020. 2
- [39] Zachary Teed and Jia Deng. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in neural information processing systems*, 34:16558–16569, 2021. 2, 6
- [40] Zachary Teed, Lahav Lipson, and Jia Deng. Deep patch visual odometry. *Advances in Neural Information Processing Systems*, 36, 2024. 6, 3
- [41] Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. Bundle adjustment—a modern synthesis. In *Vision Algorithms: Theory and Practice: International Workshop on Vision Algorithms Corfu, Greece, September 21–22, 1999 Proceedings*, pages 298–372. Springer, 2000. 5
- [42] Shashanka Venkataramanan, Mamshad Nayeem Rizve, João Carreira, Yuki M Asano, and Yannis Avrithis. Is imagenet worth 1 video? learning strong image encoders from 1 long unlabelled video. In *International Conference on Learning Representations*, 2024. 5
- [43] Jianyuan Wang, Nikita Karaev, Christian Rupprecht, and David Novotny. Visual geometry grounded deep structure from motion. *arXiv preprint arXiv:2312.04563*, 2023. 3
- [44] Shuzhe Wang, Vincent Leroy, Johann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20697–20709, 2024. 1, 2
- [45] Jeremy N. Wong, David J. Yoon, Angela P. Schoellig, and Timothy D. Barfoot. A data-driven motion prior for continuous-time trajectory estimation on se(3). *IEEE Robotics and Automation Letters*, 2020. 3
- [46] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezafofighi, and Dacheng Tao. Gmflow: Learning optical flow via global matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8121–8130, 2022. 2
- [47] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezafofighi, Fisher Yu, Dacheng Tao, and Andreas Geiger. Unifying flow, stereo and depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2, 6
- [48] Ning Xu, Linjie Yang, Yuchen Fan, Jianchao Yang, Dingcheng Yue, Yuchen Liang, Brian Price, Scott Cohen, and Thomas Huang. Youtube-vos: Sequence-to-sequence video object segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 585–601, 2018. 5
- [49] Jiazhi Yang, Shenyuan Gao, Yihang Qiu, Li Chen, Tianyu Li, Bo Dai, Kashyap Chitta, Penghao Wu, Jia Zeng, Ping Luo, et al. Generalized predictive model for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14662–14672, 2024. 5
- [50] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of*

- the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10371–10381, 2024. [2](#)
- [51] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv preprint arXiv:2406.09414*, 2024. [2](#)
- [52] Nan Yang, Rui Wang, Jorg Stuckler, and Daniel Cremers. Deep virtual stereo odometry: Leveraging deep depth prediction for monocular direct sparse odometry. In *Proceedings of the European conference on computer vision (ECCV)*, pages 817–833, 2018. [2](#)
- [53] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3d: Towards zero-shot metric 3d prediction from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9043–9053, 2023. [2](#)
- [54] Junyi Zhang, Charles Herrmann, Junhwa Hur, Varun Jampani, Trevor Darrell, Forrester Cole, Deqing Sun, and Ming-Hsuan Yang. Monst3r: A simple approach for estimating geometry in the presence of motion. *arXiv preprint arXiv:2410.03825*, 2024. [3](#), [6](#)
- [55] Zhoutong Zhang, Forrester Cole, Zhengqi Li, Michael Rubinstein, Noah Snavely, and William T Freeman. Structure and motion from casual videos. In *European Conference on Computer Vision*, pages 20–37. Springer, 2022. [2](#), [6](#)
- [56] Wang Zhao, Shaohui Liu, Hengkai Guo, Wenping Wang, and Yong-Jin Liu. Particlesfm: Exploiting dense point trajectories for localizing moving cameras in the wild. In *European Conference on Computer Vision*, pages 523–542. Springer, 2022. [2](#), [6](#)
- [57] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817*, 2018. [5](#)

AnyCam: Learning to Recover Camera Poses and Intrinsics from Casual Videos

Supplementary Material

A. Overview

In this supplementary, we explain more details on implementation in Appendix C and the test-time refinement process in Appendix D. We also consider limitations and future work in Appendix E and ethical implications in Appendix F.

B. Code & Project Page

We release our code base for training, evaluation, and visualization under github.com/Brummi/anycam. Additionally, we provide interactive 3D results and more details on our project page under fwmb.github.io/anycam.

C. Further Implementation Details

Focal Length Candidates. In our model, we configure $m = 32$ distinct focal length candidates. For every candidate f_i , we train individual prediction heads \mathcal{H}_{f_i} . Focal length is not linearly related to rotation and translation magnitudes. Empirically, we find that the following formula, which combines linear and exponential spacing, provides a good distribution of focal length candidates.

$$\delta_i = \frac{i}{m-1} \quad (11)$$

$$f_i^{\text{exp}} = \exp(\delta_i \log(f_{\min}) + (1 - \delta_i) \log(f_{\max})) \quad (12)$$

$$f_i^{\text{lin}} = \delta_i f_{\min} + (1 - \delta_i) f_{\max} \quad (13)$$

$$f_i = 0.75 \cdot f_i^{\text{exp}} + 0.25 \cdot f_i^{\text{lin}} \quad (14)$$

We define $f_{\min} = 0.1H$ and $f_{\max} = 3.5H$, where H represents the height of the input image in pixels, yielding the distribution which can be seen in Fig. 5. In this way, the model can make predictions independent of the pixel size.

Camera Pose Parametrization. The prediction heads do not directly output the $P \in \mathbb{R}^{4 \times 4}$. To ensure that the pose matrix is in SE3, we predict translation $t \in \mathbb{R}^3$ and rotation $R \in \text{SO}3$ separately. R is parametrized via the axis-angle representation, *i.e.* the model predicts three values for the different axis rotations. We find that the axis-angle representation is significantly more stable than the quaternion representation and it converges faster. When using quaternions, it usually happens that a small number of (random) prediction heads does not converge to meaningful results.

Training Stabilization. Our training datasets cover a diverse range of datasets, which all have varying scales. E.g. driving datasets depict scenes and movements much larger




Figure 5. **Focal Length Candidates.** Linear-exponential distribution of focal length candidates relative to the image height.

compared to video sequences captured from VR glasses. When naively training on all five datasets from the start, the model does not converge to a meaningful solution. We hypothesize that the different scales introduce noise that hinders the optimization process. To overcome this issue, we first undergo a warmup phase, during which datasets are introduced one-by-one. First, the model is trained for 10,000 steps on RealEstate10K, then for another 10,000 steps on RealEstate10K and EpicKitchens, and so on until all datasets have been introduced. Through this strategy, the model can already roughly estimate the camera pose and then only is adapted to a different scale.

Loss Configuration. The model is trained using the Adam optimizer at a learning rate of $\epsilon = 1e^{-4}$. After 100,000 steps, the learning rate is reduced to $\epsilon = 1e^{-5}$. We use $\lambda_{\sigma\mathbf{F}} = 1$, $\lambda_{\uparrow\downarrow} = 1$, and $\lambda_{\text{Intr}} = 1$. Since the flow loss values tend to have a very small magnitude, we set the temperature of the softmax operator in $\mathcal{L}_{\text{Intr}}$ to 100. Note that we also detach the flow losses in $\mathcal{L}_{\text{Intr}}$ and only pass gradients to the sequence head. This ensures that the different candidate heads remain independent of each other. Finally, we also apply L2 weight decay with a factor of 0.01 on the pose tokens to avoid overflow issues when training with mixed precision.

Model Architecture. We adapt the DinoV2 based DepthAnything model to predict both a pixel aligned map and tokens for pose and intrinsics. For both our backbone and UniDepth, we rely on Vit-S.

D. Test-Time Refinement Details

The main objective of our test-time refinement strategy is to reduce drift over longer time frames. To ensure geometric consistency over time, we apply bundle adjustment (BA) and optimize the camera trajectory in a sliding window fashion.

Setup. Since we primarily care about long-range dependencies, we apply BA with a stride of 3 frames. After the optimization is complete, the remaining poses are then interpolated and combined with the original predictions. During BA, we sample a uniform 16×16 grid of points per frame and then track them for 8 consecutive frames. Tracking is performed by chaining optical flow maps and we additionally accumulate uncertainty per tracked point. The uncertainty of a tracked point at a specific frame is the sum of uncertainties from all previous frames of the track. Intuitively, this means that a track only has low uncertainty as long as it does not encounter a pixel that has high uncertainty. For every point track, we optimize a single 3D point anchored in the first frame of the track and parametrize it by inverse depth. The point is initialized via the predicted depth that was also used as input by the AnyCam model. Thus, a track $T_{f,xy}$ starting in frame f at grid location x, y is defined by

$$T_{f,xy} = ((\mathbf{p}_1, \dots, \mathbf{p}_8), (\sigma_1, \dots, \sigma_8), d^{-1}) \quad (15)$$

where \mathbf{p}_j are the pixel locations in the consecutive frames, σ_j are the corresponding uncertainty values, and d^{-1} is the inverse depth of the anchor point. In total, we optimize 1) the camera poses, 2) the inverse depths of the anchor points, and 3) a single focal length value.

Optimization. The main objective of our optimization process is to minimize the reprojection error for every track. In fact, we rely on a similar formulation as the flow loss $\ell_{f,uv}^{\mathbf{F}^{i \rightarrow j}}$ used for training AnyCam. To completely filter out very dynamic objects, we define a maximum uncertainty $\sigma_{\max} = 0.05$ and ignore all points that exceed this threshold. All others are weighted accordingly in a linear fashion. Let $\mathbf{T} = \{T_{000}, \dots\}$ be the set of all tracks in the sequence:

$$\mathcal{L}_T^{Repr} = \sum_{i=2}^8 \left\| \pi_f(\mathbf{P}^{1 \rightarrow i} \pi_f^{-1}(\mathbf{p}_1, 1/d^{-1})) - \mathbf{p}_i \right\|_1 \cdot (\sigma_{\max} - \sigma_i) \quad (16)$$

$$\mathcal{L}^{Repr} = \frac{1}{|\mathbf{T}|} \sum_{T \in \mathbf{T}} \mathcal{L}_T^{Repr} \quad (17)$$

Note that the uncertainties are not optimized during test-time refinement. Additionally, we apply smoothness term to encourage straight trajectories. Let n be the total number

of frames in the sequence:

$$\mathcal{L}^{Smooth} = \frac{1}{n-2} \sum_{i=1}^{n-2} \left\| \left(\mathbf{P}_f^{i \rightarrow i+1} \right)^{-1} \mathbf{P}_f^{i+1 \rightarrow i+2} - \mathbf{I}_4 \right\|_{1,1} \quad (18)$$

The final cost function is obtained by combining both terms, with $\lambda_{Smooth} = 0.1$:

$$\mathcal{L}_{BA} = \mathcal{L}^{Repr} + \lambda_{Smooth} \mathcal{L}^{Smooth} \quad (19)$$

We implement the entire BA process in PyTorch and use the Adam optimizer with a learning rate of $1e^{-4}$.

Sliding Window. Optimizing the entire sequence at once is both costly and can lead to instabilities. Therefore, we apply BA in a sliding window fashion. We define our window to be $w = 8$ frames wide and use an overlap of $o = 6$. That means we begin by optimizing the first 8 frames, and then shift the window by $w - o = 2$ to optimize frame 3 to 10. Note that we freeze the poses that have already been optimized and only adapt poses 9 and 10. For every sliding window, the optimization is performed for 400 steps. This is repeated until the end of the sequence is reached. In the end, we perform 5000 steps of global BA, where we consider all poses.

E. Limitations & Future Work

Reliance on pretrained model. AnyCam uses both pretrained depth and optical flow models during training and inference. While UniDepth and UniMatch show really strong performance, they can fail in rare cases. Depending on the severity of the failure, the accuracy of AnyCam can then get compromised. Typical failure cases include poor optical flow predictions when there are challenging lighting conditions, or inaccurate depth predictions when the input image does not have any scene context. Note that many errors in the input can still be dealt with due to our uncertainty formulation. Similarly, even though depth prediction are very consistent in scale as UniDepth is a metric depth model, the depths can have small flickering. This becomes visible when aggregating multiple depth maps over a longer time.

For future work, it would make sense to design the model to be reliant exclusively on images as input. Furthermore, we plan to add a system which adapts the scale and shift of the depth maps to be consistent among each other to allow for more accurate 4D reconstruction.

Drift over longer time. Our test-time refinement already greatly improves the drift problem. However, even then we only use tracks of length 8. To overcome drift on a global scale, our system would require a global scene representation or other techniques like keyframes. Many existing

SLAM and SfM systems [40, 43] provide inspiration for that.

Unnatural camera motion During training, AnyCams learns to translate images, flow, and depth of a sequence to a realistic camera motion. Our training data is very diverse, covering a wide range of realistic motions, and our experiments show that AnyCam can generalize very well. Still, due to its nature as a neural network, the model can fail when encountering very uncommon / unnatural camera motions.

To improve generalization even further, we plan to train the model on even more datasets. This can be achieved easily as our training pipeline is able to digest any kind of unlabelled videos.

F. Ethical Considerations

Our training data is partially made up of videos obtained from public sources like YouTube. These videos can contain identifying information like faces, number plates, etc.. To remove such information, faces have been blurred in many datasets, *e.g.* WalkingTours, before usage in our project. Additionally, since our model only predicts camera poses and uncertainty, the output does not allow to infer the identity of persons in the input data.

While we also aim to build a data mix that covers different geographical regions and domains, it is nevertheless possible that the model learns a bias. For example, driving data in the OpenDV dataset is mostly from the US, China, and Europe. Our model might struggle in driving environments that are very different from this training data.

Finally, despite showing strong performance, we cannot provide accuracy guarantees for the predictions of AnyCam. Therefore, it should not (yet) be used in safety-critical applications.