# NEURAL CAMERA MODELS

BY
IGOR VASILJEVIC

A thesis submitted
in partial fulfillment of the requirements for
the degree of

Doctor of Philosophy in Computer Science

at the

TOYOTA TECHNOLOGICAL INSTITUTE AT CHICAGO
Chicago, Illinois

June, 2022

Thesis Committee:
Professor Gregory Shakhnarovich (Chair)
Professor Matthew Walter
Professor Matthew Turk

# NEURAL CAMERA MODELS

by

## Igor Vasiljevic

## Abstract

Modern computer vision has moved beyond the domain of internet photo collections and into the physical world, guiding camera-equipped robots and autonomous cars through unstructured environments. To enable these embodied agents to interact with real-world objects, cameras are increasingly being used as *depth sensors*, reconstructing the environment for a variety of downstream reasoning tasks. Machine-learning-aided depth perception, or *depth estimation*, predicts for each pixel in an image the distance to a scene object. While impressive strides have been made in depth estimation, significant challenges remain: (1) ground truth depth labels are difficult and expensive to collect at scale, (2) camera information is typically assumed to be known, but is often unreliable and (3) restrictive camera assumptions are common, even though a great variety of camera types and lenses are used in practice. In this thesis, we focus on relaxing these assumptions, and describe contributions toward the ultimate goal of turning cameras into truly generic depth sensors.

In Chapter 3, we show how to extend self-supervised monocular depth estimation beyond pinhole cameras, jointly estimating depth, pose and calibration for a generic parametric model. Our method obtains both state-of-the-art depth and competitive calibration, comparing favorably to traditional target-based methods. In Chapter 4, we take an even more general approach, learning a per-pixel *ray surface* model that can represent any central camera. We introduce an architecture for estimating this general model without modification on radically different camera setups, and apply it to a variety of challenging datasets where standard parametric models fail, including dashboard cameras behind windshields and marine robots exploring underwater caves.

In Chapter 5, we relax another standard assumption: the monocular camera. We tackle the problem of self-supervised learning on *multi-camera rigs*, making no assumptions about the degree of overlap between any two cameras. We test our model on recent autonomous driving datasets that include low-overlap, multi-camera setups, demonstrating that we can learn 360° point clouds with no supervision. Our metrically-scaled estimates achieve state-of-the-art results by a large margin on a challenging autonomous driving benchmark. Finally, in Chapter 6, we move beyond single-frame inference into the realm of multi-frame depth estimation. We approach this problem through the lens of input-level *inductive biases*, replacing multi-view geometry constraints with geometric encodings and view augmentations that encourage, rather than enforce, multi-view consistency. Our geometric scene representation achieves state-of-the-art results on challenging video depth benchmarks.

# Acknowledgments

I feel very lucky to have had the opportunity to work with my advisor, Gregory Shakhnarovich. His deep understanding of *everything* machine learning and vision-related has helped guide my work starting with my Master's at UChicago, through my early PhD (when my interests fluctuated wildly) to our recent fruitful collaboration with the Toyota Research Institute.

I am also lucky to have had the opportunity to work with my colleagues at TRI. My collaboration with Vitor Guizilini, Rares Ambrus and Adrien Gaidon forms the bulk of my thesis: Chapters 3, 4, 5, and 6. This research would not have been possible without their mentorship (and especially Vitor's around-the-clock guidance and collaboration, especially in Chapters 5 and 6). The opportunity to explore interesting practical problems and datasets at TRI has been invaluable in guiding my research direction.

I would also like to thank Matthew Walter, whose robotics courses (and gadget-filled lab) have inspired my interest in and guided my understanding of embodied agents (I promise to eventually find a use for that drone, Matt). I want to thank Matthew Turk for his valuable comments during my thesis proposal defense, helping guide the content of this document.

I would like to thank Jiading Fang, my TTIC colleague and collaborator. It was great to find someone whose research interests aligned closely with mine, and our collaboration forms the core of two chapters: Chapters 3 and 6. I hope this fruitful collaboration continues into the future, and I am excited to see what kinds of robotics applications you build on top of our work.

I can't forget the faculty and staff at TTIC: Madhur Tulsiani, Avrim Blum (thank you for an interesting if *challenging* qual paper), Karen Livescu, David McAllester, Michael Maire (now at UChicago), Chrissy Coleman, Mary Marre, Erica Cocom, and especially Adam Bohlander (without whom everyone's experiments would have taken far longer to run in the early days).

I also want to thank my other colleagues and friends at TTIC and UChicago: Nick Kolkin and Steven Basart for exciting early collaborations back when deep learning was (relatively) new, Ruotian Luo for his ability to debug even the most esoteric pytorch issues, Reza Mostajabi for mentoring me during my Master's and inspiring me to eventually investigate depth estimation in the first place. I never got to coauthor a paper with Sudarshan Babu, but our discussions (more often, debates) about what is interesting and meaningful to work on have guided my work for years.

I would also like to thank my other TTIC and UChicago colleagues for fruitful discussions: Haochen Wang, Falcon Dai, Chip Schaff, Gustav Larsson, Andrea Daniele, Subham Toshniwal, Takuma Yoneda, and of course the inimitable Shubhendu Trivedi.

Finally, I would like to thank my family. Mama, Majo, Katarina, Bokice, Deko, thank you for all your support and encouragement over the years.

# Contents

# List of Figures

# List of Tables

# Chapter 1

# Introduction

> Perspective is nothing else than the seeing of an object behind a sheet of glass, smooth and quite transparent, on the surface of which all the things may be marked that are behind this glass; these things approach the point of the eye in pyramids, and these pyramids are cut by the said glass.
>
> Leonardo Da Vinci.

It was a sunny June morning in Long Beach, California, and I was leaning against the wall at the back end of a standing-room-only workshop at CVPR. The last invited speaker had finished a talk about recent developments in depth estimation, and a panel session about the "future of 3D vision" was underway. I looked at my watch—if I skipped the remainder of the morning program, did I have enough time to surf at Huntington Beach before the poster session? As I started to jostle my way out of the crowd, I heard a panelist say something that immediately caught my attention:

Why don't we just *learn* the fundamental matrix? Traditional vision is based on ideal assumptions that our actual measurements cannot hope to match. In the future, our depth networks will learn geometry *implicitly* rather than enforcing geometric

constraints. We will directly predict all quantities of interest directly from image data.

To me, this statement immediately jumped out as wrong, even *absurd*. Why would we throw away decades (actually, centuries) of development in geometric vision just to learn something we should already know to be true? If we know that a geometric constraint should hold (such as epipolar relationships for multi-view depth estimation), why would we not explicitly encode this as part of the network architecture? It seemed downright wasteful to ask a neural network to learn well-known geometric relationships from scratch.

Indeed, at the time, it seemed that hybrid systems had the best of both worlds: convolutional networks for tasks like stereo depth estimation were augmented with traditional architectural elements (e.g., cost volumes), achieving state-of-the-art results ( [10, 11]). These methods outperformed hand-crafted features for a variety of geometric vision tasks (depth, flow and pose estimation), but also handily beat pure end-to-end black boxes that the panelist seemed to be advocating. Classical vision methods still held sway in the setting where no labeled data was available, but at least on benchmark datasets, a fusion of traditional techniques and learned features emerged as the dominant paradigm.

One major problem limited the advancement of learning-based 3D vision—the paucity of ground-truth annotations. The first major change in the paradigm to address this issue was the emergence of *self-supervised* depth methods: now, an end-to-end depth (or flow) architecture could be learned with no labels at all. Self-supervised networks replaced depth label supervision with a view synthesis objective: given a sequence of images, depth and pose networks are trained to warp source to target views [12], the image reconstruction loss guiding depth and pose estimation. Though these depth estimators initially performed poorly compared to traditional and supervised methods, loss [13] and architecture [9] improvements soon led to superior performance. Today, self-supervised "monodepth" techniques form the core of state-of-the-art monocular and stereo estimation [9, 14].

Self-supervised depth estimation to some extent follows the trend outlined by the CVPR panelist: rather than solving from relative pose and depth using correspondences, a pose network predicts camera transformations *directly from images*, and a depth network directly regresses a depth value for each pixel. Still, these methods rely on *explicit* geometry, particularly on accurate and pre-calibrated

(and known) cameras, limiting their practical utility.

A recent revolution in network architectures for computer vision has brought the field much closer to the panelist's vision: the proliferation of Vision Transformers (ViT) [15] and a variety of subsequent generalist networks [16, 17]. Forsaking the inductive bias of convolutional layers, Transformers (and their variants) have been applied with minimal modification to text, image, and audio data, achieving state-of-art results by encoding data-specific inductive biases (such as patch positional encodings) at the *input* level. Transformer variants have achieved impressive results on geometric vision tasks such as optical flow [17] and stereo estimation [18], solely with input-level biases (i.e., patch and camera ray encodings, respectively), outperforming much more complex, bespoke architectures that bake the inductive bias (e.g., cost volumes) into the architecture itself. These networks fully realize the vision of an end-to-end network that is free of *architectural* biases.

In the final chapter of this thesis, we will describe a Transformer-based video depth estimation architecture that outperforms hybrid methods, matching the accuracy of techniques that use a wide variety of traditional architectures and losses (e.g., cost volumes and bundle adjustment) but also outperforming them on zero-shot generalization by a wide margin. In light of the transformer revolution in vision, the comments of the CVPR panelist seem quite prescient—this new paradigm promises to upend standard practices and frameworks for geometric vision tasks in the 2020s, much as the introduction of convolutional networks did in the 2010s.

## 1.1 Thesis Outline

This thesis will present a series of methods for generalized depth estimation.

First, in Chapter 2 we will review the geometry of single and multiple cameras, briefly describing some of the techniques used in geometric vision. We then review the depth estimation task, describing the standard datasets for this task, both indoors (with ground truth obtained using noisy but dense IR cameras) and outdoors (with ground truth collected with sparse but precise LiDAR scanners). We review a novel indoor-outdoor depth dataset we collected to train a general (supervised)

depth estimator, and describe some of the challenges we faced, eventually motivating our shift to self-supervised depth estimation. We describe the self-supervised framework used throughout the thesis, and review some common metrics used to benchmark depth models.

The standard self-supervised framework described in Chapter 2 relies on calibrated, perspective cameras. In Chapters 3 and 4, our goal is to generalize self-supervised depth beyond this limited setting, proposing two *self-supervised self-calibration* frameworks: one for parametric camera models (Chapter 3), and another for non-parametric camera models (Chapter 4). In Chapter 3 we start by replacing the pinhole model with a more general parametric Unified Camera Model (UCM), and jointly estimate the camera parameters along with depth and pose. We test not only the estimated depth (achieving state-of-the-art results on the challenging EuRoC drone dataset), but also evaluate the accuracy of the estimated UCM parameters. We demonstrate for the first time *sub-pixel re-projection error* for self-supervised self-calibration. Our method compares favorably to a state-of-the-art supervised (target-based) calibration tool [19].

In Chapter 4, we consider an even more general monocular self-calibration approach, replacing parametric cameras with a non-parametric per-pixel "ray surface" model (due to Grossberg and Nayar [20]). This "generalized" camera model separately calibrates the viewing ray for each pixel, and can model any central imaging geometry. It describes equally well perspective, fisheye, and even complex camera-mirror systems (catadioptric cameras). This generality comes at the cost of highly-overparameterized calibration (where parameters must be obtained *for each pixel*); thus, adoption of this model has been limited. A further limitation is that the projection operation is significantly more complex than standard parametric models (requiring iterative optimization to find the calibration ray closest to the 3D point of interest). We show that it is possible to *learn a per-pixel ray surface with no calibration targets* as part of a self-supervised learning framework. We introduce a curriculum learning approach and a soft-max approximation for the (non-differentiable) projection operation, slowly adapting the camera model as the network learns depth and pose estimators from data. We train our "neural ray surface" model without modification on perspective, fisheye, catadioptric, underwater, and dashboard camera images, achieving state-of-the-art results and enabling for the first time self-supervised depth estimation for several challenging imaging geometries.

4

The prior two chapters describe methods that expand the class of cameras that can be used for monocular depth estimation, but they are still limited to the single-camera setting. For many robotics tasks (and especially in autonomous driving), multi-camera datasets are ubiquitous. In Chapter 5, we extend self-supervised monocular depth estimation to the multi-camera setting, proposing a model that makes no assumptions about the number of cameras. We assume that one or more cameras are rigidly attached to a generic moving rig, and these cameras can have arbitrary (or no) overlap in field-of-view. Independently predicting single-frame depth for each camera on the rig and enforcing multi-view constraints, we learn a multi-view consistent depth estimator. For camera rigs that cover the entire field of view around the vehicle, we predict $360°$ degree point cloud estimates from image data. The depth predictions have metric scale, in contrast to standard self-supervised monodepth, which is only estimated up-to-scale; in prior work metric scale requires additional information (such as velocity [9]). The full-surround point clouds are obtained from depth predictions from individual cameras and the known extrinsics, and they are well-aligned due to our consistency constraints. We achieve state-of-the-art results by a large margin on challenging autonomous driving benchmarks.

In the final chapter, we remove a core assumption in all of the prior chapters (multi-view consistency) and move beyond single frame inference. We propose a video depth estimation architecture that relies on no geometric constraints, mapping an arbitrary number of input frames to a latent space from which we can decode any number of novel depth views. Our method's starting point is similar to the ray surface in Chapter 4, encoding viewing ray directions and positions at each pixel as an implicit representation (an input-level *inductive bias* [18]) for that pixel's geometry in space. We propose a random view augmentation procedure that *encourages* rather than enforces multi-view consistency. The method significantly outperforms state-of-the-art transformer stereo networks but also outperforms traditional bespoke video depth architectures that are significantly more complex (relying on cost volumes, epipolar constraints, and even bundle adjustment). Our method also achieves state-of-the-art zero-shot generalization results by a wide margin, demonstrating that inductive bias-based learning is a promising direction for robust geometric vision.

# Chapter 2

# Background

> There are no forms in nature. Nature is a vast, chaotic collection of shapes. You as an artist create configurations out of chaos. You make a formal statement where there was none to begin with. All art is a combination of an external event and an internal event... I make a photograph to give you the equivalent of what I felt. Equivalent is still the best word.
>
> Ansel Adams.

In this chapter we describe the fundamentals of geometric vision used throughout the thesis. First, we review cameras and camera models; this will form the background for discussing our self-calibration contributions in Chapters 3 and 4. Then, we introduce the task of (monocular) depth estimation and describe the two major settings in the literature: supervised and self-supervised learning.

## 2.1   A Brief History of Cameras

In the early $11^{th}$ century, the Arab mathematician Alhazen (Ibn al-Haytham) had a problem. He wanted to measure the shape of a solar eclipse, but he lacked a device to safely observe the phenomenon [21]. Alhazen, already an accomplished geometer, would turn to an optical device and conduct one of the first camera-aided experiments. His solution was a *camera obscura*, a dark room with a small hole on one side. He identified a number of conditions and limitations of this device: [21]:

- the object of interest must be sufficiently bright;

- the *camera obscura* must be sufficiently dark;

- the camera must be sufficiently wide;

- for a sharp image to be cast, the hole must be narrow enough;

- finally, the hole must be wide enough to avoid diffraction.

The principle was simple—light travels in a straight line, and if it passes through a hole that is sufficiently small, it will not scatter but form an image (albeit an inverted one) on the opposite wall or screen.

This early "pinhole camera" allowed Alhazen to observe a partial eclipse, and the device along with his accompanying derivations on the mathematics of perspective projection pioneered the geometry of pinhole cameras many centuries before photographic cameras were developed [21]. Alhazen would go on to set the foundations of modern optics, and the principles he developed for his pinhole camera were some of the earliest steps in geometric vision.

The photographic camera as we know it today would not be developed until the $19^{th}$ century. Camera construction quickly improved, and increasingly complex camera *lenses* were developed. Many of the problems that Alhazen outlined above were solved or at least ameliorated with newer technology. The lens, a piece of shaped glass used to focus light, allowed the imaging of objects that

7

were less bright for much shorter exposure times, allowing for the capture of images that were far superior to those of pinhole-based cameras. They also reduced diffraction and distortion issues but introduced their own lens- and lens-system specific distortions that would motivate sophisticated camera models and calibration techniques.

Though they came many centuries after Alhazen's early eclipse-imaging experiments, attempts to use these new devices to recover three-dimensional geometry long predate computer vision as a field. The earliest "geometric vision" experiments followed shortly after the invention of the camera. A number of cartographers noticed that images could dramatically speed up the creation of maps. One of these researchers, the French engineer Aimé Laussedat, initiated a research program he called "iconométrie", the very earliest attempts at camera-based structure estimation. Laussedat's approach was simple: he took images of landmarks in Paris and then reconstructed them by hand using the geometry of perspective projection [22], more than 150 years before the landmark "Phototourism" [23] paper enabled large-scale automated reconstruction of landmarks from tourist images.

Édouard-Gaston Daville, a Canadian surveyor, further developed the theory of mapping from multiple images, and in the process of improving his procedures introduced some of the first camera calibration procedures for reconstruction [24]. He (and his many followers) used manual *stereoplotters* to match features from stereo pairs, and his sophisticated topographic reconstruction procedure would remain in use for decades [24].

As aerial surveys increased in frequency and importance before and during the First World War and the importance of this nascent field grew, calibration techniques were developed to carefully estimate the perspective camera parameters [25]. By the Second World War, reconnaissance and mapping applications required the use of increasingly sophisticated calibration techniques. The war spurred the development of the first sophisticated distortion models [25], and attempts were made to measure in the lab any possible imperfections of the imaging system and changes caused by variables introduced by the harsh aerial imaging environment (e.g. very low temperatures). As lab-based re-calibration was often not possible for many cameras used by the military, *self-calibration*

techniques emerged around this time. The stars were a natural target for self-calibration—estimating the "extrinsics" of the camera with respect to the stars had a long history in the practice of maritime celestial navigation (using star surveys), and there was a large enough number of these targets for relatively robust least-squares solutions [25].

In 1956, Duane Brown published a paper titled "The Simultaneous Determination of the Orientation and Lens Distortion of a Photogrammetric Camera" and initiated the study of "bundle adjustment", the joint extrinsic and intrinsic calibration and 3D structure estimation from a system of cameras (or images from a single camera). His work on bundle adjustment remains in use to this day as part of large scale structure-from-motion toolkits, and his distortion model is still one of the most commonly-used camera models. We will briefly review it in Section 2.2.2).

Today, there are many imaging devices available to us, from inexpensive aftermarket dashboard cameras, to waterproof casing underwater rigs. The proliferation of cameras, especially low-cost cameras, necessitates the use of increasingly more complicated camera models. In the later chapters of this thesis, we will discuss a variety of methods that aim to generalize modern geometric methods to the great variety of cameras available today.

## 2.2   Camera Geometry

A camera model is a representation of the relationship between world points and their images. In this section, we will first briefly discuss the most ubiquitous camera model in robotics and vision, the pinhole camera model, and the distortion parameters that are typically added to this model. Then, we will review how a camera calibration is obtained and evaluated, introduce a few relevant multi-camera relationships, and discuss viewing ray geometry.

### 2.2.1   The Pinhole Model

The aperture of Alhazen's *camera obscura* was simply a small hole in a wall—the researcher enumerated the many limitations of this approach, and in the centuries since, light-focusing *lenses* placed in front of the aperture have proliferated. These lenses transform 3D points into 2D images onto a photo-sensitive *image plane*. Modern optics has led to increasingly complex lens systems, even in consumer cameras–increasingly more complex compound lenses are common in smartphone cameras (see Figure 2.1 for an example). These designs attempt to fix many of the problems identified by Alhazen, whether it be diffraction or geometric distortion, subject to constraints (e.g. for smartphone, thickness is at a premium).

Despite the complexity of modern lens systems, the pinhole approximation remains popular in computer vision. In this section, we will briefly review the pinhole camera model, and then discuss some common distortion models that attempt to bring the estimated camera closer to the physical device.



Figure 2.1: Modern smartphone cameras use a complex series of lenses to focus light on their tiny sensors.

The pinhole model is illustrated in Figure 2.2. For a world point on an imaged object $\mathbf{P} =$
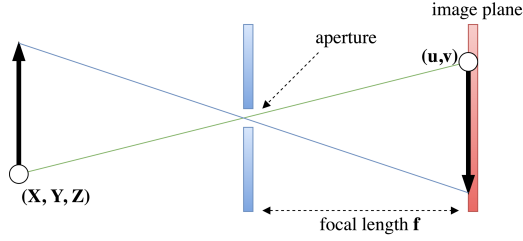
Figure 2.2: Light reflected from scene points travels through the pinhole a distance defined by the focal length to be collected on a photosensitive image plane.

$(X, Y, Z)$, reflected light travels through space and enters the camera through an aperture. We assume that this point is in camera coordinates. The size of this aperture is assumed to be infinitesimal and fixed. The light then travels a distance $f$ (the focal length) from the aperture to the image plane, intersecting it at a point $\mathbf{p} = (u, v)$. Using similar triangles we can find that $(X, Y, Z)$ projects to $(u, v)$ as:

$$u = fX/Z + c_x \quad v = fY/Z + c_y \tag{2.2.1}$$

where $(c_x, c_y)$ is the principal point of the camera, the pixel to which the center of projection is projected. The camera parameters: $(f, c_x, c_y)$ can be collected into a *intrinsics matrix*:

$$\mathbf{K} = \begin{bmatrix} f & 0 & c_x \\ 0 & f & c_y \\ 0 & 0 & 1 \end{bmatrix} \tag{2.2.2}$$

If we use *homogeneous coordinates*, which are widely used in projective geometry, the projection and unprojection operations can be written as matrix-vector products. For a pixel $\mathbf{p} = (u, v)$, we instead append a coordinate and write $\mathbf{p} = (u, v, 1)$ in the homogeneous representation, allowing projection to be a matrix-vector product. For the rest of this thesis, pixels $\mathbf{p}$ are always in homogeneous coordinates (unless indicated otherwise).

Given a scene depth $Z$, we can write *projection*, the mapping of 3D points to 2D image coordinates, as $\pi(\mathbf{P}) = \frac{1}{Z}\mathbf{KP}$. *Un-projection*, the mapping of image coordinates to 3D points, is written as

$\phi(\mathbf{p}, Z) = Z\, \mathbf{K}^{-1}\mathbf{p}$. In matrix form, for a focal length $f$ we have un-projection:

$$\phi(\mathbf{p}, Z) = Z\, \mathbf{K}^{-1}\mathbf{p} = Z \begin{bmatrix} f & 0 & c_x \\ 0 & f & c_y \\ 0 & 0 & 1 \end{bmatrix}^{-1} \begin{bmatrix} u & v & 1 \end{bmatrix}^T \tag{2.2.3}$$

and projection:

$$\pi(\mathbf{P}) = \frac{1}{Z}\mathbf{K}\mathbf{P} = \frac{1}{Z} \begin{bmatrix} f & 0 & c_x \\ 0 & f & c_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} X & Y & Z \end{bmatrix}^T, \tag{2.2.4}$$

This model (or a common variation that has a separate focal length in x and y directions: $f_x, f_y$) is often a "good enough" approximation for high quality machine vision perspective cameras, but the huge diversity of cameras and lenses today means that there are many cameras for which this is a poor approximation. More accurate variants for many lenses introduce new distortion parameters in addition to the focal length and principal point, reducing the approximation error.

### 2.2.2 Distortion Models

Distortion models allow the introduction of calibration parameters that allow the model to be a closer fit to many physical lenses. Two major types of distortion that are modeled by the Brown model [26] (introduced in Section 2.1) are *radial* distortion and *tangential* distortion. Radial distortion occurs when a lens is not rectilinear, meaning that lines are not mapped to lines but curves. In the Brown model, this distortion is typically modeled by three parameters: $(k_1, k_2, k_3)$. Tangential distortion occurs when the lens is not parallel to the image plane, and can also be modeled by two parameters: $(p_1, p_2)$.

Consider projecting a point $\mathbf{P}$ onto the image plane of a camera with lens distortions described by this model. The projection according to an un-distorted pinhole model will project the point to a pixel $\mathbf{p} = (u, v)$. Denote the distorted pixel by $\mathbf{p}_{distorted} = (u_d, v_d)$. Then, the distortion

according to the radial parameters is:

$$u_d = u(1 + k_1 r^2 + k_2 r^4 + k_3 r^6) u_d = v(1 + k_1 r^2 + k_2 r^4 + k_3 r^6) \qquad (2.2.5)$$

and tangential distortion can be written as follows:

$$u_d = u + (2p_1 uv + p_2(r^2 + 2u^2)) v_d = v + (2p_2 uv + p_1(r^2 + 2v^2)) \qquad (2.2.6)$$

Given that the radial distortion for many lenses is (approximately) radially symmetric, the parameter $r(u, v)$ is given by: $r(u, v) = \sqrt{(u_d - u)^2 + (v_d - v)^2}$.

This model is popular due to its relative simplicity, and continues to be popular in the literature, being one of the default distortion models for OpenCV [1]. In the literature on depth estimation, it has recently been used by Gordon et al. [4] as a general model for a large dataset of YouTube videos coming from many (unknown) cameras. In Chapters 3 and 4, we describe alternative distortion models that describe a much wider variety of cameras.

### 2.2.3 Calibration

The estimation of camera parameters, or *camera calibration*, is an essential first step in many robotics and vision applications. Traditionally, it is performed by capturing images of known calibration targets (such as checkerboards), which consist of patterns that allow for highly accurate estimation of correspondences between world points $\mathbf{P}_i$ and pixels $\mathbf{p}_i$. Corner detectors are able to estimate checkerboard corners with sub-pixel accuracy and automatically establish correspondence between multiple images of the target. Given these correspondences, the goal of calibration is to estimate the projection matrix $\mathbf{M} = \mathbf{K}[\mathbf{R}|\mathbf{t}]$, which includes the intrinsic parameters as well as the transformation from the camera to the world points $(\mathbf{R}, \mathbf{t})$. Using the projection equation:

$$\lambda \mathbf{p}_i = \mathbf{M} \mathbf{P}_i \qquad (2.2.7)$$

---

[1] It is used as part of the default camera model for the OpenCV calibration tutorial.

a number of algorithms have been designed for estimating $\mathbf{M}$ from $(\mathbf{p}_i, \mathbf{P}_i)$ correspondences. A popular algorithm due to Zhang [27] allows for the estimation of $\mathbf{M}$ from six correspondences, and subsequent extraction of the calibration matrix $\mathbf{K}$ and transformation $(\mathbf{R}, \mathbf{t})$. The calibration can be further refined by minimizing re-projection error (Equation 2.2.8).

This calibration procedure requires *ground-truth* correspondences, which is a strong assumption in many settings (e.g., re-calibrating a vehicle camera in the field). In this thesis, we will describe methods that require no calibration targets, but compete in calibration accuracy with methods that do (Chapter 3).

**Evaluation**   To evaluate the quality of a camera calibration, we need to measure how well the model un-projects and projects known pixel and point correspondences. Consider a projection function $\pi^{\mathbf{K}}$ for a camera with intrinsics $\mathbf{K}$. Then, points $\mathbf{P}_i$ project to pixels $\pi^{\mathbf{K}}(\mathbf{P}_i, \mathbf{R}, \mathbf{t}) = \mathbf{p}_i$. These correspondences can be AprilTag [28] positions or checkerboard corners [1]. For a set of observations $i = 1, \dots, n$, the mean re-projection error is:

$$E(\mathbf{R}, \mathbf{t}) = \frac{1}{N} \sum_{i=1}^{n} \|\pi^{\mathbf{K}}(\mathbf{P}_i, \mathbf{R}, \mathbf{t}) - \mathbf{p}_i\|^2 \tag{2.2.8}$$

This expression incorporates error in a given calibration procedure (estimation of parameters) and the camera model itself (how well does it describe the relationship between pixels and points for a given camera) and is measured in pixels.

### 2.2.4   Multi-Camera Geometry

The geometry of multiple cameras (or a single camera at multiple points in time) is at the core of geometric vision. For the purposes of this thesis, we will briefly review elements of two-view geometry that have found use in modern depth and flow estimation architectures.

Consider a set of pixel correspondences coming from two images (in homogeneous coordinates), $\mathbf{p}_0$ and $\mathbf{p}_1$. Given a common intrinsics matrix $\mathbf{K}$ and relative pose $(\mathbf{R}, \mathbf{t})$, the *epipolar constraint*

fully describes the geometry of the problem:

$$\mathbf{p}_1^T \mathbf{K}^{-T} \mathbf{R}[\mathbf{t}]_\times \mathbf{K}^{-1} \mathbf{p}_0 = \mathbf{p}_1^T \mathbf{F} \mathbf{p}_0 = \mathbf{0} \tag{2.2.9}$$

where $\mathbf{F} = \mathbf{K}^{-T} \mathbf{R}[\mathbf{t}]_\times \mathbf{K}^{-1}$ is called the *fundamental matrix*. It can be estimated from correspondences using a wide variety of algorithms, among the most popular being normalized 8-point [29].

Starting with PoseNet [30], pose estimation has also been treated as a *regression* problem, taking image-pair input and (ground truth) pose output and training a relative pose estimator directly from pixels. These supervised approaches have serious limitations [31], failing to generalize to unseen camera motions.

Direct pose regression has found an important application in self-supervised depth estimation [12], where PoseNet variants are used for generating pose predictions without any direct supervision (more details to come in Section 2.3.4).

More recently, hybrid depth-and-pose architectures have been proposed that replace PoseNets, falling back to the robustness of traditional solvers [32] and showing significant improvements in generalization for pose estimation. Optical flow networks can be used to propose two-frame correspondences which are input for an 8-point solver, or they can be used as part of an epipolar constraint [33] to encourage a model's predictions to align more closely to two-view epipolar geometry.

Towards *enforcing* two-view geometry at the loss level, a popular normalized variant is the symmetric Sampson error [34]:

$$E_{\text{sampson}} = \frac{(\mathbf{p}_1^T \mathbf{F} \mathbf{p}_0)^2}{(\mathbf{F}\mathbf{p}_0)_1^2 + (\mathbf{F}\mathbf{p}_0)_2^2 + (\mathbf{F}^T\mathbf{p}_1)_1^2 + (\mathbf{F}^T\mathbf{p}_1)_2^2} \tag{2.2.10}$$

which has found use as a loss for two-view flow estimation [35]. These loss-level constraints allow the neural network's parameters to adjust to conform to epipolar geometry, but in Chapter 6 we will describe how networks that enforce these constraints at the loss level suffer from poor generalization.

Another two-frame geometric loss, the *disparity equation*, forms the core of the self-supervised depth estimation. This relationship between pixels in two views can be used to *warp* one image to the other.

Given two views ($I_0$ and $I_1$) of the same scene with relative pose $(\mathbf{R}, \mathbf{t})$ and intrinsics $\mathbf{K}$, our goal is to warp the source ($I_0$) to the target ($I_1$) frame. If we have access to scene depth $Z_0$ at a pixel $\mathbf{p}_0$, we can un-project this pixel to a 3D point, transform it to the coordinates of $I_1$, and project it into the new image.

In this way, we can compute the coordinates of all (visible) pixels from the source to the target view:

$$\mathbf{p}_1 \sim Z_0 \mathbf{K} \mathbf{R} [\mathbf{t}]_\times \mathbf{K}^{-1} \mathbf{p}_0 \qquad (2.2.11)$$

With an estimate of $\hat{I}_1$ generated in this way, we can compare it to the (held-out) $I_1$. We will describe this loss in more detail in Section 2.3.4.

## 2.2.5 Viewing Ray Geometry

As an alternative to the pinhole camera model described above, a camera $j$ can instead be interpreted as a *bundle* of viewing rays, and each pixel $i$ can be mapped to a 3D vector in space parameterized by a ray origin $\mathbf{o}_j$ and ray direction $\mathbf{r}_{ij}$. The viewing ray represents the line in space along which the pixel samples the world. This camera representation has become popular in recent years due to its use in volume rendering, especially in neural radiance fields [36], but it has a rich history in the literature on generalized cameras [20, 37, 38].

Projection and un-projection operations can be represented in this ray view. Consider a viewing ray along a pixel in homogeneous coordinates (overloading notation a bit) $\mathbf{p}_{ij} = (u_{ij}, v_{ij}, 1)$. Denote its origin by $\mathbf{o}_j$ and its direction by $\mathbf{r}_{ij}$. For a perspective camera with a known calibration $\mathbf{K}_j$, and transformation $(\mathbf{R}_j, \mathbf{t}_j)$ with respect to the origin of the world coordinate frame, we can

find the origin:

$$\mathbf{o}_j = -\mathbf{R}_j \mathbf{t}_j \tag{2.2.12}$$

and ray direction:

$$\mathbf{r}_{ij} = \left(\mathbf{K}_j \mathbf{R}_j\right)^{-1} \mathbf{p}_{ij} \tag{2.2.13}$$

The ray direction $\mathbf{r}_{ij}$ may come from the calibration $\mathbf{K}_j$, but it may also incorporate any number of distortions; this type of model is often called the generalized or general camera model, and its earliest development is owed to the pioneering work of Grossberg and Nayar [20]. In Chapter 4, we will describe our self-supervised architecture for learning $\mathbf{r}_{ij}$.

For any pixel $\mathbf{p}_{ij}$, given a scene depth $Z$ we can un-project this pixel to a 3D point by:

$$\phi(\mathbf{p}_{ij}, Z_{ij}) = \mathbf{P}_{ij} = \mathbf{o}_j + Z_{ij}\left(\mathbf{K}_j \mathbf{R}_j\right)^{-1} \mathbf{p}_{ij} \tag{2.2.14}$$

Going in the other direction (projection) is more complex for generalized cameras. Consider a 3D point $\mathbf{P}_{ij}$ we wish to project onto the image plane. Denote the direction from the point to the camera center, i.e. $\mathbf{P}_{ij} - \mathbf{o}_j = \mathbf{q}_{ij}$. Assuming a central camera, we need to solve the following optimization problem:

$$\pi(\mathbf{P}_{ij}) = \hat{\mathbf{p}}_{ij} = \arg\max_i \mathbf{r}_{ij} \cdot \mathbf{q}_{ij} \tag{2.2.15}$$

That is, we need to find the ray on the ray surface closest to the direction to the 3D point from the camera center. This model is significantly more complex than a standard perspective projection, and Levenberg-Marquadt iteration is often used [39]. A recent neural radiance field variant [40] foregoes the need to project points through the use of volume rendering, learning the generalized camera using only the (closed form) un-projection operation.

In Chapter 4 we will describe a differentiable approximation which allows us to both project and un-project with a generalized camera, learning a differentiable approximation to Equation 2.2.5 and allowing us to learn the parameters of a wide variety of cameras.

## 2.3 Depth Estimation

Humans and most other animals rely on the ability to perceive object depth to navigate our three-dimensional world. Psychologists have amassed a collection of depth cues that humans use to turn visible objects into some notion of relative or absolute distance. These cues are typically divided into binocular, or stereo cues from seeing the object from two different perspectives (eyes), and monocular cues that only require one eye or image. The latter include occlusion, shading, perspective and texture gradients, all of which provide information about scene depth. Humans use both stereo and monocular hues to perceive depth, but traditionally geometric vision has investigated the former rather than the latter [41].

Like humans, robots need some way of perceiving the world and mapping vision-based (2D) predictions into 3D. In computer vision, a common intermediate representation of 3D shape is the *depth map*, a per-pixel representation of distance from the camera to the surface imaged at that pixel. In this section, we will describe depth estimation from the computer vision perspective. We will first describe the popular benchmark datasets that have driven much of the progress in the field over the past decade. Then, we will discuss some of the shortcomings of these datasets, and the challenges we encountered collecting our own diverse dataset attempting to address some of these limitations. Our supervised depth estimation experiments on this dataset (and challenge of obtaining accurate predictions even with a state-of-the-art supervised model) motivated our transition to self-supervised learning, which enables the use of huge un-labeled (or sparsely labeled) video datasets.

### 2.3.1 Benchmark Datasets

For any machine learning task, the most important starting point is *data*. The availability of large, diverse, and manually labeled datasets, such as ImageNet [42], Places [43], and COCO [44] has enabled the dramatic success of deep learning-based methods on recognition tasks. Annotation tools (and startups offering annotation services) have proliferated over the past few years, and as

annotation standards improved, a large number of high-quality labeled classification, detection and segmentation datasets have been collected.

In contrast, depth labels are challenging to collect. Humans can annotate relative depth (between two random points in an image, e.g., a pixel on a tree is closer to the camera than a pixel on a building behind the tree) but struggle at measuring metric depth. Work on creating annotated relative depth data [45] still relies on *RGB-D* ground truth as part of the training pipeline to train an accurate depth estimator. These RGB-D labels require specialized and often expensive 3D range sensors that have a variety of limitations (such as limited performance in direct sunlight). These include sparse but accurate LiDAR scanners typically used to collect outdoor driving datasets such as KITTI [46] and DDAD [9]) and noisy but dense IR scanners used to collect indoor datasets such as NYUv2 [47] and Scannet [48]. Both of these types of sensors have their limitations (e.g., IR cameras cannot be used outdoors and have limited range) and thus datasets are typically split between outdoor (streets) and indoor (apartment and office) datasets. These limitations seriously limit the diversity of possible training datasets.

One possible solution is the creation of synthetic data, and a large number of synthetic datasets have been proposed (e.g., Synthia [49], Virtual KITTI [50]) and SUN-CG [51], all with varying degrees of photorealism. With a simulator, perfect per-pixel depth maps can be generated, but this solution introduces a domain gap between synthetic training data and real-world test data (also known as the "sim2real" [52] gap, a very active area of research).

Next we will briefly describe the depth datasets used throughout this thesis, and then describe a dataset we collected to try and rectify the limitations of prior data.

- **KITTI** [53]. The KITTI dataset is the standard benchmark for depth and visual odometry evaluation. It was captured in and around the German city of Karlsruhe, and despite its relatively low diversity remains among the most popular depth estimation datasets. The training protocol for this dataset and the splits were introduced in Eigen et al. [54]. With the filtering steps described by Zhou et al. [12] to remove static frames, there are $39,810$ images for training, $4,424$ for validation and $697$ for evaluation.

- **EuRoC** [3]. The dataset consists of a set of indoor stereo image sequences captured by a micro-aerial vehicle. The dataset contains accurate motion-capture pose and ground truth depth captured with a Leica MS50 laser scanner. The scenes consist of small calibration rooms and a university mechanical room. The dataset consists of a set of indoor MAV sequences with general six-DoF motion. Consistent with recent work [4], in Chapters 3 and 4 we train networks on this dataset using center-cropping and down-sample the images to a $384 \times 256$ resolution.

- **OmniCam** [55]. OmniCam is a driving sequence (a single scene with $12,607$ frames) taken with an omnidirectional catadioptric camera, providing ground truth odometry. The content of the dataset is similar to KITTI: residential streets in the German city of Karlsruhe.

- **Multi-FOV** [56]. Multi-FOV is a small driving dataset recorded in a simulated environment, providing ground truth depth in a single synthetic scene for three different cameras —pinhole, fisheye, and catadioptric. The rendering is non-photorealistic, but the dataset is unique for having ground-truth for several high-distortion camera types (and in Chapter 4 we use it as a quantitative test of our non-parametric model on catadioptric data).

- **ScanNet** [48]. ScanNet is an RGB-D video dataset containing 2.5 million views from around $1500$ scenes. The data consists of real-world videos captured in apartments and offices, and is the indoor counterpart of KITTI, having found much use especially in multi-frame depth estimation papers. There are typically two different training (and evaluation) settings for this dataset: *stereo* and *video* depth estimation. For the stereo experiments, the setting is described by Kusupati et al. [57]: downsample scenes temporally by a factor of 20, and use a custom split to create stereo pairs, resulting in $94212$ training and $7517$ test samples. For the video experiments, the evaluation protocol was described by Teed et al. [58], a total of $1405$ scenes is used for training.

- **7-Scenes** [59]. This dataset, collected through KinectFusion [60], consists of $640 \times 480$ images in seven indoor office settings (similar in character to the Scannet office scenes), with a variable number of scenes in each setting. There are $500$–$1000$ images in each scene. This dataset is often used to test the out-of-domain generalization of models trained on Scannet.

Figure 2.3: Samples from DIODE, including (top) RGB images, (middle) depth maps, and (bottom) surface normal maps. The depth and surface normal maps are false colored according to ![legend] and ![sphere], respectively. Note that invalid depth values and normals are rendered as black and that the maximum depth value is image specific.

- **DDAD** [9]. DDAD is a modern, diverse driving dataset captured in multiple cities with high-resolution cameras and LiDAR ground truth. It contains a set of six cameras, and has a total of 12650 training samples, from which we consider all six cameras for a total of $75,900$ images. The validation set contains 3950 samples (23700 images) and ground-truth depth maps, used only for evaluation. Following the standard evaluation procedure [9], in our experiments input images were downsampled to a $640 \times 384$ resolution, and for evaluation we considered distances up to 200m without any cropping.

- **nuScenes** [61]. The nuScenes dataset is another driving dataset containing scenes from multiple cities. It is a popular benchmark for 2D and 3D object detection, as well as semantic and instance segmentation. It contains images from a synchronized six-camera array, comprised of 1000 scenes with a total of $1.4$ million images. It is a challenging dataset for depth estimation because of the relatively low resolution of the images, very small overlap between the cameras, high diversity of weather conditions and time of day, and unstructured environments. In our experiments, the raw images are $1600 \times 900$, which are downsampled to $768 \times 448$, and evaluated at distances up to 80m without any cropping.

## 2.3.2  DIODE: Dense Indoor/Outdoor Depth Dataset

The benchmark datasets we reviewed in Section 2.3.1 involve either indoor or outdoor scenes, but not both. Furthermore, they tend to lack diversity, particularly for outdoor datasets, which are usually collected with autonomous driving in mind and thus consist of street scenes. Indoor RGB-D datasets are usually collected using structured light cameras, which provide dense, but noisy, depth maps up to approximately 10 m, limiting their application to small indoor environments (e.g., home and office environments). Outdoor datasets typically have a narrow scope, often limited to driving scenarios, and are generally acquired with customized rigs consisting of monocular cameras and LiDAR scanners. Typical LiDAR scanners have a high sample rate, but relatively low spatial resolution. Consequently, the characteristics of available indoor and outdoor depth maps are quite different, and models trained on one domain typically generalize poorly to the other [62].

In this section we introduce the DIODE (Dense Indoor/Outdoor DEpth) dataset in an effort to address the aforementioned limitations of existing RGB-D datasets. DIODE is a large-scale dataset of diverse indoor and outdoor scenes collected using a survey-grade laser scanner (FARO Focus S350 [63]). Figure 2.3 presents a few representative examples from DIODE, illustrating the diversity of the scenes and the quality of the 3D measurements. This quality allows us to produce not only depth maps of unprecedented density and resolution, but also to provide surface normals for indoor and outdoor scenes with a level of accuracy and coverage not possible with existing datasets. The most important feature of DIODE is that it is the first dataset that covers both indoor and outdoor scenes in the same sensing and imaging setup.

We designed and acquired the DIODE dataset with three primary goals in mind. First, the dataset should include a diverse set of indoor outdoor scenes. Second, it should provide dense depth maps, with accurate short-, mid-, and long-range depth measurements for a large fraction of image pixels. Third, the depth measurements should be highly accurate.

After describing the dataset collection process, we will review some of the applications the dataset has found in the computer vision literature. Though some works have used DIODE as a depth estimation benchmark, it has also found a few interesting and surprising applications. Then, we

Figure 2.4: Illustrations of image/depth map cropping and rectification process. Left: a panorama (approx. 8K×20K pixels) produced by the scanner; every pixel with a valid return is associated with a depth value. Four colored frames indicate a subset of FOV crops projected on the panorama. Right: the corresponding frames after reprojection (to correct artifacts like those on the soffit in the purple frame), downsampling to 768×1024 pixels, and rectification. The depth map associated with the panorama is cropped and rectified in the same way, yielding aligned RGB-D pairs.

introduce depth estimation and point out some of the challenges we encountered when trying to both *train* and evaluate an (at the time) state-of-the-art depth estimator on DIODE. Our results on those experiments (showing a difficulty of generalizing) motivate us to pivot to self-supervised learning as a means of training depth estimators, transitioning from static images to video data as a means of training depth estimators without laborious (and expensive) large-scale ground truth depth collection.

**Data acquisition**     The aforementioned dataset qualities (dense, accurate, equally good indoors and outdoors) preclude measuring depth using structured light cameras, and instead requires using LiDAR. We collected our dataset using a FARO Focus S350 scanner [63]. The FARO is an actuated survey-grade phase-shift laser scanner that provides highly accurate depth measurements over a large range (0.6 m to 350 m) with error as low as 1 mm, with high angular resolution (0.009°). The FARO includes a color camera mounted coaxially with the depth laser, and produces a high-resolution panorama that is automatically aligned with the FARO's depth returns. These attributes give the FARO a variety of advantages over the more frequently used Velodyne LiDAR with a separate RGB camera, or Kinect depth cameras:

- the scanner is equally well suited for indoor and outdoor scanning;

- the point clouds are orders of magnitude more dense; and
- there is virtually no baseline between the depth sensor and the RGB camera.

**Scanning parameters** The FARO allows for the customization of various parameters that govern the scanning process. These include the resolution of the resulting depth scan (i.e., the number of points), the color resolution of the RGB panorama (i.e., standard or high definition), and the quality of the scan (i.e., the integration time of each range measurement). We chose the following scanning settings:

- $1\times$ quality: single scanning pass for every azimuth;
- $360°$ degree horizontal FOV, $150°$ vertical FOV;
- $^1/_2$ resolution: $\approx 170$M points; and
- $3 \times$ HDR: low exposure, regular, high exposure bracketing for RGB.

These settings result in a scan time of approximately 11 minutes. The intermediate output of a scan is a $20700 \times 8534$ (approximately) RGB panorama and a corresponding point cloud, with each 3D point associated with a pixel in the panorama (and thus endowed with color). As with other LiDAR sensors, highly specular objects as well as those that are farther than $350$ m (including the sky) do not have an associated depth measurement. Another limitation of the scanner for RGB-D data collection is that the LiDAR "sees" through glass or in darkness, resulting in detailed depth maps for image regions that lack appearance information.

**Scanning Locations** We chose scan locations to ensure diversity in the dataset as well a similar number of indoor and outdoor scenes. The scenes include small student offices, large residential buildings, hiking trails, meeting halls, parks, city streets, and parking lots, among others. The scenes were drawn from three different cities. Given the relatively long time required for each scan (approximately 11 min) and the nature of the scanning process, we acquired scans when we could avoid excessive motion and dynamic changes in the scene. However, occasional movement through the scenes is impossible to avoid entirely.

24

The resulting scans exhibit diversity not just between the scenes themselves, but also in the scene composition. Some outdoor scans include a large number of nearby objects (compared to KITTI, where the majority of street scans have few objects near the car), while some indoor scenes include distant objects (e.g., as in the case of large meeting halls and office buildings with large atria), in contrast to scenes in other indoor datasets collected with comparatively short-range sensors.
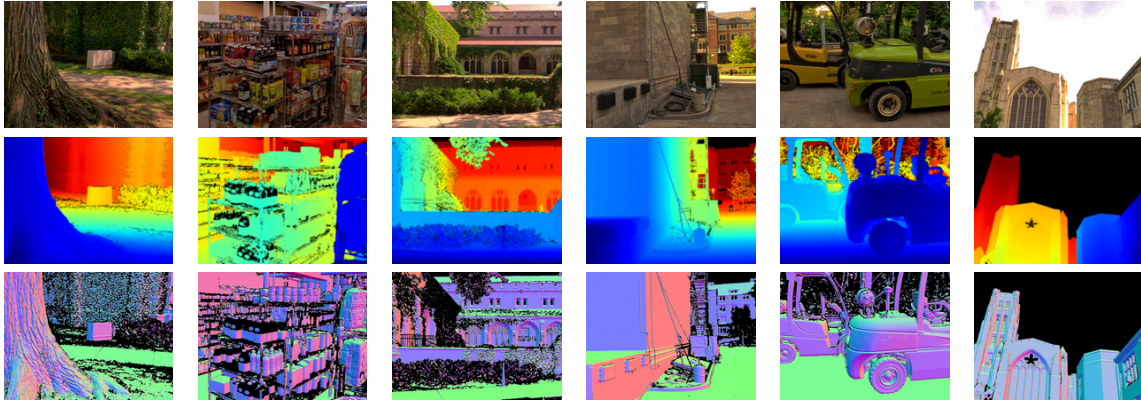


Figure 2.5: More samples from DIODE, including (top) RGB images, (middle) depth maps, and (bottom) surface normal maps. The depth and surface normal maps are false colored according to the same legends in Figure 1.

**Image Extraction**    We process the scans to produce a set of rectified RGB images (henceforth referred to as "crops") at a resolution of $768 \times 1024$. The crops correspond to a grid of viewing directions at four elevation angles ($-20°$, $-10°$, $0°$, $10°$, $20°$, and $30°$), and at regular $10°$ azimuth intervals, yielding 216 viewing directions. We rectify each crop, corresponding to a $45°$(vertical) $\times$ $60°$(horizontal) FOV.

Curved sections of the panorama that correspond to each viewing frustum must be undistorted to form each rectified crop, i.e., a rectangular image with the correct perspective. To accomplish this, we associate each pixel in the rectified crop with a viewing ray in the canonical coordinate frame of the scanner. We use this information to map from panorama pixels and the 3D point cloud to crop pixels.

For each pixel $\mathbf{p}_{ij}$ in the desired $768 \times 1024$ crop, let the ray passing through the pixel be $\mathbf{r}_{ij}$. We assign the RGB value of $\mathbf{p}_{ij}$ to the average of the RGB values of the nearest five pixels in terms

of the angular distance between their rays and $\mathbf{r}_{ij}$. We employ a similar procedure to generate a rectified depth map. For each ray $\mathbf{r}_{ij}$, we find the set of 3D points $\mathbf{X}_{ij}$ whose rays are nearest to $\mathbf{r}_{ij}$ in angular distance. We discard points for which the angular distance is greater than $0.5°$. We then set the depth of pixel $\mathbf{p}_{ij}$ to the robust mean of the depth of points in $\mathbf{X}_{ij}$, using the median $80\%$ of depth values. In the event that the set $\mathbf{X}_{ij}$ is empty, we record $\mathbf{p}_{ij}$ as having no return (coded as depth 0).

To compute normals, for each spatial index $(i, j)$, we then take the set of 3D points $\hat{\mathbf{X}}_{ij}$ indexed by the $11 \times 11$ grid centered at $(i, j)$. Using RANSAC [64], we find a plane that passes through the median of the set $\hat{\mathbf{X}}_{ij}$, and for which at least $40\%$ of the points in $\hat{\mathbf{X}}_{ij}$ have a residual less than an adaptive threshold, equal to approximately $0.2\,\text{mm}$ for points at $1\,\text{m}$ depth and $6\text{cm}$ at $300\,\text{m}$. We define the normal at position $(i, j)$ to be the vector normal to this plane, facing towards the origin (we set it to o if no such plane is found). Finally, we rotate these normals for each crop according to the camera vector, and rectify the normals using the same procedure used for the depth map. Note that the normals resulting from this procedure are based on the point cloud which is about ten times denser than the D channel in the RGB-D. Thus they are much more accurate and stable than if one were to derive them from the depth map.

**Crop selection**    The scanner acquires the full 3D point cloud before capturing RGB images. This, together with the relatively long scan duration can result in inconsistencies between certain RGB image regions and the corresponding depth for dynamic elements of the scene (e.g., when a car is stationary during the 3D acquisition, but moves before the RGB images are acquired). Additionally, some crops might have almost no returns (e.g., an all-sky crop for an outdoor scan). We manually curated the dataset to remove such crops, as well as those dominated by flat, featureless regions (e.g., a bare wall surface close to the scanner).

**Masking**    Though the depth returns are highly accurate and dense, the scanner produces occasional erroneous returns on specular objects and "sees" through glass, resulting in inconsistencies between RGB and depth. We use a robust median filter to identify and reject spurious returns,

and manually mask additional regions in the *validation* and *test* sets with spurious or inconsistent returns. We provide the masks and the raw depth returns to allow users to implement alternative masking or inpainting schemes [65].

**Standard split**    We establish a train/validation/test split in order to ensure the reproducibility of our results as well as to make it easy to track progress of methods using DIODE. The validation set consists of curated crops from 15 indoor and 14 outdoor scans, while the test set consists of crops from 20 indoor and 20 outdoor scans. When curating scans in the validation and test partitions, we do not allow the fields-of-view of the selected crops to overlap by more than $20°$ in azimuth for validation scans, and $40°$ for test scans. No such restriction is used when selecting train crops.
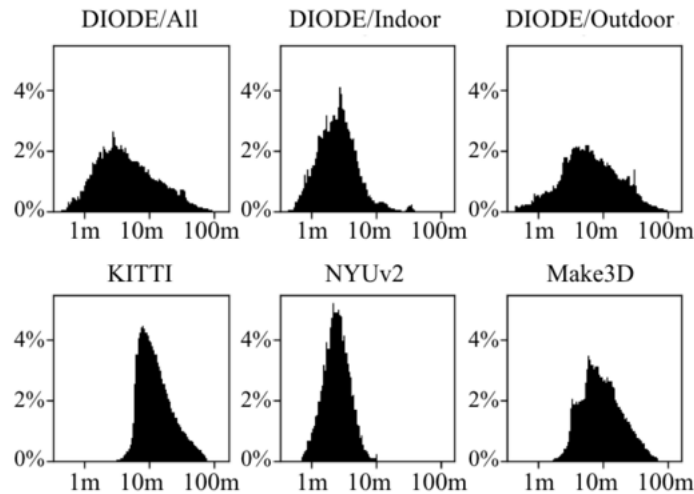


Figure 2.6: Distribution of measured depth values for DIODE (train set) and other popular RGB-D datasets.

**Dataset statistics**    Figure 2.6 compares the distribution of values in the depth maps in popular datasets to DIODE (values beyond 100 m are only found in DIODE and thus we clip the figures for DIODE for ease of comparison). Note that given that there are often objects both near and far from the camera in outdoor scans, the distribution of depth values is more diffuse in DIODE/outdoor than in KITTI. Only the much smaller and lower resolution Make3D is close to matching the diversity of DIODE depth values.

One important consequence of the highly accurate point clouds captured by the FARO scanner is that very accurate surface normal maps can be constructed; rather than obtaining normals from downsampled depth maps (as in other datasets such as NYUv2) we obtain our surface normals from the original point cloud, improving accuracy and detail.

**Applications of DIODE**   Though originally intended as a dataset for training of general (indoor/outdoor) depth and surface normal estimators, DIODE continues to find applications as an evaluation dataset: it has been used as a "zero-shot" test set for echo-location [66], ordinal depth regression [67], depth and scale prediction [68], and novel ranking losses for depth estimation [69]. The per-pixel, highly precise nature of DIODE data has also found an interesting "unintended" use —it was used to generate a synthetic dataset of physically-based images rendered behind glass to *train* a monocular reflection removal network [70]; the authors' architecture achieved impressive results with real-world reflection-distorted test data.  In the next section, we will describe some of the challenges we experienced when using DIODE as a *training* dataset for supervised depth estimation, and some of the limitations this exposed in a fully-supervised approach to this task.

### 2.3.3   Supervised Monocular Depth Estimation

Supervised depth estimation—predicting ground truth 3D structure from images using machine learning—dates back to the pioneering work of Saxena et al. [71], where a Markov Random Field was trained to predict depth using a small labeled dataset. Since then, deep learning-based architectures trained on large-scale benchmark datasets have proliferated [72, 54, 73, 74, 75]. Papers have proposed new loss functions [73] and architectures [75] and dataset-mixing strategies [76] to improve the robustness and accuracy of supervised depth estimators.

In this section, we will describe the common metrics for evaluating the quality of depth predictions, then describe a series of experiments on the DIODE dataset that reveal some of the limitations of a purely supervised approach to this task.

### 2.3.3.1 Metrics

A variety of metrics have been proposed to compare ground-truth depth maps [2] with neural network predictions (for an early review, see Cadena et al. (2016) [77]). Most popular metrics average over all valid ground-truth pixels (call that number $T$). For a pixel $i$, if there exists a ground-truth depth value $d_i$, call the predicted depth $\hat{d}_i$.

- absolute error scaled by the reciprocal of the groundtruth depth (abs rel): $\frac{1}{T} \sum_i \frac{|d_i - \hat{d}_i|}{d_i}$. This metric is among the most popular in comparing different models on datasets such as NYUv2, KITTI and DDAD
- square root of the mean squared error (rmse): $\sqrt{\frac{1}{T} \sum_i (\hat{d}_i - \hat{d}_i)^2}$
- mae and rmse between the log of predicted depth and log of ground-truth depth (mae $\log_{10}$ and rmse $\log_{10}$)
- percentage of predicted depth $\hat{d}_i$ within $\mathrm{thr}$ relative to ground-truth depth $d_i$, i.e., $\delta = \max(\frac{\hat{d}}{d}, \frac{d}{\hat{d}}) < \mathrm{thr}$

A major limitation of all of these metrics is that they favor methods that predict accurate depth over large and potentially uninteresting "background" regions. As an illustrative example, imagine a relatively small tire lying in the middle of a road, with a large tree right next to the road. A depth estimator that very accuracy predicts depth on the surface of the road but completely missing the tire would have a higher accuracy than a method that accurately reconstructs the tire on the road but has a noisier prediction on the tree (which represents many more pixels than the tire). In most applications, predicting a possible road hazard is much more important than predicting details on a non-driveable surface. As a solution, a number of methods that attempt to more accurately predict depth on certain classes (e.g., vehicles) include segmentation-based object- and category-level variants of the above metrics [78], but whole-image metrics are still common and will be used throughout this thesis.

---

[2]Depth "ground truth" can come from stereo depth, IR cameras, synthetic data, or projected LiDAR point clouds, but the metrics for comparing predictions with any of these sources of ground truth are typically the same regardless of the source of ground truth.

### 2.3.3.2 DIODE experiments

We evaluate the performance of a strong baseline monocular depth estimator on the DIODE dataset, and highlight the challenge of predicting high-resolution depth maps. We investigate the generalization between indoor and outdoor data on DIODE. As our baseline predictor, we use the DenseDepth [79] architecture, which achieved near-state-of-the-art results on both the NYUv2 and KITTI datasets, and thus served as a simple baseline to test the performance of neural networks on our indoor/outdoor dataset. These experiments are intended to provide a "sanity check" and assess the difficulty and utility of DIODE using currently available models.

**Model**    For monocular depth estimation, we use the same objective as in DenseDepth [79], which consists of L1 loss, depth gradient, and structural similarity (SSIM) [80]. The architecture uses a DenseNet-169 [81] pretrained on ImageNet as an encoder as well as a simple decoder with no batch normalization.

We train three models on the indoor (DIODE/Indoor) and outdoor (DIODE/Outdoor) subsets of DIODE, as well as the entire dataset (DIODE/All). During training, all networks are trained with the batch size of 4 for 30 epochs using Adam [82]. We start with a learning rate of $0.0001$ and decrease it by one-tenth after 20 epochs. The CNN is fed with a full-resolution image ($1024 \times 768$) and outputs a prediction at half of the resolution ($512 \times 384$). We apply the validity mask on ground truth during training. We employ random horizontal flips and random channel swaps for data augmentation. Note that no inpainting is applied to the ground truth before training. We do not fine-tune the model on DIODE/Indoor or DIODE/Outdoor after training on DIODE/All. During final evaluation, we apply $2\times$ upsampling to the prediction to match the size of the ground truth. The metrics are only evaluated on valid pixels.

The weight on each loss term is set as the same as that in the original model. We set the maximum depth to be $350\,\mathrm{m}$, the farthest range in the DIODE dataset. Other settings are identical to the original DenseDepth model. We evaluate the performance of the model on the validation set using standard pixel-wise error metrics and three pixel-accuracy metrics [54].
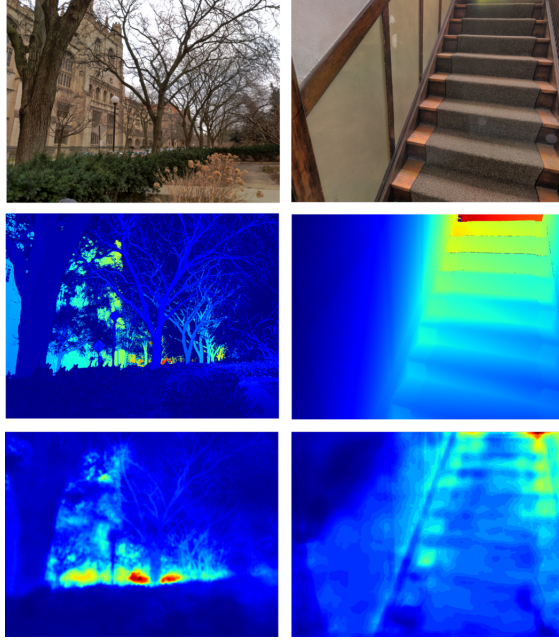
Figure 2.7: RGB frames, ground-truth depth, and predictions from a model trained on DIODE/All on indoor and outdoor validation data. Note that the model struggles to capture fine features in the outdoor image and reconstruct the stairs for the indoor frame.

| Experimental Setting | | Lower is better | | | | |
|---|---|---|---|---|---|---|
| Validation | Train | mae (m) | rmse (m) | abs rel | mae $\log_{10}$ | rmse $\log_{10}$ |
| | DIODE/Indoor | 1.2731 | 1.6093 | **0.3668** | **0.1494** | **0.1719** |
| DIODE/Indoor | DIODE/Outdoor | 2.2809 | 2.8876 | 1.1335 | 0.3076 | 0.3470 |
| | DIODE/All | **1.2280** | **1.5636** | 0.4506 | 0.1576 | 0.1830 |
| | DIODE/Indoor | 14.6225 | 20.1914 | 0.6637 | 0.5938 | 0.6553 |
| DIODE/Outdoor | DIODE/Outdoor | **7.6357** | **13.1280** | **0.4085** | **0.2132** | **0.3407** |
| | DIODE/All | 8.7813 | 14.1335 | 0.4533 | 0.2602 | 0.3751 |
| | DIODE/Indoor | 8.0641 | 11.0622 | 0.5178 | 0.3755 | 0.4178 |
| DIODE/All | DIODE/Outdoor | **5.0050** | 8.0970 | 0.7647 | 0.2596 | 0.3438 |
| | DIODE/All | 5.0704 | **7.9581** | **0.4520** | **0.2098** | **0.2807** |

Table 2.1: Baseline performance of monocular depth estimation for different training and validation sets.

**Analysis** Table 2.1 presents the results of generalization across indoor/outdoor scene types. For indoor validation, the model trained on DIODE/All outperforms the model trained on DIODE/Indoor in the metrics of mae and rmse, which indicates fewer outliers occur during evaluation. This may be explained by the larger size (roughly $2\times$ the images) of the outdoor dataset as well as the fact that outdoor scans capture many diverse objects at a wide range of distances (including near the scanner).

The performance slightly degrades on the outdoor validation when training on DIODE/All, this may be because most of a typical indoor scene is within approximately 20 m of the camera.

The model trained on DIODE/All outperforms the model trained only on DIODE/Indoor on indoor data, showing that the outdoor data (with many returns near the camera) allow the model to generalize better. The same model does not have much degradation on the outdoor val set, showing that it can predict reasonable depth both indoors and outdoors.

Note that absolute levels of accuracy are low despite use of a near state-of-the-art model at the time. In Figure 2.7, the low quality of depth predictions is evident–the model trained on the full dataset is a poor depth predictor in both the indoor setting (failing to predict the shape of the stairs) and the outdoor setting (failing to reconstruct most fine structures). Though the dataset includes tens of thousands of training frames, this data is not yet at the scale needed to *train* a quality monocular depth estimator.

**Conclusions**    Though the DIODE dataset has found interesting applications as an evaluation dataset, its main strength (the diversity of scenes) is also a disadvantage when it comes to training for depth estimation. We struggled to train an accurate depth estimator with even a near state-of-the-art model, though the same model performed well when trained and evaluated on NYUv2 (indoor) and KITTI (outdoor), showing that it is far more challenging to train on diverse data. Clearly, a much larger dataset would be required to train and test an accurate depth estimator on the same dataset. However, scaling to more scenes would have taken many more person-hours, given that despite its relatively modest size, DIODE was very labor-intensive to collect and process. This motivated us to consider a different paradigm, one where large-scale video datasets could be used to train depth estimators *with no ground truth*: self-supervised depth estimation.

### 2.3.4   Self-Supervised Depth Estimation

In Section 2.3.2, we reviewed the labor-intensive nature of depth label collection, and the difficulty in training a diverse depth estimator. Clearly, the requirement of (expensive) depth labels was a hindrance for training effective depth estimators, limiting the accuracy of learning-based approaches.

To alleviate the requirement of depth labels, Garg et al. [83] and Godard et al. [84] introduced an alternative strategy that involved training a monocular depth network with stereo images, leveraging Spatial Transformer Networks [85] to warp the right image into a synthesized version of the left. The resulting loss between synthesized and original left images can be defined in a fully-differentiable way [86], thus allowing the depth network to be self-supervised in an end-to-end fashion. Following Godard et al. [84], Zhou et al. [12] extended this self-supervised training to a purely monocular setting, where depth and pose networks are learned simultaneously from unlabeled video sequences obtained from a pre-calibrated pinhole camera.

Depth estimation is now cast as a view synthesis task, using an image reconstruction objective. Early results were blurry, and lacked the ability to reconstruct fine detail like thin structures. However, further improvements in the view synthesis loss [87, 88] and network architectures [9] led to large improvements, and eventually the accuracy of self- and semi-supervised approaches were competitive with fully supervised networks, eventually exceeding their accuracy [13, 88, 4, 9, 14]. These learned depth estimators have found numerous applications, including monocular 3D object detection, where "pseudo-LIDAR" [89] point cloud estimates obtained from monocular depth maps are used to predict 3D bounding boxes. Rather than focus on applications, in this thesis we will describe various expansions of the original method of Zhou et al. [12] to novel camera settings.

Next, we will describe the self-supervised depth learning framework based on the work of Zhou et al. [12] used in Chapters 3, 4 and 5.

**The framework**   Consider a video obtained from a moving camera, and for a given timestep, define a *target* frame $I_t$. Our goal is to predict a per-pixel depth map for this image. Next, define a

set of *context* images for this target frame, typically these are temporally nearby frames (e.g., if the target frame is sampled at time step $t$, the context images could be the images sampled at $t-1$ and $t+1$). The typical self-supervised depth architecture consists of two parts:

- A depth network that produces depth maps $\hat{D}_t$ for a target image $I_t$;

- A pose network that predicted the relative pose for pairs of target $t$ and context $c$ frames,
  $$\hat{\mathbf{X}}^{t\to c} = \left( \begin{smallmatrix} \hat{\mathbf{R}}^{t\to c} & \hat{\mathbf{t}}^{t\to c} \\ \mathbf{0} & \mathbf{1} \end{smallmatrix} \right) \in \mathrm{SE}(3)$$

To supervise these two networks, a reconstruction $\hat{I}_t$ of image $I_t$ is made using the predicted depth $\hat{d}^t$ and predicted relative pose $\hat{\mathbf{X}}^{t\to c}$. A warping map is used to resample context images $I_c$ into a prediction of the target $\hat{I}_t$. To obtain the warp map for this reconstruction, we use STN [85] via grid sampling with bilinear interpolation. This view synthesis operation is thus fully differentiable, enabling gradient back-propagation for end-to-end training. We define the pixel-warping operation as:

$$\hat{\mathbf{p}}^t = \pi\big(\hat{\mathbf{R}}^{t\to c}\phi(\mathbf{p}^t, \hat{d}^t, \mathbf{K}) + \hat{\mathbf{t}}^{t\to c}, \mathbf{K}\big) \tag{2.3.1}$$

where $\phi(\mathbf{p}, \hat{d}, \mathbf{K}) = \mathbf{P}$ is the unprojection of a pixel in homogeneous coordinates $\mathbf{p}$ to a 3D point $\mathbf{P}$ for a given estimated depth $\hat{d}$. Denote the projection of a 3D point back onto the image plane as $\pi(\mathbf{P}, \mathbf{K}) = \mathbf{p}$. Both operations require the camera parameters, which for the standard pinhole model [90] is defined by the $3 \times 3$ intrinsics matrix $\mathbf{K}$. The camera parameters are typically assumed to be known, and this restricts the self-supervised learning framework to sequences coming from calibrated, perspective cameras. The self-supervised loss to be minimized is of the form:

$$\mathcal{L}(I_t, \hat{I}_t) = \mathcal{L}_p(I_t, I_C) + \lambda_d \, \mathcal{L}_d(\hat{D}_t), \tag{2.3.2}$$

which is the combination of an appearance-based photometric loss $\mathcal{L}_p$ and a weighted depth smoothness loss $\mathcal{L}_d$, described below in more detail. This loss is then averaged per pixel and batch during training to produce the final value to be minimized.

In this thesis we do not explicitly model dynamic objects. The pixel-warping operation for the photometric loss is only valid for *static* scenes, and any moving objects are treated as noise, typically

through loss masking. Explicit modeling of dynamic objects (so that a valid target exists for both the static and dynamic scene) [4, 91] could further improve experimental results.

**Appearance-Based Loss.** Similar to Godard et al. [84] and Zhou et al. [12], the similarity between target $I_t$ and warped $\hat{I}_t$ images is estimated at the pixel level using Structural Similarity (SSIM) [86]. Unlike Mean Squared Error (MSE), which compares raw pixels, SSIM is a perceptual metric, decomposing error into luminance, contrast and structure components, computing over image windows computing local statistics (mean, variance, and covariance). It is combined with an L1 loss term:

$$\mathcal{L}_p(I_t, \hat{I}_t) = \alpha \, \frac{1 - \text{SSIM}(I_t, \hat{I}_t)}{2} + (1 - \alpha) \, \|I_t - \hat{I}_t\|. \tag{2.3.3}$$

To increase robustness against parallax or the presence of dynamic objects, we follow Godard et al. [13] and consider only the minimum pixel-wise photometric loss value for each context image in $I_C$. Similarly, we mask out static pixels by removing those with warped photometric loss $\mathcal{L}_p(I_t, \hat{I}_t)$ higher than the original photometric loss $\mathcal{L}_p(I_t, I_c)$.

**Depth Smoothness Loss.** To regularize the depth in textureless image regions, we incorporate an edge-aware term similar to Godard et al. [84], that penalizes high depth gradients in areas with low color gradients:

$$\mathcal{L}_s(\hat{D}_t) = |\delta_x \hat{D}_t| e^{-|\delta_x I_t|} + |\delta_y \hat{D}_t| e^{-|\delta_y I_t|}, \tag{2.3.4}$$

In the following chapters, we extend this basic framework to allow for increasingly diverse camera models and geometries.

# Chapter 3

# Learning Parametric Camera Models

<div style="text-align: right">

Your first 10,000 photographs are your worst.

—————————————————

Henri Cartier-Bresson.

</div>

Camera calibration is integral to robotics and computer vision algorithms that seek to infer geometric properties of the scene from visual input streams. In practice, calibration is a labor-intensive procedure requiring specialized data collection and careful tuning. This process must be repeated whenever the parameters of the camera change, which can be a frequent occurrence for mobile robots and autonomous vehicles. In contrast, self-supervised depth estimation approaches can bypass explicit calibration, typically by inferring per-frame projection models that optimize a view-synthesis objective.

In this chapter, we extend this self-supervised approach to explicitly self-calibrate a wide range of cameras from raw videos in the wild. We propose a learning algorithm to learn *per-sequence* calibration parameters using an efficient family of general camera models. Our procedure achieves calibration results with *sub-pixel re-projection error*, and outperforms other learning-based methods on a challenging depth benchmark. We validate our approach on a wide variety of camera geometries, including perspective, fisheye, and catadioptric.

## 3.1 Introduction

Cameras provide rich information about the scene, while being small, lightweight, inexpensive, and power efficient. Despite their wide availability, camera calibration largely remains a manual, time-consuming process that typically requires collecting images of known targets (e.g., checkerboards) as they are deliberately moved in the scene [27]. While applicable to a wide range of camera models [92, 93, 20], this process is tedious and has to be repeated whenever the camera parameters change. A number of methods perform calibration "in the wild" [94, 95, 96]. However, they rely on strong assumptions about the scene structure, which cannot be met during deployment in unstructured environments. Learning-based methods relax these assumptions, and regress camera parameters directly from images, either by using labelled data for supervision [97] or by extending the framework of self-supervised depth and ego-motion estimation [83, 12] to also learn per-frame camera parameters [4, 98].

While these methods enable learning accurate depth and ego-motion without calibration, they are either over-parameterized [98] (described in Chapter 4) or limited to near-pinhole cameras [4]. In contrast, in this chapter we describe a self-supervised camera calibration algorithm capable of learning expressive models of different camera geometries in a computationally efficient manner. In particular, our approach adopts a family of general camera models [1] that scales to higher resolutions than previously possible, while still being able to model highly complex geometries such as catadioptric lenses. Furthermore, our framework learns camera parameters on a per-sequence rather than per-frame basis, resulting in calibrations that are more accurate (and more stable) than those achieved using contemporary learning methods. We evaluate the re-projection error of our approach compared to conventional target-based calibration routines, showing comparable sub-pixel performance despite only using raw videos at training time.

The contributions described in this chapter can be summarized as follows:

- We propose to self-calibrate a variety of generic camera models from raw video using self-supervised depth and pose learning as a proxy objective, providing for the first time a calibra-

(a) Input　　　　　(b) Predicted depth　　　　(c) Rectified image



(d) Camera parameter re-calibration
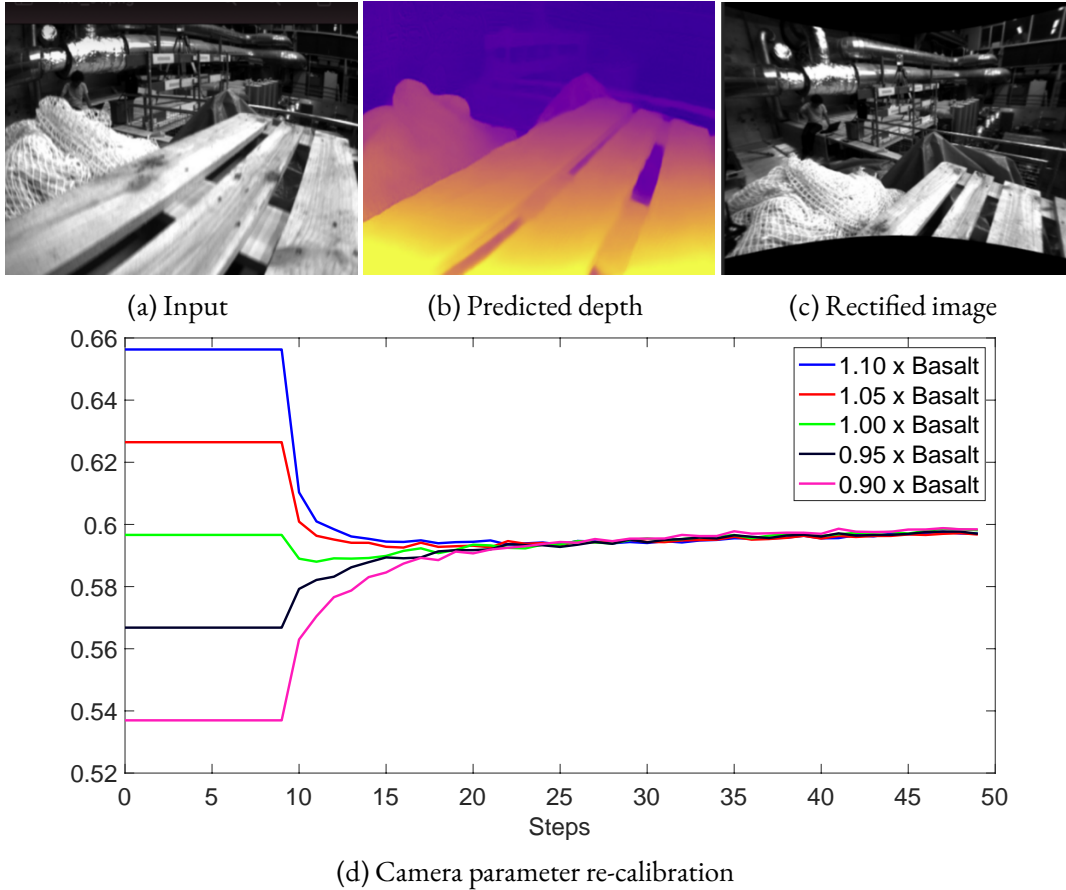
Figure 3.1: **Our self-supervised self-calibration procedure** can recover accurate parameters for a wide range of cameras using a structure-from-motion objective on raw videos (EuRoC dataset, top), enabling on-the-fly re-calibration and robustness to intrinsics perturbation (bottom).

tion evaluation of camera model parameters **learned purely through self-supervision**.

- We demonstrate the utility of our framework on challenging and radically different datasets, learning depth and pose on perspective, fisheye, and catadioptric images without architectural changes.

- We achieve **state-of-the-art depth evaluation results on the challenging EuRoC MAV dataset** by a large margin, using our proposed self-calibration framework.

## 3.2  Related Work

This chapter builds on the self-supervised depth and pose learning approach introduced by Zhou et al. [12] and reviewed in Section 2.3.4. Next we review prior work in traditional and learning-based camera calibration.

**Camera Calibration.** Traditional calibration for a variety of camera models uses targets such as checkerboards or AprilTags to generate 2D-3D correspondences, which are then used in a bundle adjustment framework to recover relative poses as well as intrinsics [27, 34]. Targetless methods typically make strong assumptions about the scene, such as the existence of vanishing points and known (Manhattan world) scene structure [94, 95, 96]. While highly accurate, these techniques require a controlled setting and manual target image capture to re-calibrate. Several models are implemented in OpenCV [99], kalibr [100]. These methods require specialized settings to work, limiting their generalizability.

**Camera Models.** The pinhole camera model is ubiquitous in robotics and computer vision [101, 102] and is especially common in recent deep learning architectures for depth estimation [12]. There are two main families of models for high-distortion cameras. The first is the "high-order polynomial" distortion family that includes pinhole radial distortion [103], omnidirectional [92], and Kannala-Brandt [93]. The second is the "unified camera model" family that includes the Unified Camera Model (UCM) [104], Extended Unified Camera Model (EUCM) [105], and Double Sphere Camera Model (DS) [1]. Both families are able to achieve low reprojection errors for a variety of different camera geometries [1], however the unprojection operation of the "high-order polynomial" models requires solving for the root of a high-order polynomial, typically using iterative optimization, which is a computationally expensive operation. Further, the process of calculating gradients for these models is non-trivial. In contrast, the "unified camera model" family has an easily computed, closed-form unprojection function. While our framework is applicable to high-order polynomial models, we choose to focus on the unified camera model family in this chapter.

**Learning Camera Calibration.** Work in learning-based camera calibration can be divided into

39

two types: *supervised* approaches that leverage ground-truth calibration parameters or synthetic data to train single-image calibration regressors; and *self-supervised* methods that utilize only image sequences. Our proposed method falls in the latter category, and aims to self-calibrate a camera system using only image sequences. Early work on applying CNNs to camera calibration focused on regressing the focal length [106] or horizon lines [107]; synthetic data was used for distortion calibration [108] and fisheye rectification [109]. Using panorama data to generate images with a wide variety of intrinsics, Lopez et al. [110] are able to estimate both extrinsics (tilt and roll) and intrinsics (focal length and radial distortion). DeepCalib [97] takes a similar approach: given a panoramic dataset, generate projections with different focal lengths. Then, they train a CNN to regress from a set of synthetic images $I$ to their (known) focal lengths $f$. Typically, training images are generated by taking crops of the desired focal lengths from 360 degree panoramas [111, 112]. While this can be done for any kind of image, and does not require image sequences, it does require access to panoramic images. Furthermore, the warped "synthetic" images are not the true 3D-2D projections. This approach has been extended to pan-tilt-zoom [113] and fisheye [109] cameras.

Other recent work [33, 4, 114] relaxes the assumption of a known camera matrix by learning the intrinsics in a self-supervised depth and ego-motion framework. These architectures allow training on completely uncalibrated videos in the wild, and can adapt to different focal lengths from different cameras because the camera parameters themselves are predicted from image frames in a fully self-supervised way. However, these methods are limited to a few fixed parametric camera models (usually the pinhole model or pinhole and distortion parameters) and cannot be trained on a wide variety of "cameras in the wild" (e.g. catadioptric cameras). Methods also exist for specialized problems like undistorting portraits [115], monocular 3D reconstruction [68], and rectification [116, 117].

## 3.3 Methodology

In Section 2.3.4 we described the self-supervised monocular depth learning framework that we use as proxy for self-calibration. Next we will describe the family of unified camera models we consider and how we learn their parameters end-to-end.

### 3.3.1 End-to-End Self-Calibration

UCM [104] is a parametric global central camera model that uses only five parameters to represent a diverse set of camera geometries, including perspective, fisheye, and catadioptric. A 3D point is projected onto a unit sphere and then projected onto the image plane of a pinhole camera, shifted by $\frac{\alpha}{1-\alpha}$ from the center of the sphere (Fig. 3.2). EUCM and DS are two extensions of the UCM model. EUCM replaces the unit sphere with an ellipse as the first projection surface, and DS replaces the one unit sphere with two unit spheres in the projection process. We self-calibrate all three models (in addition to a pinhole baseline) in our experiments. For brevity, we only describe the original UCM and refer the reader to Usenko et al. [1] for details on the EUCM and DS models.
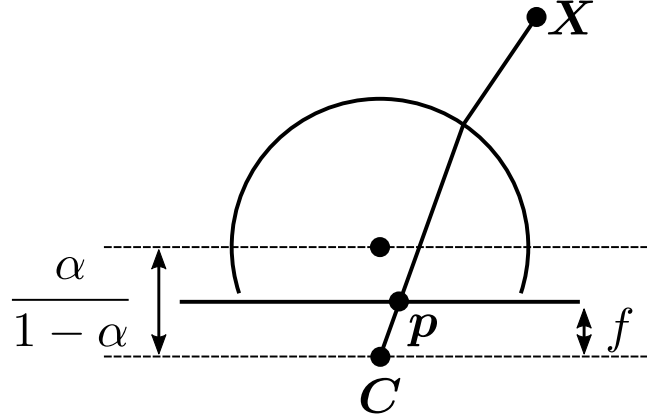


Figure 3.2: **The Unified Camera Model [1] used in our self-calibration pipeline**. Points are projected onto a unit sphere before being projected onto an image plane of a standard pinhole camera offset by $\frac{\alpha}{1-\alpha}$ from the sphere center.

There are multiple parameterizations for UCM [104], and we use the one from Usenko et al. [1] since it has better numerical properties. UCM extends the pinhole camera model $(f_x, f_y, c_x, c_y)$ with only one additional parameter $\alpha$. The 3D-to-2D projection of $\boldsymbol{P} = (x, y, z)$ is defined as

$$\boldsymbol{\pi}(\boldsymbol{P}, \boldsymbol{i}) = \begin{bmatrix} f_x \frac{x}{\alpha d + (1-\alpha)z} \\ f_y \frac{y}{\alpha d + (1-\alpha)z} \end{bmatrix} + \begin{bmatrix} c_x \\ c_y \end{bmatrix} \tag{3.3.1}$$

where the camera parameters are $\boldsymbol{i} = (f_x, f_y, c_x, c_y, \alpha)$ and $d = \sqrt{x^2 + y^2 + z^2}$
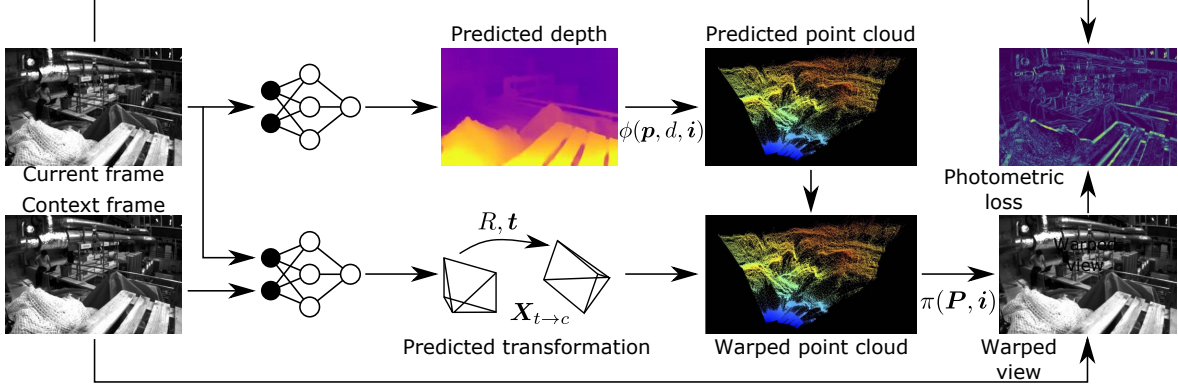
Figure 3.3: **Our self-supervised self-calibration architecture.** We use gradients from the photometric loss to update the parameters of a unified camera model (Fig. 3.2).

The unprojection operation of pixel $\boldsymbol{p} = (u, v, 1)$ at estimated depth $\hat{d}$ is:

$$\phi(\boldsymbol{p}, \hat{d}, \boldsymbol{i}) = \hat{d} \frac{\xi + \sqrt{1 + (1 - \xi^2)r^2}}{1 + r^2} \begin{bmatrix} m_x \\ m_y \\ 1 \end{bmatrix} - \begin{bmatrix} 0 \\ 0 \\ \hat{d}\zeta \end{bmatrix} \tag{3.3.2}$$

where

$$m_x = \frac{u - c_x}{f_x}(1 - \alpha) \qquad m_y = \frac{v - c_y}{f_y}(1 - \alpha) \tag{3.3.3a}$$

$$r^2 = m_x^2 + m_x^2 \qquad \zeta = \frac{\alpha}{1 - \alpha} \tag{3.3.3b}$$

As shown in Equations 3.3.1 and 3.3.2, the UCM camera model provides closed-form projection and unprojection functions that are both differentiable. Therefore, the overall architecture is end-to-end differentiable with respect to both neural network parameters (for pose and depth estimation) and camera parameters. This enables learning self-calibration end-to-end from the aforementioned view synthesis objective alone. At the start of self-supervised depth and pose training, rather than pre-calibrating the camera parameters, we initialize the camera with "default" values based on image shape only (for a detailed discussion of the initialization procedure, please see Section 3.4.5). Although the projection 3.3.1 and unprojection 3.3.2 are initially inaccurate, they quickly converge to highly accurate camera parameters with sub-pixel re-projection error (see Table 3.1).

As we show in our experiments, our method combines flexibility with computational efficiency. Indeed, our approach enables learning from heterogeneous datasets with potentially vastly differing sensors for which separate parameters $i$ are learned. As most of the parameters (in the depth and pose networks) are shared thanks to the decoupling of the projection model, this enables scaling up in-the-wild training of depth and pose networks. Furthermore, our method is efficient, with only one extra parameter relative to the pinhole model. This enables learning depth for highly-distortd catadioptric cameras at a much higher resolution than the method we introduce in Chapter 4; $(1024 \times 1024$ vs. $384 \times 384)$ this efficiency comes at the cost of generality. Note that, in contrast to other learning-based self-calibration methods [4, 98], we learn intrinsics per-sequence rather than per-frame. This increases stability compared to per-frame methods that exhibit frame-to-frame variability [98], and can be used over sequences of varying sizes.

## 3.4 Experiments

In this section we describe two sets of experimental validations for our architecture: (i) calibration, where we find that the re-projection error of our learned camera parameters compares favorably to target-based traditional calibration toolboxes; and (ii) depth evaluation, where we achieve state-of-the-art results on the challenging EuRoC MAV dataset.

### 3.4.1 Datasets

Self-supervised depth and ego-motion learning uses monocular sequences [12, 88, 4, 9] or rectified stereo pairs [88, 118] from forward-facing cameras [46, 9, 61]. Given that our goal is to learn camera calibration from raw videos in challenging settings, we use the standard KITTI dataset as a baseline, and focus on the more challenging and distorted EuRoC [3] fisheye sequences. For more information on these datasets, please refer to Section 2.3.1.

- **KITTI** [46]. We use this dataset to show that our self-calibration procedure is able to

accurately recover pinhole intrinsics alongside depth and ego-motion.

- **EuRoC** [3]. For calibration evaluation, we follow Usenko et al. [1] and use the calibration sequences from the dataset. We evaluate the UCM, EUCM and DS camera models in terms of re-projection error.

- **OmniCam** [55]. Since this catadioptric sequence does not provide ground-truth depth information, we only provide qualitative results.

## 3.4.2 Training Protocol

We implement the group of unified camera models described in Usenko et al. [1] as differentiable PyTorch [119] operations, modifying the self-supervised depth and pose architecture of Godard et al. [88] to jointly learn depth, pose, and the unified camera model intrinsics. We use a learning rate of 2e-4 for the depth and pose network and 1e-3 for the camera parameters. We use a StepLR scheduler with $\gamma = 0.5$ and a step size of 30. All of the experiments are run for 50 epochs. The images are augmented with random vertical and horizontal flip, as well as color jittering. We train our models on a Titan X GPU with 12 GB of memory, with a batch size of 16 when training on images with a resolution of $384 \times 256$.

| Method | Mean Reprojection Error | |
| --- | --- | --- |
| | Target-based | Learned |
| Pinhole | 1.950 | 2.230 |
| UCM [104] | 0.145 | 0.249 |
| EUCM [105] | 0.144 | 0.245 |
| DS [1] | 0.144 | 0.344 |

Table 3.1: **Mean reprojection error on EuRoC** at $256 \times 384$ resolution for UCM, EUCM and DS models using (left) AprilTag-based toolbox calibration Basalt [2] and (right) our self-supervised learned (L) calibration. Note that despite using no ground-truth calibration targets, our self-supervised procedure produces sub-pixel reprojection error.

### 3.4.3   Camera Self-Calibration

We evaluate the results of the proposed self-calibration method on the EuRoC dataset; detailed depth estimation evaluations are provided in Section 3.4.6. To our knowledge, ours is the first direct calibration evaluation of self-supervised intrinsics learning; although Gordon et al. [4] compare *ground-truth* calibration to their per-frame model, they do not evaluate the re-projection error for their learned parameters.

Following Usenko et al. [2], we evaluate our self-supervised calibration method on the family of unified camera models: UCM, EUCM, and DS, as well as the perspective (pinhole) model. As a lower bound, we use the Basalt [2] toolbox and compute camera calibration parameters for each unified camera model using the calibration sequences of the EuRoC dataset. We note that unlike Basalt, our method regresses the intrinsic calibration parameters directly from raw videos, without using any of the calibration sequences.

| Method | $f_x$ | $f_y$ | $c_x$ | $c_y$ | $\alpha$ | $\beta$ | $\xi$ | $w$ |
|---|---|---|---|---|---|---|---|---|
| UCM (L) | 237.6 | 247.9 | 187.9 | 130.3 | 0.631 | — | — | — |
| UCM (B) | 235.4 | 245.1 | 186.5 | 132.6 | 0.650 | | | |
| EUCM (L) | 237.4 | 247.7 | 186.7 | 129.1 | 0.598 | 1.075 | — | — |
| EUCM (B) | 235.6 | 245.4 | 186.4 | 132.7 | 0.597 | 1.112 | | |
| DS (L) | 184.8 | 193.3 | 187.8 | 130.2 | 0.561 | — | -0.232 | — |
| DS (B) | 181.4 | 188.9 | 186.4 | 132.6 | 0.571 | | -0.230 | |

Table 3.2: **Intrinsic calibration evaluation of different methods** on the EuRoC dataset, where B denotes intrinsics obtained from Basalt, and L denotes learned intrinsics.
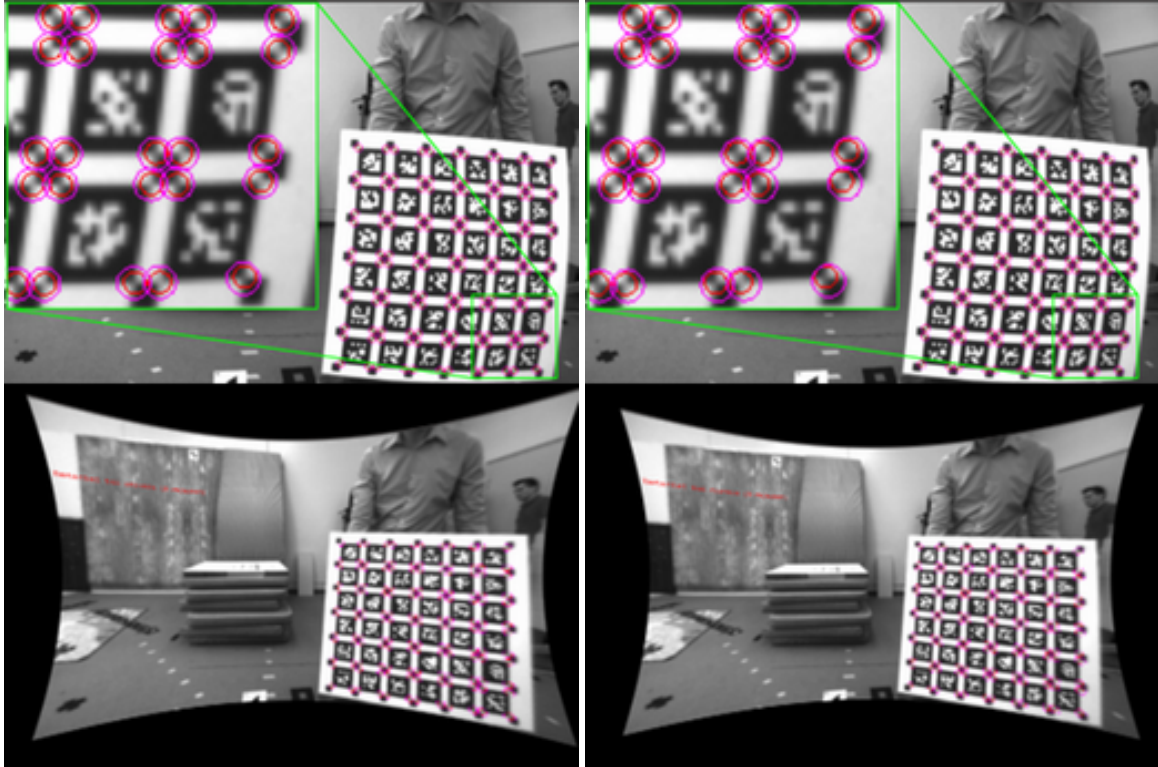
Figure 3.4: **EuRoC rectification results** using images from the calibration sequences. Each column visualizes the results rendered using (left) the Basalt calibrated intrinsics and (right) our learned intrinsics. The top row shows that detected (small circles) and reprojected (big circles) corners are close using both calibration methods. The bottom row shows the same images after rectification.

Table 3.1 summarizes our re-projection error results. We use the EuRoC AprilTag [28] calibration sequences with Basalt to measure re-projection error using the full estimation procedure (Table 3.1 — *Target-based*) and learned intrinsics (Table 3.1 — *Learned*). For consistency, we optimize for both intrinsics and camera poses for the baselines and only for the camera poses for the learned intrinsics evaluation. Note that with learned intrinsics, UCM, EUCM and DS models all achieve sub-pixel mean projection error despite the camera parameters having been learned from raw video data.

Table 3.2 compares the target-based calibrated parameters to our learned parameters for different camera models trained on the *camo* sequences of the EuRoC dataset. Though the parameter vectors were initialized with no prior knowledge of the camera model and updated purely based on gradients from the reprojection error, they converge to values very close to the output of a procedure that uses bundle adjustment on calibrated image sequences.

| Perturbation | $f_x$ | $f_y$ | $c_x$ | $c_y$ | $\alpha$ | $\beta$ | **MRE** |
|---|---|---|---|---|---|---|---|
| $I_{1.10}$ init | 242.3 | 253.6 | 189.5 | 130.7 | 0.5984 | 1.080 | 0.409 |
| $I_{1.05}$ init | 241.3 | 252.3 | 188.5 | 130.5 | 0.5981 | 1.078 | 0.367 |
| $I_c$ init | 240.2 | 251.4 | 187.9 | 130.0 | 0.5971 | 1.076 | 0.348 |
| $I_{0.95}$ init | 239.5 | 250.9 | 187.8 | 129.2 | 0.5970 | 1.076 | 0.332 |
| $I_{0.90}$ init | 238.8 | 249.6 | 187.7 | 129.1 | 0.5968 | 1.071 | 0.298 |
| $I_c$ | 235.6 | 245.4 | 186.4 | 132.7 | 0.597 | 1.112 | 0.144 |

Table 3.3: **EUCM perturbation test results.** With perturbed initialization, all intrinsic parameters achieve sub-pixel convergence for mean re-projection error (**MRE**), with only a small offset to the Basalt calibration numbers.

### 3.4.4 Image Rectification

Using our learned camera parameters, we rectify calibration sequences on the EuRoC dataset to demonstrate the quality of the calibration. EuRoC was captured with a fisheye camera and exhibits a high degree of radial distortion that causes the straight edges of the checkerboard grid to be curved. In Figure 3.4, we see that our learned parameters allow for the rectified grid to track closely to the true underlying checkerboard.

### 3.4.5 Re-calibration: Perturbation Experiments

Thus far, we have assumed to have no prior knowledge of the camera calibration. In many real-world robotics settings, however, one may want to re-calibrate a camera based on a potentially incorrect prior calibration. Generally, this requires the capture of new calibration data. Instead, we can initialize our parameter vectors with this initial calibration (in this setting, a perturbation of Basalt calibration of the EUCM model) and see the extent to which self-supervision can nudge the parameters back to their "true value".

Given Basalt parameters $I_c = [f_x, f_y, c_x, c_y, \alpha, \beta]$, we perturb them as $I_{1.1} = 1.1 \times I_c$, $I_{1.05} = 1.05 \times I_c$, $I_{0.95} = 0.95 \times I_c$, $I_{0.9} = 0.9 \times I_c$ and initialize the camera parameters at the beginning of training with these values. All runs have warm start, i.e., freezing the gradients for the intrinsics

(a) $f_x$       (b) $f_y$

(c) $c_x$       (d) $c_y$

(e) $\alpha$       (f) $\beta$
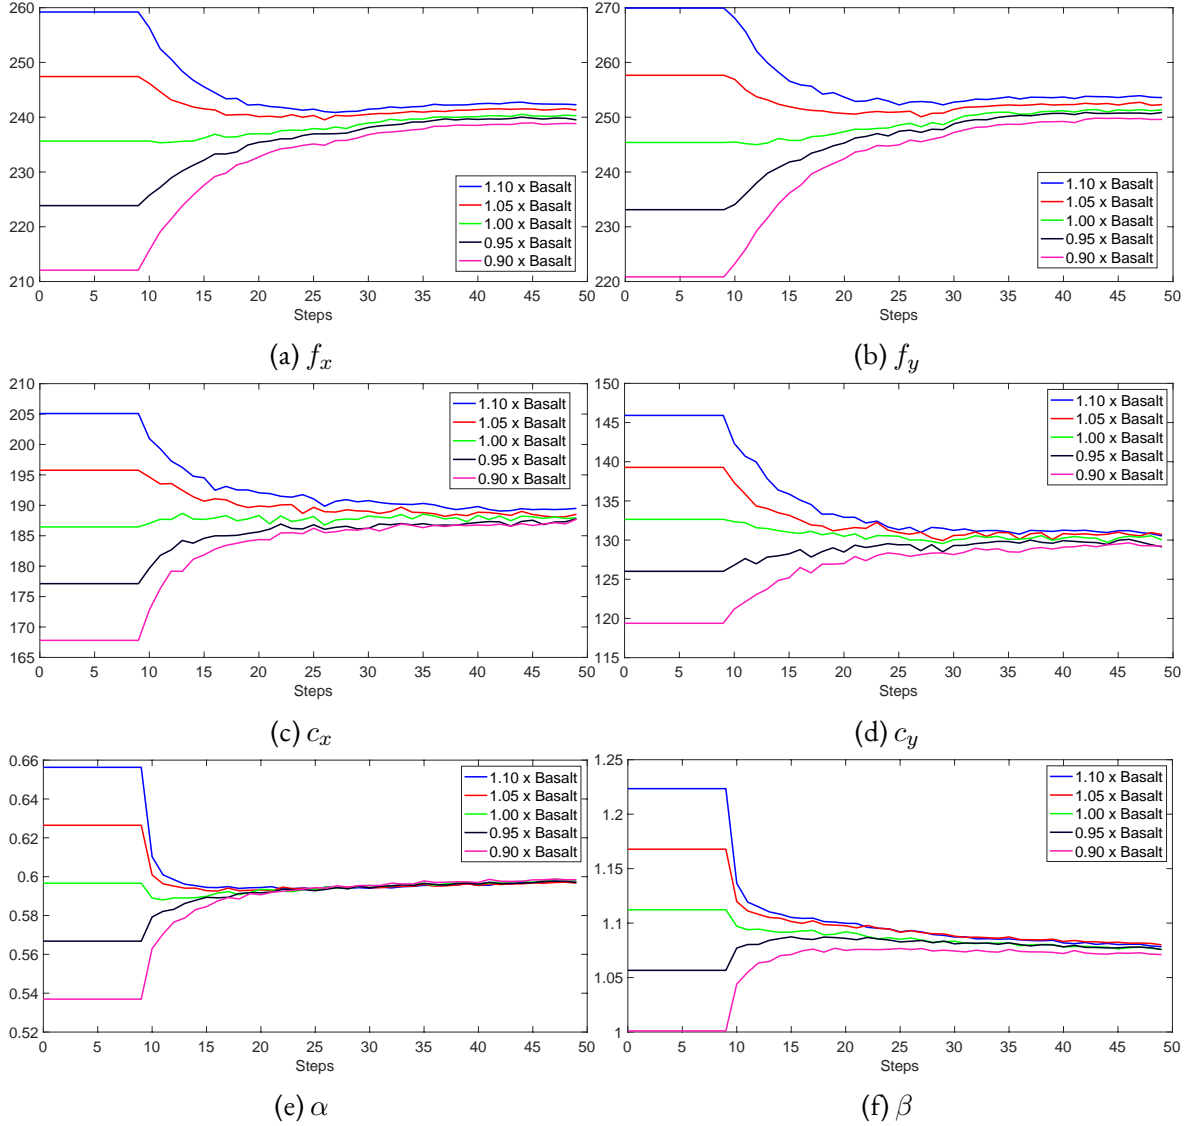
Figure 3.5: **EuRoC perturbation test**, showing how our proposed learning-based method is able to recover from changes in camera parameters for online self-calibration.

for the first 10 epochs while we train the depth and pose networks. As Figure 3.5 shows, our method converges to within $3\%$ of the Basalt estimate for each parameter. Table 3.3 provides the values of the converged parameters along with the mean projection error (MRE) for each experiment.

| Method | Camera | Abs Rel↓ | Sq Rel↓ | RMSE↓ | $\delta_{1.25}$ ↑ |
|---|---|---|---|---|---|
| Gordon et al. [4] | K | 0.129 | 0.982 | 5.23 | 0.840 |
| Gordon et al. [4] | L(P) | 0.128 | 0.959 | 5.23 | 0.845 |
| Vasiljevic et al. [98] | K(NRS) | 0.137 | 0.987 | 5.33 | 0.830 |
| Vasiljevic et al. [98] | L(NRS) | 0.134 | 0.952 | 5.26 | 0.832 |
| Ours | L(P) | 0.129 | **0.893** | 4.96 | 0.846 |
| Ours | L(UCM) | **0.126** | 0.951 | **4.89** | **0.858** |

Table 3.4: **Quantitative depth evaluation on the KITTI [3] dataset**, using the standard *Eigen* split and the *Garg* crop, for distances up to 80m (with median scaling). K and L(·) denote known and learned intrinsics, respectively. P means pinhole model.

| Method | Camera | Abs Rel↓ | Sq Rel↓ | RMSE↓ | $\alpha_1$ ↑ |
|---|---|---|---|---|---|
| Gordon et al. [4] | PB | 0.332 | 0.389 | 0.971 | 0.420 |
| Vasiljevic et al. [98] | NRS | 0.303 | 0.056 | 0.154 | 0.556 |
| Ours | UCM | 0.282 | 0.048 | 0.141 | 0.591 |
| Ours | EUCM | **0.278** | **0.047** | **0.135** | **0.598** |
| Ours | DS | **0.278** | 0.049 | 0.141 | 0.584 |

Table 3.5: **Quantitative depth evaluation of different methods on the EuROC [3] dataset**, using the evaluation procedure in Gordon et al. [4] with center cropping. The training data consists of "Machine Room" sequences and the evaluation is on the "Vicon Room 201" sequence (with median scaling). PN means Plumb Bob model.

## 3.4.6 Depth Estimation

While we use depth and pose estimation as proxy tasks for camera self-calibration, the unified camera model framework allows us to achieve meaningful results compared to prior camera-learning-based approaches (see Figures 3.6 and 3.7).

**KITTI results.** Table 3.4 presents the results of our method on the KITTI dataset. We note that our approach is able to model the simple pinhole setting, achieving results that are on par with approaches that are tailored specifically to this camera geometry. Interestingly, we see an increase in performance using the UCM model, which we attribute to the ability to further account for and correct calibration errors.

**EuRoC results.** Compared to KITTI, EuRoC is a significantly more challenging dataset that

| Dataset | Abs Rel↓ | Sq Rel↓ | RMSE↓ | $\alpha_1 \uparrow$ | $\alpha_2 \uparrow$ | $\alpha_3 \uparrow$ |
|---|---|---|---|---|---|---|
| EuRoC [4] | 0.265 | **0.042** | 0.130 | 0.600 | 0.882 | **0.966** |
| EuRoC+KITTI | **0.244** | 0.044 | **0.117** | **0.742** | **0.907** | 0.961 |

Table 3.6: **Quantitative multi-dataset depth evaluation** on EuRoC (without cropping and with median scaling).

involves cluttered indoor sequences with six-DoF motion. Compared to the per-frame distorted camera models of Gordon et al. [4] and Vasiljevic et al. [98] (Chapter 4), we achieve significantly better absolute relative error, especially with EUCM, where the error is reduced by 16% (see Table 3.5).
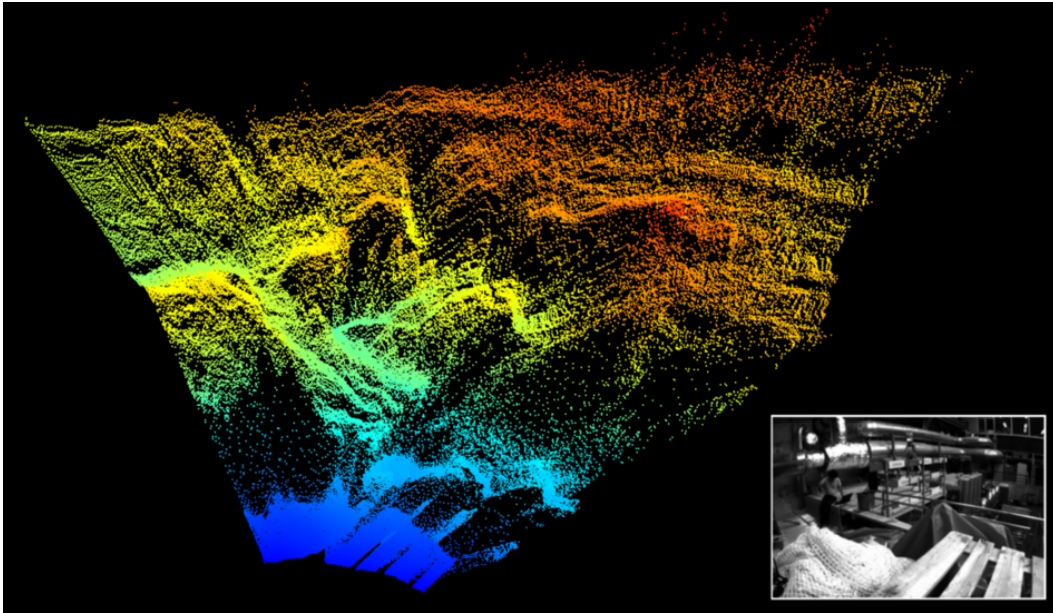


Figure 3.6: **Self-supervised monocular pointcloud** for EuRoC, obtained by unprojecting predicted depth with our learned camera parameters (input image on the bottom right).

**Combining heterogeneous datasets.** One of the strengths of the unified camera model is that it can represent a wide variety of cameras without prior knowledge of their specific geometry. As long as we know which sequences come from which camera, we can learn separate calibration vectors that share the same depth and pose networks. This is particularly useful as a way to improve performance on smaller datasets, since it enables one to take advantage of unlabeled data from other sources. To evaluate this property, we experimented with mixing KITTI and EuRoC. In this experiment, we

reshaped the KITTI images to match those in the EuRoC dataset (i.e., $384 \times 256$). As Table 3.6 shows, our algorithm is able to take advantage of the KITTI images to improve performance on the EuRoC depth evaluation.
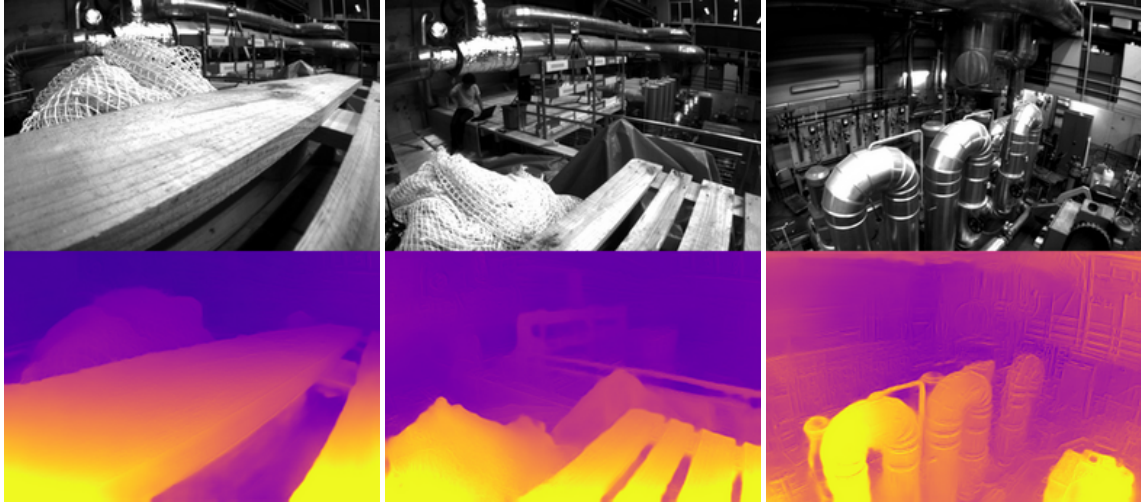


(a) EuRoC
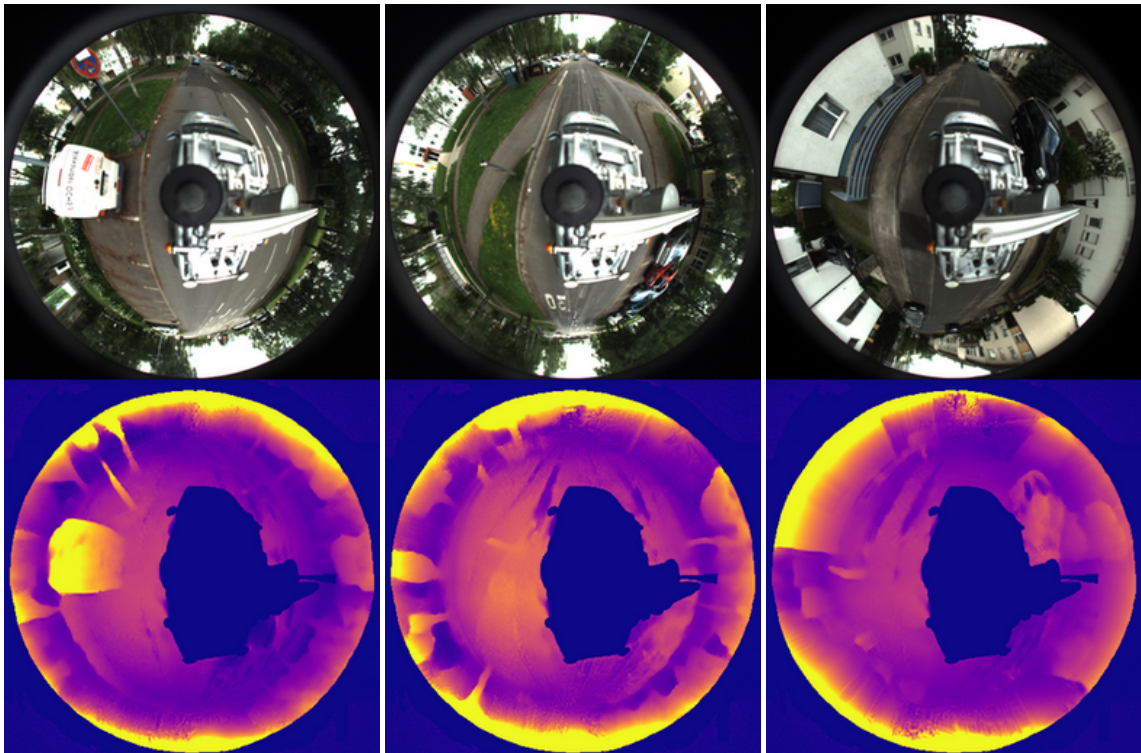


(b) OmniCam

Figure 3.7: **Qualitative depth estimation results** on non-pinhole datasets with (a) fisheye and (b) catadioptric images.

## 3.5 Conclusion

We proposed a procedure to self-calibrate a family of general camera models using self-supervised depth and pose estimation as a proxy task. We rigorously evaluated the quality of the resulting camera models, demonstrating sub-pixel calibration accuracy comparable to manual target-based toolbox calibration approaches. Our approach generates per-sequence camera parameters, and can be integrated into any learning procedure where calibration is needed and the projection and un-projection operations are interpretable and differentiable. As shown in our experiments, our approach is particularly amenable to online re-calibration, and can be used to combine datasets of different sources, learning independent calibration parameters while sharing the same depth and pose network.

# Chapter 4

# Learning Non-Parametric Camera Models

Art is both the taking and giving of beauty; the turning out to the light the inner folds of the awareness of the spirit. It is the recreation on another plane of the realities of the world; the tragic and wonderful realities of earth and men, and of all the inter-relations of these.

Ansel Adams.

In Chapter 3, we introduced a a procedure for self-supervised self-calibration of a family of parametric camera models. This parametric model is a poor approximation in some cases (i.e., a perspective camera underwater suffers from significant non-linear distortion caused by the camera-water interface). In this chapter, we go one step further, showing that self-supervision can be used to model *any central viewing geometry*. Inspired by the "raxel" model of Grossberg and Nayar [120], we introduce Neural Ray Surfaces (NRS) [98], a network that learns per-pixel viewing rays, learned entirely from unlabeled raw video sequences. We demonstrate the use of NRS for self-supervised learning of visual odometry and depth from videos obtained using a wide variety of camera setups, including underwater imaging, cameras behind windshields, and catadioptric cameras.
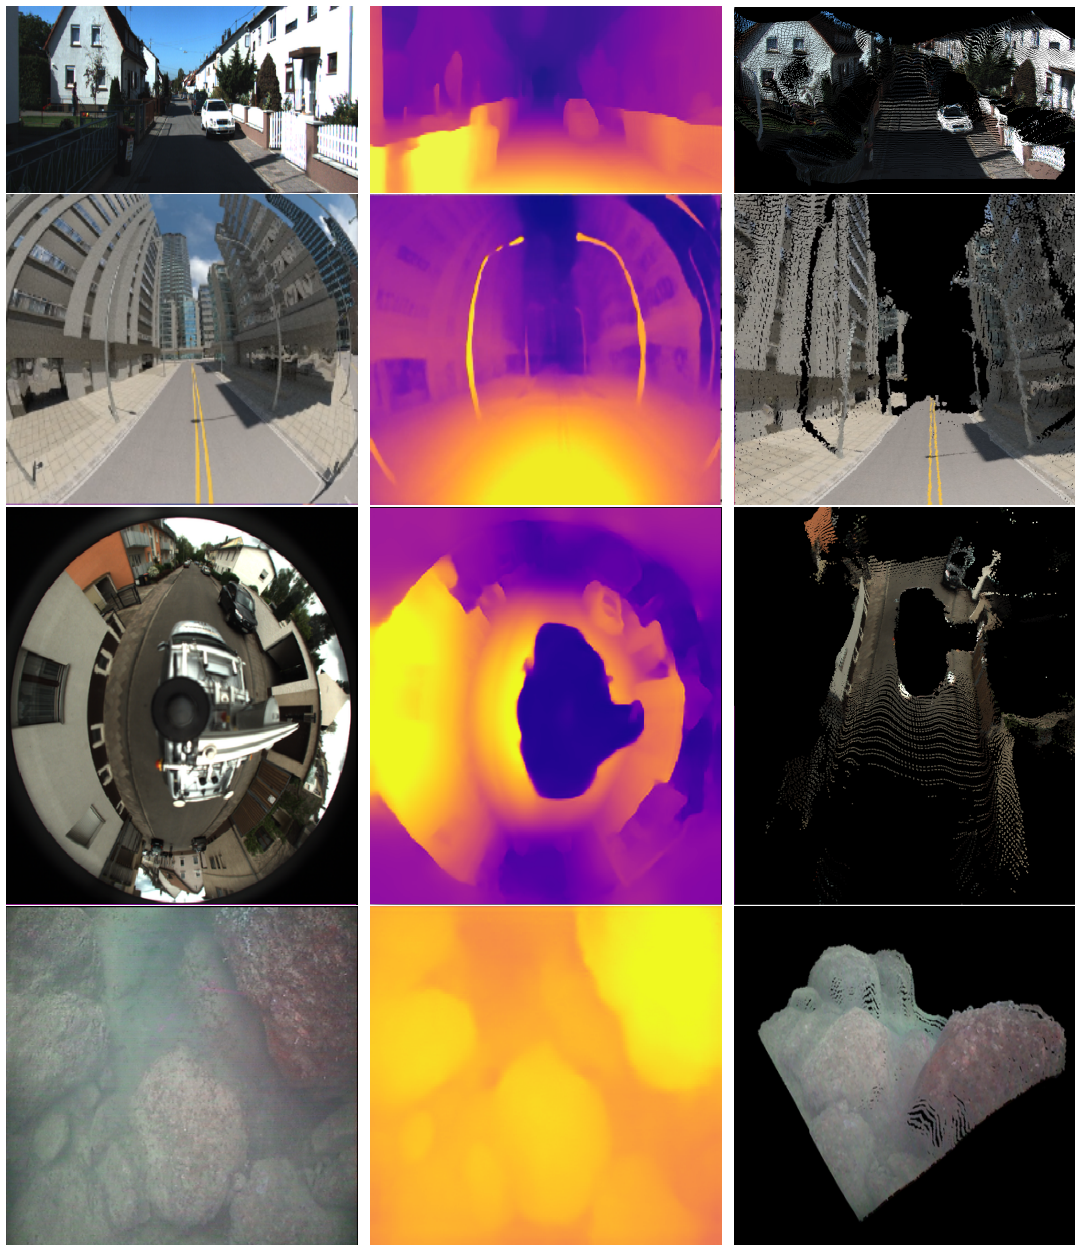
|              |              |               |
|:------------:|:------------:|:-------------:|
| (a) Input image | (b) Depth map | (c) Pointcloud |

Figure 4.1: **Our self-supervised Neural Ray Surfaces** can learn a wide variety of projection geometries purely from video sequences, including **pinhole** (top row, *KITTI*); **fisheye** (second row, *Multi-FOV*); **catadioptric** (third row, *OmniCam*); and **underwater** (bottom row, *Underwater Caves*).

## 4.1 Introduction

The perspective pinhole camera model [90] is ubiquitous due to its simplicity—it has few parameters and is easy to calibrate. Recently, deep neural architectures that rely on the pinhole assumption with geometric constraints have led to major advances in tasks such as monocular 3D detection [89] and depth estimation [12]. These networks are generally trained on curated and rectified image datasets where the pinhole assumption is appropriate. Recent work [4] has shown that the parameters for the pinhole camera model can be learned in a fully self-supervised way, thus enabling self-supervised learning on videos where calibration might not be available and mixing data from different cameras during training.

In Chapter 3, we reviewed a method for self-calibrating a more general family of camera models that encompass a wide variety of possible distortions. [121]. Despite these advances, there are a variety of settings where the pinhole assumption does not hold—from fisheye and catadioptric lenses to physical arrangements that break the pinhole assumption (e.g., a dashboard camera behind a windshield [39], or a camera underwater [122]).

The pinhole model allows for closed-form projection and unprojection operations, and thus can be easily used as a module in deep architectures, either fixed and precomputed or learned [4]. Parametric distortion models for pinhole cameras as well as models for more complex lens designs [93, 123] can also be adapted for deep architectures, but adapting these models to learn depth and ego-motion has three major disadvantages: (1) distortion models are generally a simplification of complex lens distortion, leading them to only be approximately correct; (2) a new differentiable projection architecture needs to be created for each camera model; and (3) there are settings where standard parametric models are not applicable, such as cameras behind a windshield or underwater.

Instead of adapting individual camera models [123], we propose the *end-to-end self-supervised learning of a differentiable projection model* from raw un-calibrated videos, in addition to depth and ego-motion. The generic camera model of Grossberg and Nayar [20] directly relates pixels to viewing rays, allowing for a per-pixel ray surface that can model a wide variety of distortions and lens

systems. The representational power of this model comes at the cost of complexity, leading to a large literature on generic camera calibration [124, 37, 125, 39]. In particular, the projection operation is considerably more complex than in the perspective model, generally requiring a computationally expensive optimization step to project 3D points to pixels.

Our **Neural Ray Surface** (NRS) model is differentiable and resource-efficient, allowing its use as a geometric module in the standard self-supervised depth and ego-motion setting of Zhou et al. [12]. In contrast to the pinhole intrinsics prediction module in Gordon et. al. [4], our model can be trained on datasets captured with radically different (unknown) cameras (Figure 6.1). We demonstrate learning depth and ego-motion on pinhole, fisheye, and catadioptric datasets, showing that our model can learn accurate depth maps and odometry where the standard perspective-based architecture, which is an incorrect model for non-pinhole lenses, diverges. We evaluate the strength of our model on several depth and visual odometry tasks that until now were considered beyond what is possible for learning-based self-supervised monocular techniques.

Our main contributions are as follows:

- We show that it is possible to learn a **pixel-wise projection model directly from video sequences** without the need for any prior knowledge of the camera system.

- We devise a **differentiable extension** for the unprojection and projection operations that define a generic ray surface model, thus allowing the **end-to-end learning of ray surfaces** for a given target task.

- We replace the standard pinhole model in the self-supervised monocular setting with our proposed ray surface model, thus enabling the learning of depth and pose estimation for many camera types, including for the first time on **catadioptric cameras.**

## 4.2 Related Work

**Generic Camera Models**    The differentiable ray surface model in our architecture is inspired by the general camera model of Grossberg and Nayar [20]. This model directly relates pixels with viewing rays, treating the camera as a black box [124]. It is applicable to many different imaging systems, including omnidirectional catadioptric cameras, fisheye cameras, pinhole cameras behind refractive surfaces such as windshields, etc. Despite the appealing generality of these camera models, calibration with such a large number of parameters remains challenging.

There exist multiple variations of this model and techniques for calibration [124, 37, 125, 39], as well as investigations into distortion calibration [126, 127] and multi-view geometry [128, 129]. Recent works have explored spline-based ray surface models to simplify calibration, reducing the number of parameters to be estimated [130, 131, 39].

Our NRS model shares the same projection model as that of Grossberg and Nayar [20], however our focus in this chapter is on using NRS as a tool for end-to-end learning of monocular depth and pose with arbitrary cameras, rather than calibration. We leave the investigation of self-supervised learning as a *calibration* tool for general cameras to future work.

In Chapter 3, we introduced a self-calibration procedure for a simpler (though still fairly general [1]) family of camera models. That model can well approximate perspective, fisheye, and catadioptric cameras, and has a significant reduction in parameters compared to NRS. The gains in computational efficiency come at the cost of generality; NRS is particularly applicable to settings where parametric camera models are less suitable. Indeed, the per-frame nature of NRS (compared to the per-sequence model in Chapter 3) is a strength in settings where the projection model is changing, for example in underwater imaging (see Section 4.4.6).

## 4.3 Methodology

As discussed in Chapter 2, a camera model is defined by two operations: the *unprojection* from image pixels $\mathbf{p}$ to 3D points $\mathbf{P}$, i.e., $\phi(\mathbf{p}, d) = \mathbf{P}$; and the *projection* of 3D points onto the image plane, i.e., $\pi(\mathbf{P}) = \mathbf{p}$. The standard pinhole perspective model [90] provides simple closed-form solutions to these two operations, as matrix-vector products (Figure 4.2a).

In the generic camera model of Grossberg and Nayar [20], the camera model consists of a ray surface that associates each pixel with a corresponding direction, offering a non-parametric association between 3D points and image pixels. In this model, although unprojection is simple and can be computed in closed form, the projection operation has no closed-form solution and is non-differentiable, which makes it unsuitable for learning-based applications (Figure 4.2b). Below we describe our variant of this generic camera model that is differentiable (see Figure 6.2a for the full architecture), and thus amenable to end-to-end learning in a self-supervised monocular setting.

**Notation** As described in Section 2.2.5, we think of general cameras as a bundle of *viewing rays*. For pixel $i$, this viewing ray is represented by a camera origin $\mathbf{o}_c$ and viewing direction $\mathbf{r}_i$. In our experiments, we assume that the monocular camera is central (so that the ray origin is the same for each pixel [37]), and without loss of generality place it at the origin of the reference coordinate system. Thus, $\mathbf{o}_c = \mathbf{0}$ and the camera model is fully described by the viewing directions $\mathbf{r}_i$. We will denote the array of viewing directions, or *ray surface*, by $\mathbf{Q}$, so that for a pixel $\mathbf{p}_i$, $\mathbf{Q}(\mathbf{p}_i) = \mathbf{r}_i$.

### 4.3.1 Unprojection

Given the above definition, for any pixel $\mathbf{p}_i$ and a predicted depth $\hat{d}_i$, we can easily obtain its corresponding 3D point $\mathbf{P}_i$ as follows:

$$\mathbf{P}_i = \hat{d}_i \hat{\mathbf{Q}}(\mathbf{p}_i) \tag{4.3.1}$$
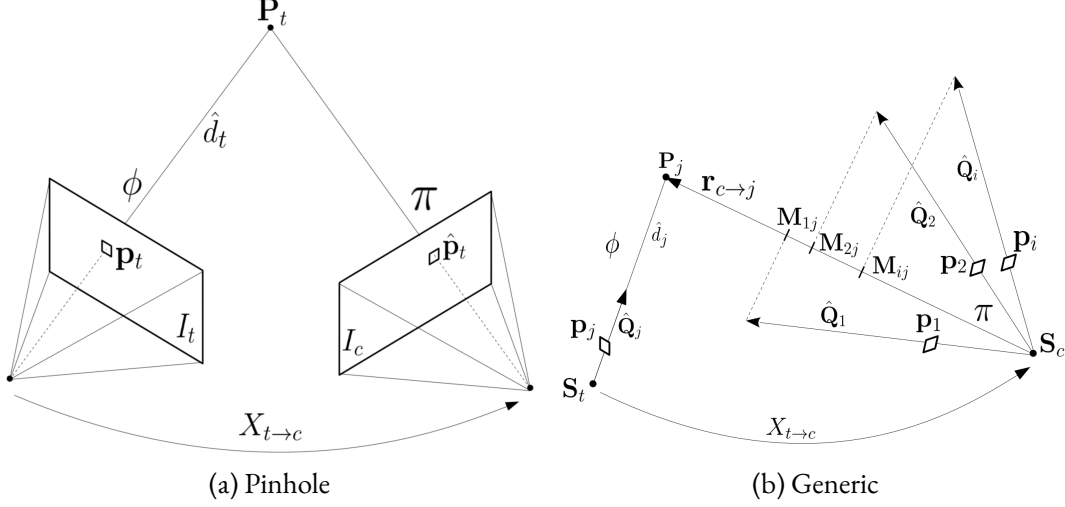
Figure 4.2: **Unprojection $\phi$ and projection $\pi$ operations** for (a) the standard pinhole; and (b) our proposed neural ray surface, for a single pixel $\mathbf{p}_j$ considering target $I_t$ and context $I_c$ images. Straight arrows in (b) represent unitary ray surface vectors $\mathbf{Q}(\mathbf{p})$, drawn out of scale to facilitate visualization. In this example, $\mathbf{p}_1$ is associated to $\mathbf{p}_j$, since it satisfies Equation 4.3.3.

In other words, we scale the predicted ray vector $\hat{\mathbf{r}}_i$ by the predicted depth $\hat{d}_i$.

### 4.3.2 Projection

Consider $\mathcal{P}_t = \{\mathbf{P}_j\}_{j=1}^{HW}$, produced by unprojecting pixels from $I_t$ as 3D points. In the standard pinhole camera model, projection is a simple matrix-vector product (Equation 2.2.4). For the proposed neural ray surface, however, for each 3D point $\mathbf{P}_j$ we must find the corresponding pixel $\mathbf{p}_i \in I_c$ with ray direction $\hat{\mathbf{r}}_i = \hat{\mathbf{r}}_c(\mathbf{p}_i)$ that most closely matches the direction of $\mathbf{P}_j$ to the camera center $\mathbf{S}_c = \mathbf{0}$ (see Figure 4.2b). Call this direction $\mathbf{r}_{c \to j} = \mathbf{P}_j - \mathbf{S}_c = \mathbf{P}_j$. Thus, we must find $\mathbf{p}_i^*$ such that:

$$\mathbf{p}_i^* = \arg\max_{\mathbf{p_i} \in I_c} \langle \hat{\mathbf{Q}}_c(\mathbf{p_i}), \mathbf{r}_{c \to j} \rangle \tag{4.3.2}$$

Solving this problem requires searching over the entire ray surface $\hat{\mathbf{Q}}_c$ and can be computationally expensive: a camera producing images of resolution $H \times W$ would require $(HW)^2$ evaluations, as each 3D point from $\mathcal{P}_t$ can be associated with any pixel from $I_c$. Additionally, the *argmax* operation
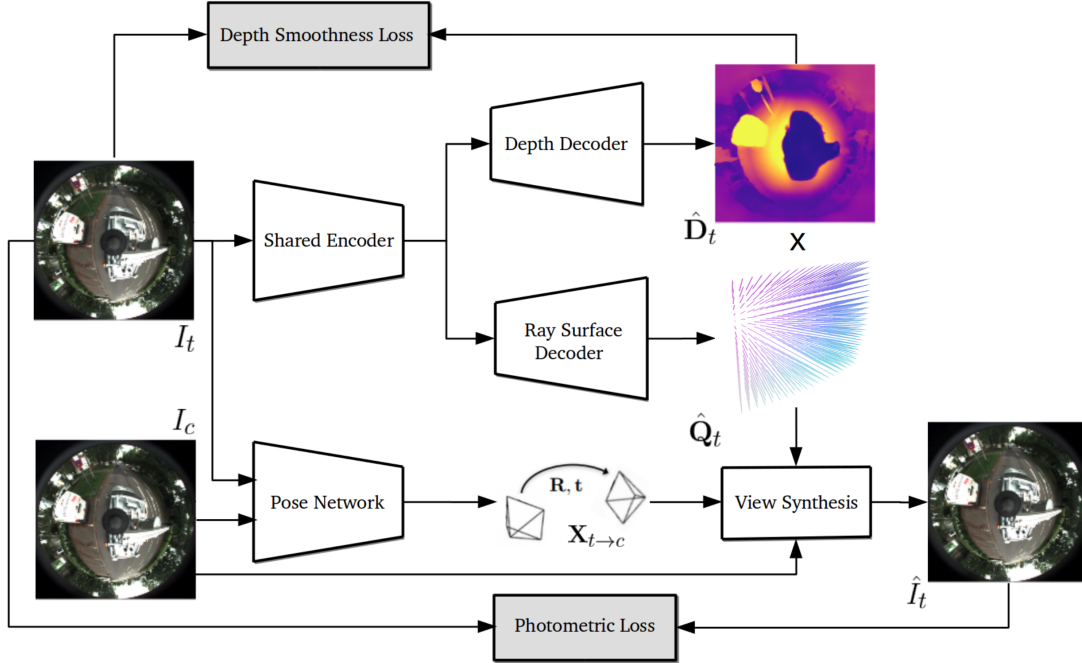
Figure 4.3: **Proposed self-supervised monocular depth, pose, and ray surface estimation framework**. Both depth and ray surface decoders share the same encoder backbone, and by combining the predicted depth map $\hat{D}_t$ with the predicted ray surface $\hat{Q}_t$, we are able to perform the view synthesis necessary for self-supervised learning.

is non-differentiable, which precludes its use in an end-to-end learning-based setting. We describe solutions to each of these issues below, that in conjunction enable the simultaneous learning of depth, pose and our proposed neural ray surface in a fully self-supervised monocular setting.

**Softmax Approximation**    To project the 3D points $\mathcal{P}_t$ onto context image $I_c$, we need to find for each $\mathbf{P}_j \in \mathcal{P}_t$ the corresponding pixel $\mathbf{p}_i \in I_c$ with surface ray $\hat{\mathbf{Q}}_i$ closest to the direction $\mathbf{r}_{c \to j} = \mathbf{P}_j$. Taking the dot product of each direction $\mathbf{r}_{c \to j}$ with each ray vector $\hat{\mathbf{Q}}_i$, we obtain a $(H \times W)^2$ tensor $\mathbf{M}$ where each coefficient $\mathbf{M}_{ij} = \langle \hat{\mathbf{Q}}_i, \mathbf{r}_{c \to j} \rangle = \mathbf{M}(\mathbf{p}_i, \mathbf{P}_j)$ represents the similarity between $\hat{\mathbf{Q}}_i$ and $\mathbf{r}_{c \to j}$. With this notation, projection for our proposed neural ray surface is given by selecting the $i^*$ index for each $\mathbf{P}_j$ with:

$$i^* = \arg\max_i \mathbf{M}(\mathbf{p}_i, \mathbf{P}_j) \tag{4.3.3}$$

To make this projection operation differentiable, we substitute *argmax* with a *softmax* with temperature $\tau$, thus obtaining a new tensor $\tilde{\mathbf{M}}$ defined as:

$$\tilde{\mathbf{M}}(\mathbf{p}_i, \mathbf{P}_j) = \frac{\exp(\mathbf{M}(\mathbf{p}_i, \mathbf{P}_j)/\tau)}{(\sum_i \exp(\mathbf{M}(\mathbf{p}_i, \mathbf{P}_j)/\tau))} \tag{4.3.4}$$

We anneal the temperature over time during training, so that the tensor approaches approximately one-hot for each pixel. We obtain the 2D-3D association used for projection by multiplying with a vector of pixel indices. Thus, projection can now be implemented in a fully differentiable way using STNs [85].

**Residual Ray Surface Template**    In the structure-from-motion setting, learning a randomly initialized ray surface is similar to learning 3D scene flow [132], which is a challenging problem when no calibration is available, particularly when considering self-supervision [133, 134]. To avoid this random initialization, we can instead learn a *residual* ray surface $\hat{\mathbf{Q}}_r$, that is added to a fixed ray surface template $\mathbf{Q}_0$ to produce $\hat{\mathbf{Q}} = \mathbf{Q}_0 + \lambda_r \hat{\mathbf{Q}}_r$. The introduction of this template allows the injection of geometric priors into the learning framework, since if some form of camera calibration is known – even if only an approximation – we can generate its corresponding ray surface, and use it as a starting point for further refinement using the learned ray surface residual. If no such information is available, we initialize a pinhole template based on approximate "default" calibration parameters, unprojecting a plane at a fixed distance and normalizing its surface.

For stability, we start training only with the template $\mathbf{Q}_0$ and gradually introduce the residual $\hat{\mathbf{Q}}_r$, by increasing the value of $\lambda_r$. We find that this *pinhole prior* significantly improves training stability and convergence speed even in a decidedly non-pinhole setting (i.e., catadioptric cameras). Predicting ray surface residuals on a per-frame basis allows for training on multiple datasets (with images obtained from different cameras) as well as adapting a pre-trained model to a new dataset.

Additionally, there are settings where frame-to-frame variability is expected even with a single camera (e.g. underwater imaging in a turbid water interface, rain droplets on a lens) but per-frame prediction may introduce unwanted frame-to-frame variability in settings where we would expect

a stable ray surface (i.e. all images come from the same camera). In the experiments section we evaluate the stability of ray surface predictions for a converged KITTI model, and find minimal frame-to-frame variability.

**Patch-Based Data Association**    In the most general version of our proposed neural ray surface model, rays at each pixel are independent and can point in completely different directions.

Because of that, Equation 4.3.3 requires searching over the entire image, which quickly becomes computationally infeasible at training time even for lower resolution images, both in terms of speed and memory footprint. To alleviate such heavy requirements, we restrict the optimal projection search (Equation 4.3.4) to a small $h \times w$ grid in the context image $I_c$ surrounding the $(u, v)$ coordinates of the target pixel $\mathbf{p}_t$. The motivation is that, in most cases, camera motion will be small enough to produce correct associations within this neighborhood, especially when using the residual ray surface template described above. To further reduce memory requirements, the search is performed on the predicted ray surface at half-resolution, which is then upsampled using bilinear interpolation to produce pixel-wise estimates. At test-time none of these approximations are necessary, and we can predict a full-resolution ray surface directly from the input image.

## 4.4   Experiments

In this section we demonstrate that our proposed neural ray surface model can be trained without any architectural changes on datasets containing video sequences captured with a variety of different cameras, while still achieving competitive results with other methods that rely on pre-calibrated or learned pinhole models.

To that end, we evaluate our framework on the standard rectified KITTI benchmark, a fisheye dataset (Multi-FOV) for depth evaluation and a catadioptric dataset (OmniCam) for visual odometry evaluation. We then discuss our experiments on two challenging datasets, an internal dashboard camera sequence and an underwater cave sequence.
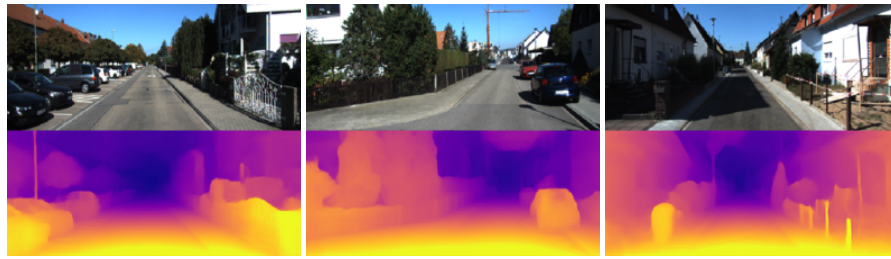
(a) Pinhole (KITTI)
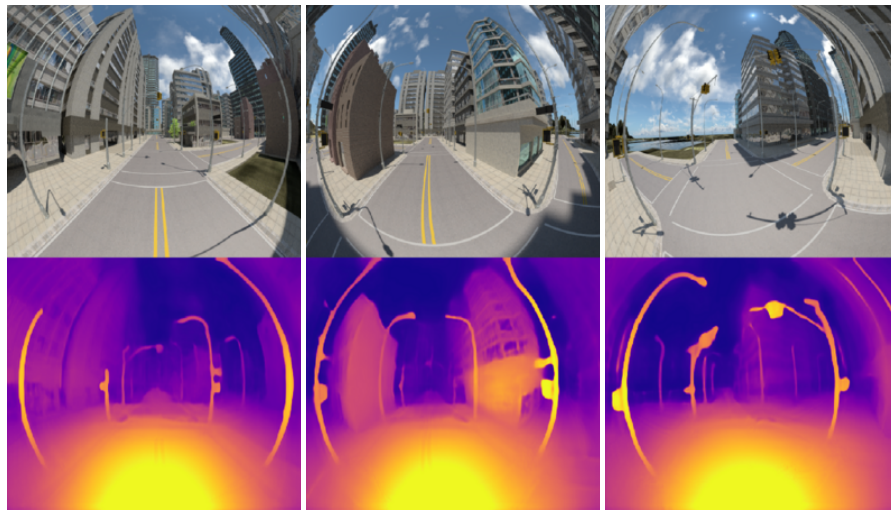


(b) Fisheye (Multi-FOV)
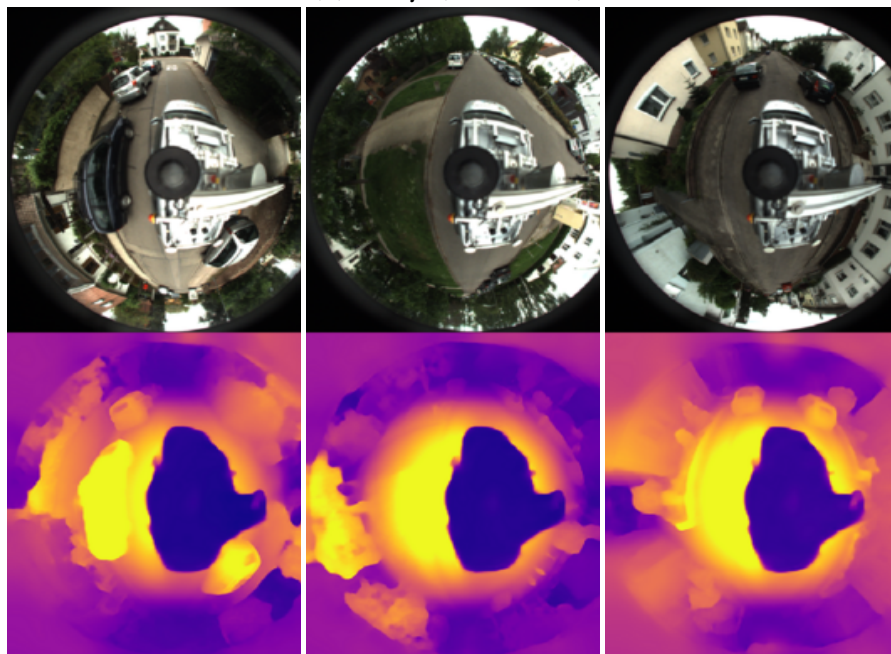


(c) Catadioptric (OmniCam)

Figure 4.4: **Qualitative depth estimation results for different camera geometries** using our proposed NRS model. Note that all these results were obtained using the same architecture and hyper-parameters (Figure 6.2a); the only modification are the sequences used for training and inference.

### 4.4.1 Datasets

We evaluate NRS on the KITTI depth and odometry benchmarks, and on the OmniCam odometry data. For more information on these datasets, please refer to Section 2.3.1.

- **KITTI** [53]. KITTI images are rectified, so we use this dataset to show that our proposed NRS model does not degrade results when the pinhole assumption is still valid.

- **Multi-FOV** [56]. To our knowledge, this dataset provides the only fisheye sequence with ground-truth depth maps, and it serves as a test of our model on fisheye cameras.

- **OmniCam** [135]. From OmniCam, we are particularly interested in the ground truth odometry sequence (given the lack of ground-truth depth maps).

### 4.4.2 Implementation Details

Our models were implemented using Pytorch [119] and trained across eight V100 GPUs. To highlight the flexibility of our proposed framework, all experiments used the same training hyper-parameters: Adam optimizer [82], with $\beta_1 = 0.9$ and $\beta_2 = 0.999$; batch size of 4 with learning rate of $2 \cdot 10^{-4}$ for 20 epochs; the previous $t - 1$ and subsequent $t + 1$ images are used as temporal context; color jittering and horizontal flipping as data augmentation; SSIM weight of $\alpha = 0.85$; and depth smoothness weight of $\lambda_d = 0.001$.

Furthermore, we used $41 \times 41$ patches for ray surface data association during projection. The ray surface template $\mathbf{Q}_0$ was initialized from a pinhole camera model with $f_x = c_x = W/2$ and $f_y = c_y = H/2$, increasing $\lambda_r$ from 0 to 1 over the course of 10 epochs. For the *depth network*, we experiment with two alternatives: a simpler *ResNet* architecture described by Godard et al. [13] and a more complex *PackNet* architecture described by Guizilini et al. [9]. For the *pose network*, we use the standard variant introduced by Zhou et al. [136] without the explainability mask.

### 4.4.3 Depth Evaluation

For depth estimation, we evaluate our framework on datasets containing pinhole (KITTI) and fisheye (Multi-FOV) cameras. Qualitative depth results for these datasets, and for OmniCam[1], are shown in Figure 4.4.

**KITTI** We evaluate our framework on a rectified, close-to-pinhole dataset as a sanity check on our model to answer the question: *is ray surface prediction comparable to predicting pinhole intrinsics when the projection model is known to be approximately pinhole?*

To this end, we perform the following ablation studies, as shown in Table 1: $PH - K$, where NRS is used with only a pinhole template initialized from known intrinsics; $RS - K$, where a ray surface network is learned with a pinhole template initialized from known intrinsics; and $RS - L$, where a ray surface network is learned with a pinhole template initialized from dummy intrinsics ($f_x = c_x = W/2$ and $f_y = c_y = H/2$).

The results in Table 1 suggest that, even though our framework is much more flexible, it still achieves competitive results with the recent "in the wild" self-supervision framework of [4]. In fact, our experiments showed small improvements when the ray surface model was used instead of the pinhole model, most likely due to small calibration and rectification errors that our neural framework is able to model accurately. Additionally, we hypothesize that, because the same encoder has to learn both depth and camera features, our framework benefits from a larger number of learned parameters, which is corroborated by the significant improvement obtained by the *PackNet* architecture. Like the pinhole prediction baseline in [4], our ray surface network operates on a per-frame basis. We measured the stability of the converged NRS-ResNet model by computing the coefficient of variation (a measure of dispersion) across the test set for KITTI, finding it to be less than $2.5\%$, showing that the predicted surface is very stable frame to frame.

---

[1]Projected depth maps were not available for a quantitative depth evaluation of OmniCam.

| Method | Camera | Abs Rel↓ | Sq Rel↓ | RMSE↓ | $\delta_{1.25}$ ↑ |
|---|---|---|---|---|---|
| Gordon [4] | $K$ | 0.129 | 0.982 | 5.230 | 0.840 |
| Gordon [4] | $L$ | 0.128 | 0.959 | 5.230 | **0.845** |
| NRS-ResNet | $PH - K$ | 0.137 | 0.969 | 5.377 | 0.821 |
| NRS-ResNet | $RS - K$ | 0.137 | 0.987 | 5.337 | 0.830 |
| NRS-ResNet | $RS - L$ | 0.134 | 0.952 | 5.263 | 0.832 |
| NRS-PackNet | $RS - L$ | **0.127** | **0.667** | **4.049** | 0.843 |

Table 4.1: **Quantitative depth evaluation for different methods on the KITTI dataset**, for distances up to 80m. In the *Camera* column, $PH$ indicates a pinhole template and $RS$ a ray surface network, with $K$ representing *known* parameters and $L$ *learned* parameters. We compare with another method that proposes the simultaneous learning of pinhole camera parameters [4].

**Multi-FOV**    Using the fisheye sequence in Multi-FOV, we compare NRS to the standard pinhole model for self-supervised depth estimation. This dataset deviates significantly from the pinhole assumption, and we can see in Table 4.2 that the our Neural Ray Surface-based model leads to a substantial improvement over the standard pinhole model: from $0.441$ absolute relative error down to $0.225$, a decrease of $51\%$. These results demonstrate that NRS is flexible enough to adapt to both pinhole (KITTI) and fisheye geometries (Multi-FOV) without any hyper-parameter changes.

| Model | Abs Rel↓ | Sq Rel↓ | RMSE↓ | $\delta_{1.25}$ ↑ |
|---|---|---|---|---|
| Pinhole | 0.441 | 4.211 | 7.352 | 0.336 |
| NRS-ResNet | **0.225** | **1.165** | **4.848** | **0.593** |

Table 4.2: **Quantitative depth evaluation on the Multi-FOV dataset**, for distances up to 80m using NRS-ResNet.

### 4.4.4   Visual Odometry

Real-world driving sequences for autonomous driving applications are captured with a wide variety of cameras, and a recent work [56] showed that large field-of-view cameras benefit traditional visual odometry methods, thanks to their ability to track more features across frames.

To this end, in addition to KITTI, we also evaluate our proposed framework on the OmniCam dataset, containing catadioptric images that give a full $360°$ field-of-view around the vehicle. To the best of our knowledge, NRS is the first self-supervised monocular method able to learn visual
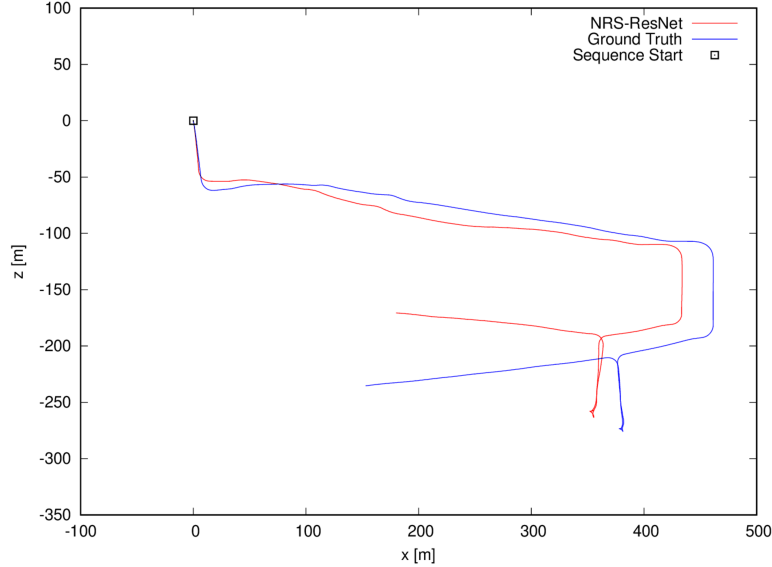
Figure 4.5: **Predicted trajectory for the first 2000 frames of the OmniCam dataset**, compared to the ground truth IMU/GPS trajectory, using NRS-ResNet.

odometry on catadioptric videos.

We plot the predicted trajectory from our pose network on the OmniCam dataset in Figure 4.5, comparing our predictions to the first 2,000 frames of OmniCam GPS/IMU ground truth. Even though the trajectory experiences global drift—it is worth noting that inference is performed on a two-frame basis, without loop-closure or any sort of bundle adjustment—it is remarkably accurate locally, especially given the fact that standard pinhole-based architectures completely diverge when applied to this dataset. Our NRS-ResNet model achieves an ATE of **0.035** on this dataset, while the same framework trained with a pinhole projection model produced significantly worse results, with an ATE of $0.408$.

For the KITTI dataset, we adopt the standard evaluation procedure, training on sequences 00-08 and testing on sequences 09 and 10, with the scale alignment procedure introduced in [136]. We report the 5-snippet ATE metric in Table 4.3, achieving comparable results to calibrated pinhole-based models, even though we do not require any prior knowledge of the camera system and do not perform any postprocessing or trajectory correction.

| | Seq. 09 | Seq. 10 |
|---|---|---|
| Zhou [12] | $0.0210 \pm 0.0170$ | $0.0200 \pm 0.015$ |
| Mahjourian [137] | $0.0130 \pm 0.0100$ | $0.0120 \pm 0.011$ |
| GeoNet [138] | $0.0120 \pm 0.0070$ | $0.0120 \pm 0.009$ |
| Godard [84] | $0.0230 \pm 0.0130$ | $0.0180 \pm 0.014$ |
| Struct2Depth [139] | $0.0110 \pm 0.0060$ | $0.0110 \pm 0.010$ |
| Gordon - known [4] | $\mathbf{0.009} \pm 0.0015$ | $0.008 \pm 0.011$ |
| Gordon - learned [4] | $0.0120 \pm 0.0016$ | $0.0100 \pm 0.010$ |
| Gordon - corrected [4] | $0.0100 \pm 0.0016$ | $\mathbf{0.007} \pm 0.009$ |
| NRS-ResNet | $0.0150 \pm 0.0301$ | $0.0103 \pm 0.0073$ |

Table 4.3: **Absolute trajectory error (ATE) on the KITTI dataset**, over five-frame snippets.

## 4.4.5 Pointcloud Reconstructions
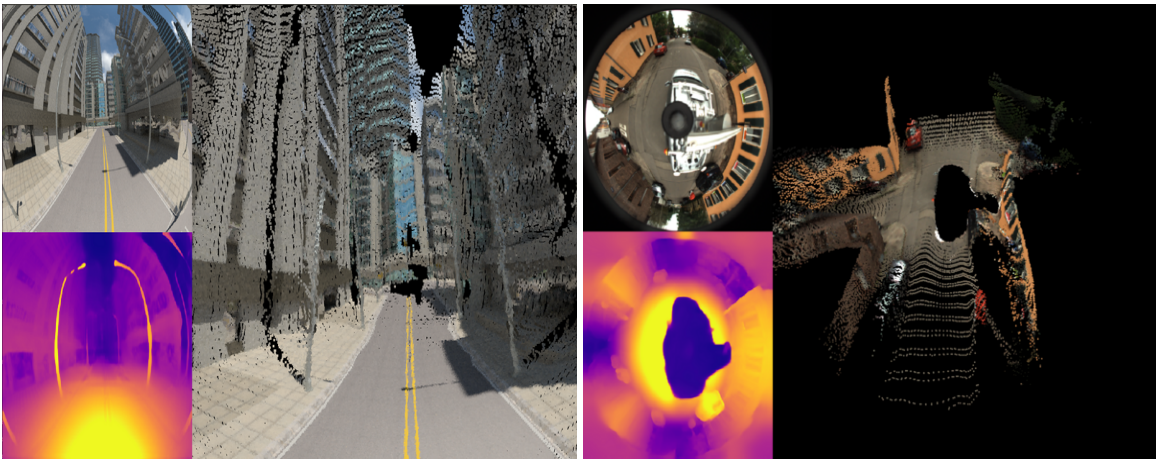


Figure 4.6: **Estimated pointclouds for the Multi-FOV (top) and OmniCam (bottom) datasets**. Our NRS framework enables the generation of geometrically accurate pointclouds from highly distorted images, in a self-supervised monocular setting.

Examples of reconstructed pointclouds using our proposed self-supervised NRS framework are shown in Figure 4.6, for the Multi-FOV and OmniCam datasets. These pointclouds are produced by multiplying the predicted depth map with the predicted ray surface. Note how, for Multi-FOV, we are able to correctly reconstruct straight 3D structures (i.e. buildings and poles) from a highly distorted image. Similarly, for OmniCam we can reconstruct the entire scene surrounding the vehicle, generating a 360° pointcloud from a single image in a fully self-supervised monocular setting.

### 4.4.6 Challenging Datasets

Datasets for self-supervised depth and ego-motion (mainly composed of street scenes for autonomous driving applications) are usually rectified to conform to the pinhole assumption. Thus, the use of camera models that conform to this assumption is generally adequate and able to produce accurate predictions. However, there are many settings in which the pinhole assumption is not appropriate, even when a near-pinhole camera is used.

The generality of NRS allows us to train in settings where a standard parametric model is not appropriate, without any changes in architecture. In this section, we describe in further detail experiments on two datasets mentioned in the paper—our internal DashCam dataset and a publicly available underwater caves dataset.

#### 4.4.6.1 Dashboard Camera

DashCam is an internal dataset containing video sequences taken with a fisheye camera behind a windshield. This capture setting is not modeled by standard parametric camera models, making it a good candidate for the application of generic camera models [130]. Furthermore, these images were heavily compressed to facilitate wireless transmission, which poses an additional challenge for the self-supervised photometric loss due to texture degradation.

There is no available ground-truth for this dataset, however the camera is calibrated and the distortion parameters are available. In Figure 4.9 we compare depth maps obtained from training our NRS model on the raw sequence to depth maps produced by a standard pinhole-based self-supervised model. We find that depth maps produced by the pinhole-based model on the rectified data are qualitatively significantly degraded compared to the NRS-based model trained on raw data.

We attribute this behavior to the rectification process, that degrades the information used to generate appearance-based features for monocular depth estimation. While rectification generally does not significantly affect results [53], the presence of compression artifacts and windshield distortions leads to significant degradation. Our NRS model, on the other hand, does not require

any rectification and therefore is able to use raw image information, leading to more accurate depth estimation even under such conditions.



| (a) Input image | (b) Depth map | (c) Pointcloud |
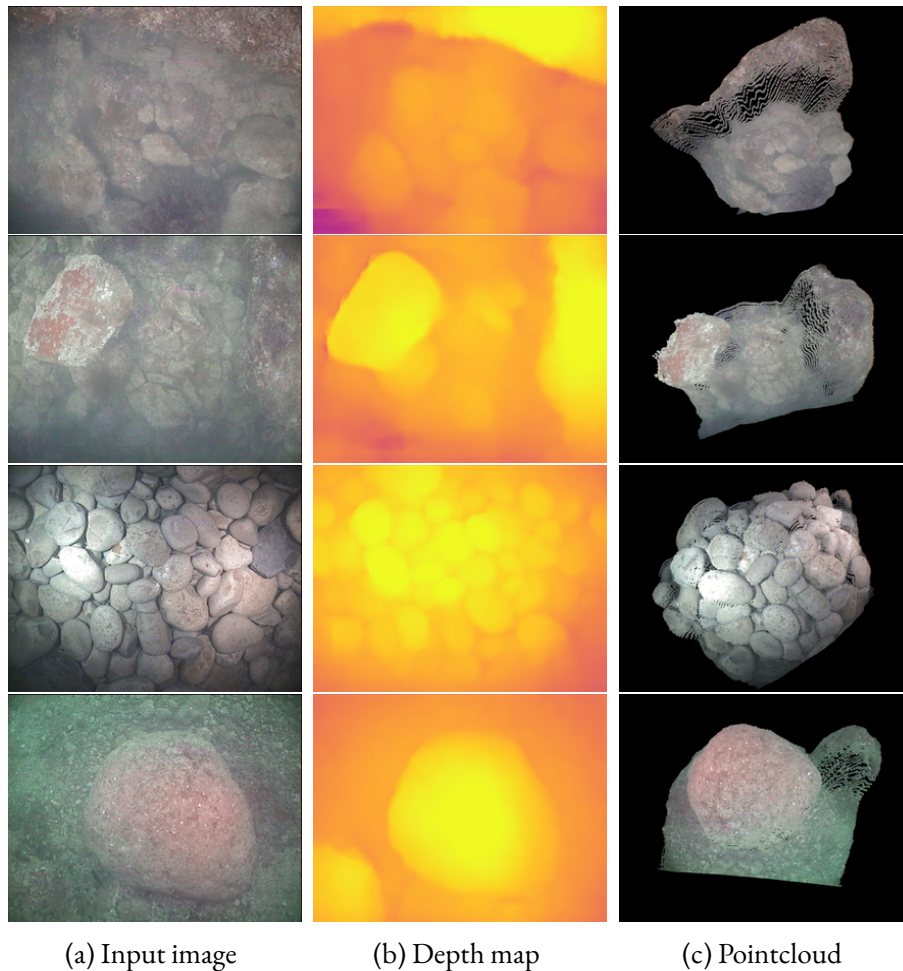
Figure 4.7: **Qualitative depth results on the *Underwater Caves* dataset**, using our proposed NRS model for self-supervised monocular depth and ego-motion estimation.

#### 4.4.6.2  Underwater Caves

Another challenging setting where the standard pinhole model is inappropriate is underwater vision, where refraction at the camera-water interface renders the standard pinhole model inaccurate [122]. This causes off-the-shelf structure-from-motion algorithms that rely on the pinhole assumption to produce inaccurate reconstructions [140]. We tested our proposed NRS model on the *Underwater Caves* dataset [141], a challenging visual odometry dataset taken in an underwater cave complex.
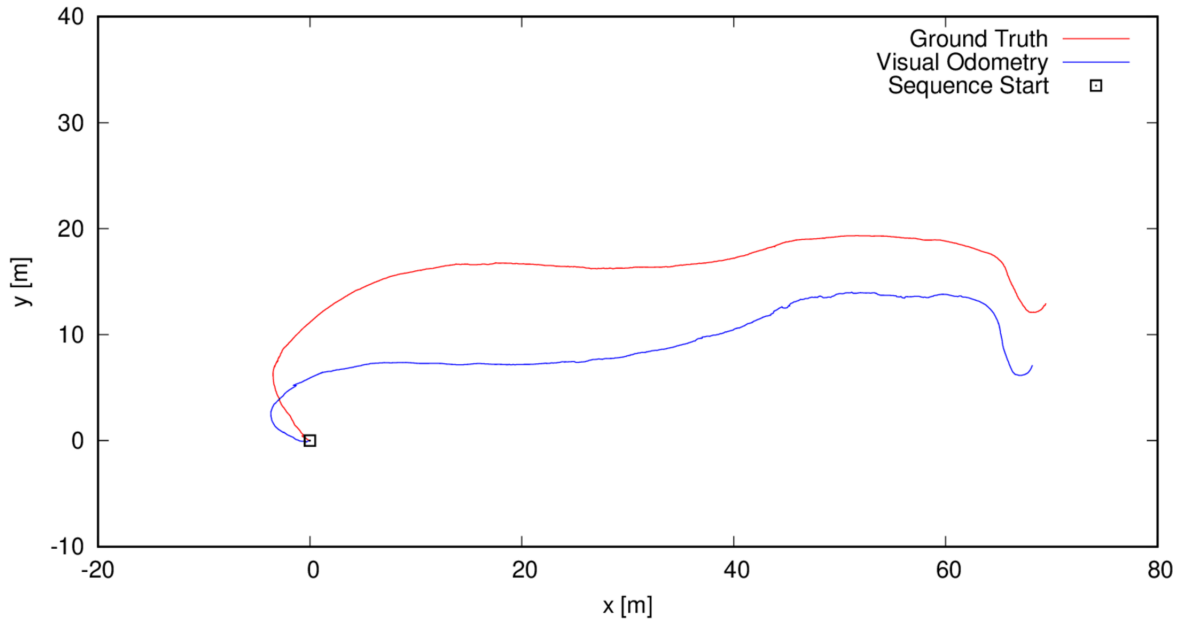
Figure 4.8: **Predicted trajectory on the *Underwater Caves* dataset** (last 2000 frames), obtained by accumulating predicted poses on a two-frame basis. Predicted trajectory scale obtained from ground truth scale.

The dataset is relatively limited in size (10k frames) and includes a variety of extremely challenging environments (low lamp illumination in a dark underwater cave, large levels of turbidity, etc.). Unsurprisingly, our baseline with a pinhole camera model [13] fails to learn meaningful depth and ego-motion predictors in this setting.

However, our NRS-based model is able to learn reasonable depth and odometry predictions on this data (see Figure 4.7), despite the fact that this is a challenging setting with many unstructured objects (rather than the manmade objects and surfaces common in datasets such as KITTI [53] and NYUv2 [65]). To our knowledge, this is the first demonstration of meaningful qualitative depth estimation for a dataset of natural objects. We also used the pose network to evaluate odometry predictions compared to the ground truth odometry, achieving an ATE of 0.0415 (see Figure 4.8). To our knowledge, this is the first demonstration of learning-based visual odometry in an underwater environment. Note that only raw videos were used at training time, without any ground truth or prior knowledge of camera model.

<div align="center">(a) Neural Ray Surfaces            (b) Pinhole</div>
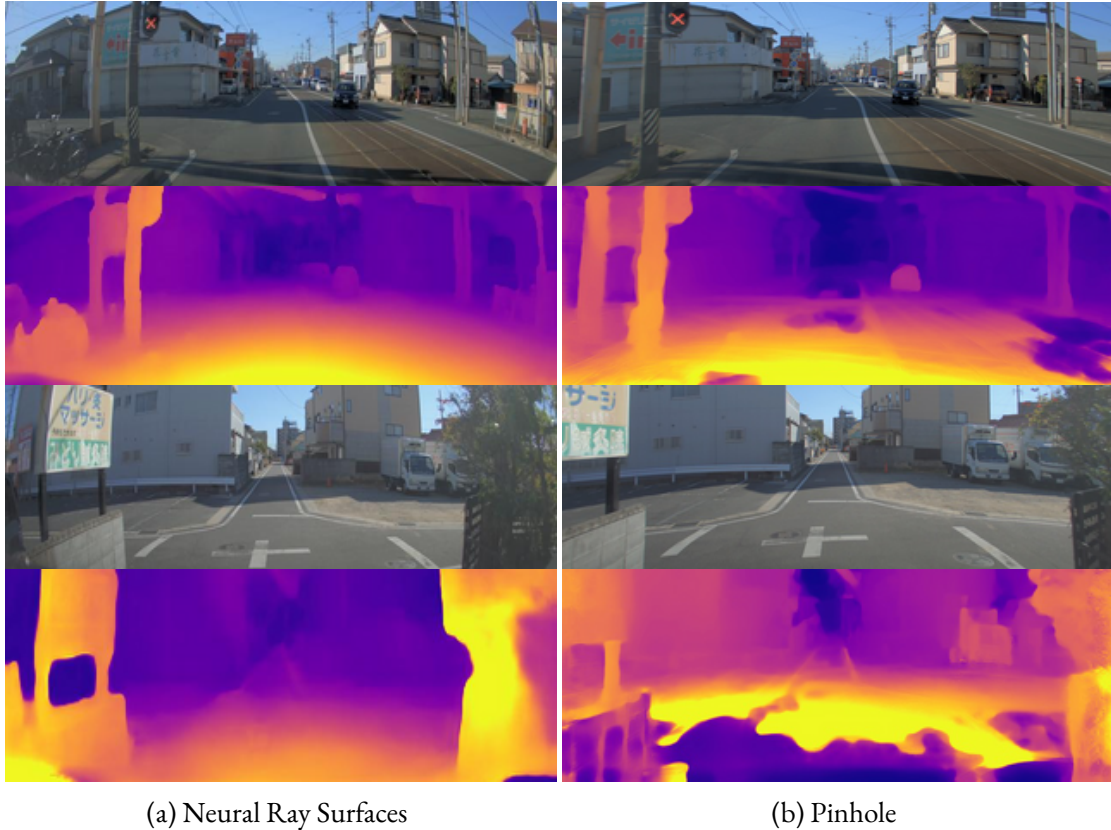
Figure 4.9: **Qualitative depth results on the DashCam dataset**. The left column shows raw RGB images and corresponding depth maps using our proposed Neural Ray Surfaces (NRS) model. The right column shows rectified RGB images and corresponding depth maps using a pinhole camera model. Note how NRS consistently leads to qualitatively better depth estimates, even though it uses as input raw unrectified images.

## 4.4.7   Network Architectures

In Table 4.4 we describe in details the networks used in our experimental evaluation. The depth network (*ResNet* and *PackNet*) receives a single RGB image as input and is composed of a shared encoder with two decoders: one for depth and one for the ray surface. The pose network (*PoseNet*) receives two concatenated RGB images as input, and produces as output the transformation between frames. Note that our proposed NRS model does not rely on any particular architecture, and others can be readily incorporated for potential improvements in speed and performance.

| | Layer Description | K | S | Out. Dim. |
|---|---|---|---|---|
| | **ResidualBlock (K, S)** | | | |
| #A | Conv2d→BN→ReLU | K | 1 | |
| #B | Conv2d→BN→ReLU | K | S | |
| | **UpsampleBlock (#skip)** | | | |
| #C | Conv2d→BN→ReLU→Upsample | 3 | 1 | |
| #D | Conv2d (#C ⊕ #skip)→BN→ReLU | 3 | 1 | |
| #0 | Input RGB image | - | - | 3×H×W |
| | **Encoder** | | | |
| #1 | Conv2d→BN→ReLU | 7 | 1 | 64×H×W |
| #2 | Max. Pooling | 3 | 2 | 64×H/2×W/2 |
| #3 | ResidualBlock (x2) | 3 | 2 | 64×H/4×W/4 |
| #4 | ResidualBlock (x2) | 3 | 2 | 128×H/8×W/8 |
| #5 | ResidualBlock (x2) | 3 | 2 | 256×H/16×W/16 |
| #6 | ResidualBlock (x2) | 3 | 2 | 512×H/32×W/32 |
| | **Depth Decoder** | | | |
| #7 | UpsampleBlock (#5) | 3 | 1 | 256×H/16×W/16 |
| #8 | UpsampleBlock (#4) | 3 | 1 | 128×H/8×W/8 |
| #9 | UpsampleBlock (#3) | 3 | 1 | 64×H/4×W/4 |
| #10 | UpsampleBlock (#2) | 3 | 1 | 32×H/2×W/2 |
| #11 | UpsampleBlock (#1) | 3 | 1 | 32×H×W |
| #12 | Conv2d→Sigmoid | 3 | 1 | 1×H×W |
| | **Ray Surface Decoder** | | | |
| #13 | UpsampleBlock (#5) | 3 | 1 | 256×H/16×W/16 |
| #14 | UpsampleBlock (#4) | 3 | 1 | 128×H/8×W/8 |
| #15 | UpsampleBlock (#3) | 3 | 1 | 64×H/4×W/4 |
| #16 | UpsampleBlock (#2) | 3 | 1 | 32×H/2×W/2 |
| #17 | UpsampleBlock (#1) | 3 | 1 | 32×H×W |
| #18 | Conv2d→Tanh | 3 | 1 | 3×H×W |

(a) Depth/Ray Surface Network (ResNet) [13].

| | Layer Description | K | S | Out. Dim. |
|---|---|---|---|---|
| | **ResidualBlock (K, S)** | | | |
| #A | Conv2d→GN→ELU | K | 1 | |
| #B | Conv2d→GN→ELU | K | 1 | |
| #C | Conv2d→GN→ELU→Dropout | K | S | |
| | **UpsampleBlock (#skip)** | | | |
| #D | Unpacking | 3 | 1 | |
| #E | Conv2d (#D ⊕ #skip)→GN→ELU | 3 | 1 | |
| #0 | Input RGB image | - | - | 3×H×W |
| | **Encoder** | | | |
| #1 | Conv2d→GN→ELU | 5 | 1 | 64×H×W |
| #2 | Conv2d→GN→ELU→Packing | 7 | 1 | 64×H×W |
| #3 | ResidualBlock (x2)→Packing | 3 | 1 | 64×H/4×W/4 |
| #4 | ResidualBlock (x2)→Packing | 3 | 1 | 128×H/8×W/8 |
| #5 | ResidualBlock (x3)→Packing | 3 | 1 | 256×H/16×W/16 |
| #6 | ResidualBlock (x3)→Packing | 3 | 1 | 512×H/32×W/32 |
| | **Depth Decoder** | | | |
| #7 | UpsampleBlock (#5) | 3 | 1 | 512×H/16×W/16 |
| #8 | UpsampleBlock (#4) | 3 | 1 | 256×H/8×W/8 |
| #9 | UpsampleBlock (#3) | 3 | 1 | 128×H/4×W/4 |
| #10 | UpsampleBlock (#2) | 3 | 1 | 64×H/2×W/2 |
| #11 | UpsampleBlock (#1) | 3 | 1 | 64×H×W |
| #12 | Conv2d→Sigmoid | 3 | 1 | 1×H×W |
| | **Ray Surface Decoder** | | | |
| #13 | UpsampleBlock (#5) | 3 | 1 | 512×H/16×W/16 |
| #14 | UpsampleBlock (#4) | 3 | 1 | 256×H/8×W/8 |
| #15 | UpsampleBlock (#3) | 3 | 1 | 128×H/4×W/4 |
| #16 | UpsampleBlock (#2) | 3 | 1 | 64×H/2×W/2 |
| #17 | UpsampleBlock (#1) | 3 | 1 | 64×H×W |
| #18 | Conv2d→Tanh | 3 | 1 | 3×H×W |

(b) Depth/Ray Surface Network (PackNet) [9].

| | Layer Description | K | S | Out. Dim. |
|---|---|---|---|---|
| #0 | Input 2 RGB images | - | - | 6×H×W |
| #1 | Conv2d→GN→ReLU | 3 | 2 | 16×H/2×W/2 |
| #2 | Conv2d→GN→ReLU | 3 | 2 | 32×H/4×W/4 |
| #3 | Conv2d→GN→ReLU | 3 | 2 | 64×H/8×W/8 |
| #4 | Conv2d→GN→ReLU | 3 | 2 | 128×H/16×W/16 |
| #5 | Conv2d→GN→ReLU | 3 | 2 | 256×H/32×W/32 |
| #6 | Conv2d→GN→ReLU | 3 | 2 | 256×H/64×W/64 |
| #7 | Conv2d→GN→ReLU | 3 | 2 | 256×H/128×W/128 |
| #8 | Conv2d | 1 | 1 | 6×H/128×W/128 |
| #9 | Global Pooling | - | - | 6 |

(c) Pose Network [136].

Table 4.4: **Neural network architectures used in our proposed NRS framework**, for the joint self-supervised learning of depth, pose and ray surfaces from monocular images. The depth network outputs $1 \times H \times W$ tensors with predicted inverse depth values, that are scaled between the minimum and maximum depth ranges. The ray surface network outputs $3 \times H \times W$ tensors, that are normalized to produce unitary vectors. The pose network outputs a 6-dimensional vector, representing $(x, y, z)$ translation and $(roll, pitch, yaw)$ Euler angles. *BN* stands for Batch Normalization [5], *GN* for Group Normalization [6], *Dropout* is described in [7], *Upsample* doubles spatial dimensions using bilinear interpolation, *ReLU* are Rectified Linear Units and *ELU* are Exponential Linear Units [8]. The symbol ⊕ indicates feature concatenation.

## 4.5 Discussion

Our experiments demonstrate that NRS achieves comparable results to the standard pinhole-model based architectures on near-pinhole data, while also enabling for the first time self-supervised depth and pose learning on challenging "in the wild" non-pinhole datasets (such as the catadioptric OmniCam dataset).



(a) Pinhole (KITTI)                    (b) Catadioptric (OmniCam)
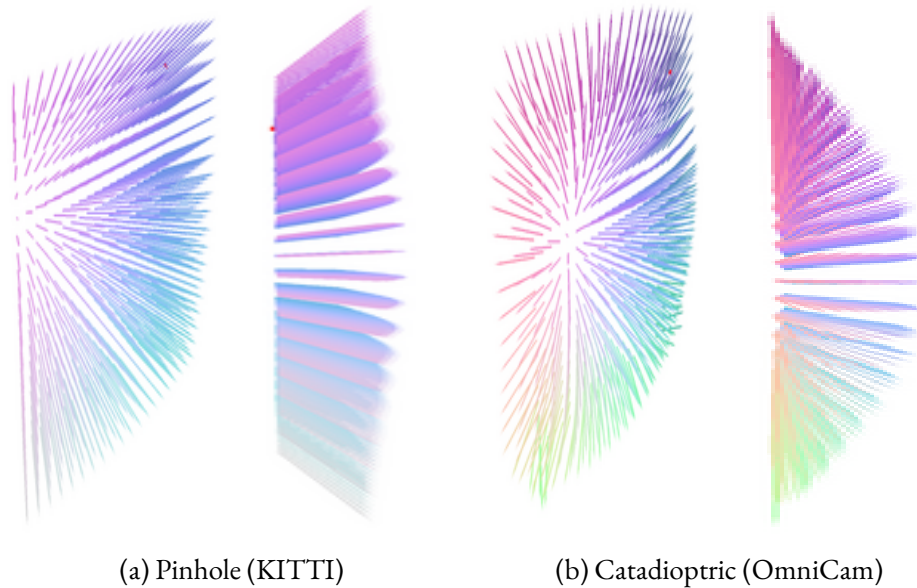
Figure 4.10: **Learned KITTI and OmniCam ray surfaces**, visualized as unitary 3D vectors for sub-sampled pixels (perspective and side view). Rays are colored by their directions for clarity. Note that NRS is able to adjust the ray surface on a per-pixel level in order to learn a projection model for two very different camera geometries.

OmniCam is particularly challenging because catadioptric image formation is substantially different from the pinhole projection model. Figure 4.10 visualizes the learned KITTI pinhole ray surface compared to the learned OmniCam catadioptric ray surface—both learned with the same architecture. The flexibility of NRS allows per-pixel updates to the pinhole template, learning ray surfaces that facilitate depth and ego-motion estimation for very different ray geometries and fields of view.

We also tested the ability of NRS to model ray geometries in two other challenging settings—an internal dataset consisting of driving sequences taken by a dashboard camera behind a windshield,

and a publicly-available sequence from an underwater cave environment [141]. In both of these settings, refraction (for the former, caused by the curved windshield, and for the latter, the water-camera interface) renders the standard parametric pinhole camera model inappropriate. In fact, we find that a standard pinhole-based self-supervised model trained on a rectified variant of these datasets fails to produce meaningful predictions, while NRS manages to predict reasonable depth and pose estimates without any changes to its original architecture (an example of depth prediction on the underwater dataset can be found in Figure 6.1).

## 4.6  Conclusion

We introduce Neural Ray Surfaces (NRS), a novel self-supervised learning framework capable of jointly estimating depth, pose, and per-pixel ray surface vectors in an end-to-end differentiable way. Our method can be trained on raw unlabeled videos captured from a wide variety of camera geometries without any calibration or architectural modification, thus broadening the use of self-supervised learning in the wild. We experimentally show on three different datasets that our methodology can tackle visual odometry and depth estimation on pinhole, fisheye, and catadioptric cameras without any architecture modifications.

# Chapter 5

# Learning Multi-Camera Monodepth

I am always surprised when I see several cameras, a gaggle on lenses, filters, meters, et cetera, rattling around in a soft bag with a complement of refuse and dust. Sometimes the professional is the worst offender!

Ansel Adams.

Self-supervised depth estimation architectures are largely restricted to the monocular or rectified stereo setting, capturing only a small fraction of the scene around the vehicle. Thus far, we have investigated how to expand the class of cameras that can be used for self-supervised learning, but we have maintained the restriction of *monocular* cameras for training. In this chapter, we extend monocular self-supervised depth and ego-motion estimation to large-baseline multi-camera rigs. Using generalized spatio-temporal contexts, pose consistency constraints, and carefully-designed photometric loss masking, we learn a single network generating dense, consistent, and scale-aware point clouds that cover the same full surround 360° field of view as a typical LiDAR scanner. We also propose a new scale-consistent evaluation metric more suitable to multi-camera settings. Experiments on two challenging benchmarks illustrate the benefits of our approach over strong

baselines.

## 5.1 Introduction

As we have seen in Chapters 3 and 4, self-supervision forms an integral part of modern state-of-the-art depth estimation architectures. Though recently released datasets contain multi-camera data that cover the same full 360° field of view as LiDAR [61, 9], research has focused on forward-facing cameras or stereo pairs. In this chapter, we extend self-supervised depth and ego-motion learning to *the general multi-camera setting*, where cameras can have different intrinsics and minimally overlapping regions, as required to minimize the number of cameras on the platform while providing full 360° coverage. We describe why stereo-based learning techniques do not apply in this setting, and show that batching cameras independently does not effectively leverage all information available in a multi-camera dataset.

We propose instead to leverage *cross-camera temporal contexts* via *spatio-temporal photometric constraints* to increase the amount of overlap between cameras thanks to the system's ego-motion. By exploiting known extrinsics between cameras, and enforcing *pose consistency constraints* to ensure all cameras follow the same rigid body motion, we are able to learn *scale-aware models* without any ground-truth depth or ego-motion labels. Furthermore, our multi-camera constraints enable the prediction of consistent point clouds across images, as reflected in our proposed *shared median-scaling* evaluation protocol. Finally, we find that masking out non-overlapping and self-occluded areas during photometric loss calculation has a drastic impact on performance.

In summary, our contributions are as follows:

- We demonstrate, for the first time, self-supervised learning of scale-aware and consistent depth networks in wide-baseline multi-camera settings, which we refer to as **Full Surround Monodepth (FSM)**.

- We introduce key techniques to extend self-supervised depth and ego-motion learning to wide-
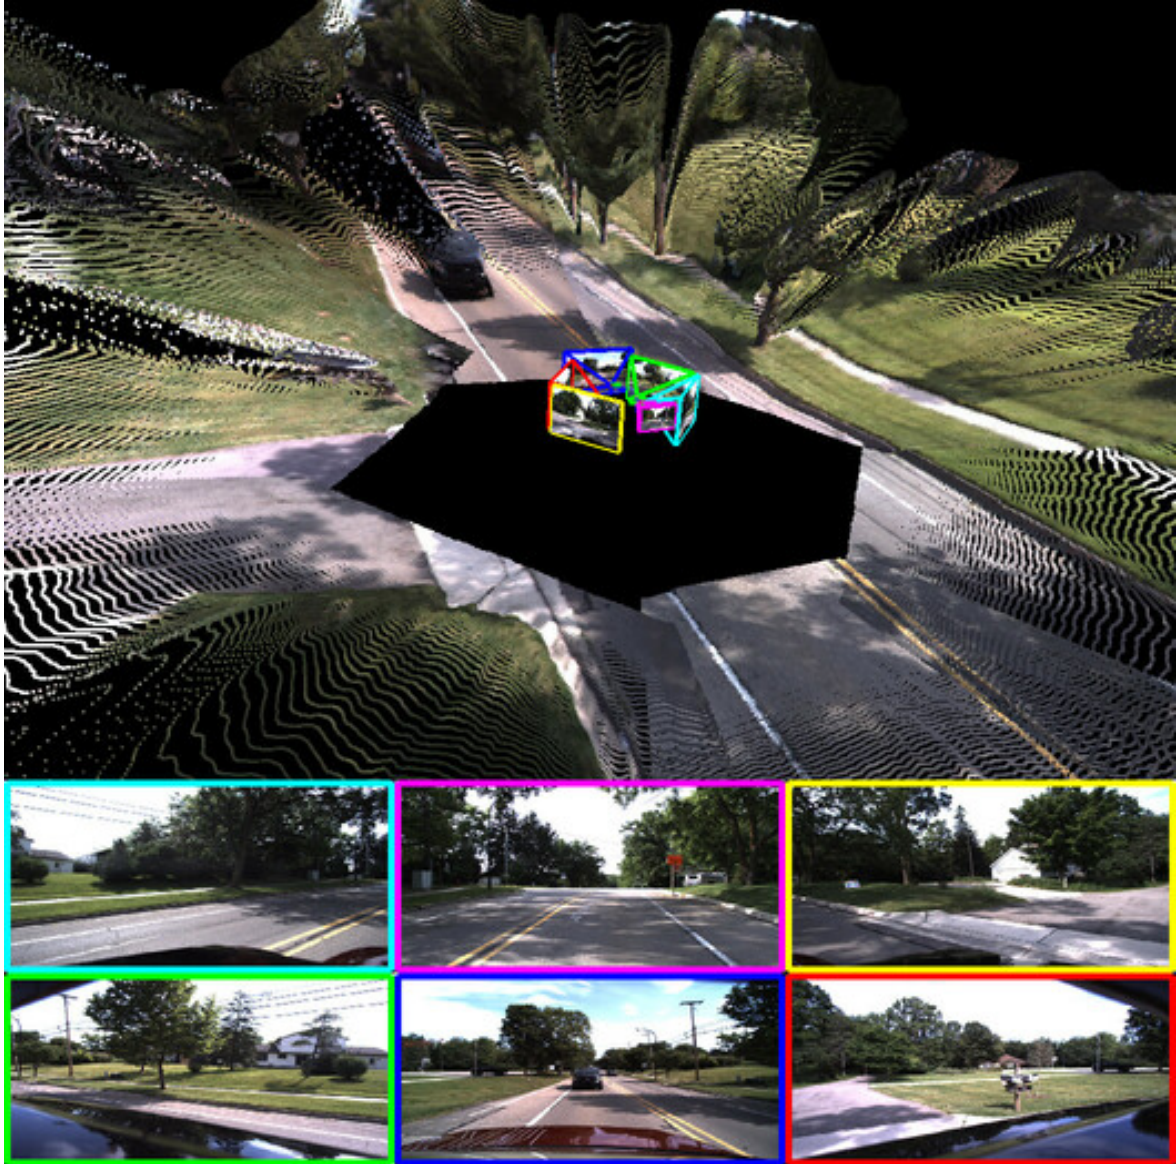
Figure 5.1: **Consistent scale-aware Full Surround Monodepth (FSM) pointcloud** from multiple cameras.

baseline multi-camera systems: multi-camera spatio-temporal contexts and pose consistency constraints, as well as study the impact of non-overlapping and self-occlusion photometric masking in this novel setting.

- We ablate and show the benefits of our proposed approach on two publicly available multi-camera datasets: *DDAD* [9] and *nuScenes* [61], achieving *state-of-the art results by a wide margin.*

78

## 5.2   Related Work

**Learning with Stereo Supervision.** Depth estimation from a rectified stereo pair is a classical task in computer vision [142, 41]. In this setting, the 2D matching problem is greatly simplified to a 1D disparity search. In recent years, supervised stereo depth estimation methods [143], as well as self-supervised techniques [84], have emerged as competitive learning-based approaches to this task. Self-supervised methods take advantage of rectified stereo training data with large overlap to train a disparity estimation network. Our proposed method is intended for multi-camera configurations with very large baselines (and thus minimal image overlap) where stereo rectification, and by extension disparity estimation, is not feasible.

**Omnidirectional Depth Estimation.** A popular approach to 360° depth estimation is through equirectangular or omnidirectional images [144, 145, 146]. These methods operate on panoramic images to estimate depth, either monocular or through stereo [147]. For robotics tasks, these images suffer from major disadvantages as a visual representation: (1) annotated datasets generally consist of perspective images, making transfer difficult; (2) specialized architectures are necessary; and (3) network training is limited by GPU memory, so resolution must be sacrificed to train using images with such a large field of view.

In Chapter 4, we produce 360° point clouds using (monocular) catadioptric omnidirectional images. However, the resolution of catadioptric images drops dramatically at range, while our proposed approach in this chapter generates much higher resolution pointclouds using systems of perspective cameras. Recently, a multi-camera training architecture has been proposed for fisheye cameras [148]; however, this work does not include cross-camera geometric consistency constraints and the training data is not available for comparison.

**Deep Multi-view Stereo.** Our multi-camera setting is related to the multi-view stereo (MVS) learning literature, which are generally supervised approaches where learned matching allows a network to predict matching cost-volumes [149, 150]. Khot *et al.* [151] relax the supervision requirements and propose a self-supervised MVS architecture, taking insights from self-supervised

monocular depth estimation and using a photometric loss. However, their proposed setting assumes a large collection of images surrounding a single object with known relative pose and large overlap for cost volume computation, and is thus very different from our setting—our architecture is designed to work with image sequences from any location and with arbitrarily small overlapping between cameras.

## 5.3  Methodology

Please refer to Section 2.3.4 for an introduction to single camera monocular self-supervised depth and ego-motion learning. In this section, we extend the description to our multi-camera setting and detail our three technical contributions.

### 5.3.1  Multi-Camera Spatio-Temporal Contexts

Multi-camera approaches to self-supervised depth and ego-motion are currently restricted to the stereo setting with rectified images that enable predicting disparities [84], which are then converted to depth through a known baseline. Although methods have been proposed that combine stereo and monocular self-supervision [88, 152], directly regressing depth also from stereo pairs, these still assume the availability of highly-overlapping images, from datasets such as KITTI [46]. Our proposed approach differs from the stereo setting in the sense that it *does not require stereo-rectified or highly-overlapping images*, but rather is capable of exploiting small overlaps (as low as $10\%$) between cameras with arbitrary locations as a way to both *improve individual camera performance* and generate *scale-aware* models from known extrinsics. Multi-camera rigs with such low overlap are common, e.g., in autonomous driving as a cost-effective solution to 360° vision [61, 9].

Let $C_i$ and $C_j$ be two cameras with extrinsics $\mathbf{X}_i$ and $\mathbf{X}_j$, and intrinsics $\mathbf{K}_i$ and $\mathbf{K}_j$. Denoting the relative extrinsics as $\mathbf{X}_{i \to j}$ and abbreviating $\phi_i(\mathbf{p}, \hat{d}) = \phi(\mathbf{p}, \hat{d}, \mathbf{K}_i)$ and $\pi_i(\mathbf{P}) = \pi(\mathbf{P}, \mathbf{K}_i)$, we can use Equation 2.3.1 to warp images from these two cameras:

$$\hat{\mathbf{p}}_i = \pi_j\big(\mathbf{R}_{i \to j}\phi_i(\mathbf{p}_i, \hat{d}_i) + \mathbf{t}_{i \to j}\big) \qquad\qquad (5.3.1)$$

Note that the above equation is purely *spatial*, since it warps images between different cameras taken at the same timestep. Conversely, Equation 2.3.1 is purely *temporal*, since it is only concerned with warping images from the same camera taken at different timesteps.

Therefore, for any given camera $C_i$ at a timestep $t$, a context image can be either temporal (i.e., from adjacent frames $t-1$ and $t+1$) or spatial (i.e., from any camera $j$ that overlaps with $i$). This allows us to further generalize the concept of "context image" in self-supervised learning to also include temporal contexts from other overlapping cameras. This is done by warping images between different cameras taken at different timesteps using a composition of known extrinsics with predicted ego-motion:

$$\hat{\mathbf{p}}_i^t = \pi_j\big(\mathbf{R}_{i \to j}\big(\hat{\mathbf{R}}_j^{t \to c}\phi_j(\mathbf{p}_j^t, \hat{d}_j^t) + \hat{\mathbf{t}}_j^{t \to c}\big) + \mathbf{t}_{i \to j}\big) \qquad\qquad (5.3.2)$$

A diagram depicting such transformations can be found in Figure 5.2, and Figure 5.3 shows examples of warped images and corresponding photometric losses using the *DDAD* dataset [9]. Particularly, the fifth and sixth rows show examples of multi-camera photometric losses using purely spatial contexts (Equation 5.3.1) and our proposed spatio-temporal contexts (Equation 5.3.2). As we can see, *spatio-temporal contexts (STC) promote a larger overlap between cameras and smaller residual photometric loss*, due to occlusions and changes in brightness and viewpoint. This improved photometric loss leads to *better self-supervision* for depth and ego-motion learning in a multi-camera setting, as validated through our experiments.

### 5.3.2   Multi-Camera Pose Consistency Constraints

Beyond cross-camera constraints due to image overlap, there are also natural pose constraints due to the fact that all cameras are rigidly attached to a single vehicle (i.e., relative camera extrinsics are constant and fixed). Specifically, the pose network is used to predict independent poses for
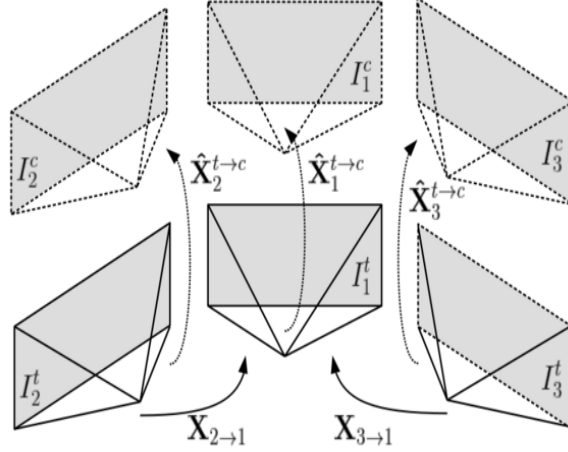
Figure 5.2: **Multi-camera spatio-temporal transformation matrices**. Solid cameras are *target* (current frames), and dotted cameras are *context* (adjacent frames). Spatial transformations ($\mathbf{X}_{i \to 1}$) are obtained from known extrinsics, and temporal transformations ($\hat{\mathbf{X}}_i^{t \to c}$) from the pose network.

each camera, even though they should correspond to the same transformation, just in a different coordinate frame. For a given camera $i$, the pose network predicts its transformation $\hat{\mathbf{X}}_i^{t \to t+1}$ from the current frame $t$ to the subsequent frame $t + 1$. In order to obtain predictions from different cameras that are in the same coordinate frame, we transform this prediction to the coordinate frame of a "canonical" camera $C_j$. We denote $\hat{\mathbf{X}}_i^{t \to t+1}$ in $C_j$ coordinates as $\tilde{\mathbf{X}}_i^{t \to t+1}$.

To convert a predicted transformation $\hat{\mathbf{X}}_i^{t \to t+1}$ from the coordinate frame of camera $C_i$ to camera $C_j$, we can use the extrinsics $\mathbf{X}_i$ and $\mathbf{X}_j$ to generate $\tilde{\mathbf{X}}_i^{t \to t+1}$ as follows:

$$\tilde{\mathbf{X}}_i^{t \to t+1} = \mathbf{X}_j^{-1} \mathbf{X}_i \hat{\mathbf{X}}_i^{t \to t+1} \mathbf{X}_i^{-1} \mathbf{X}_j \tag{5.3.3}$$

where $\tilde{\mathbf{X}}_i^{t \to t+1} = \begin{pmatrix} \tilde{\mathbf{R}}_i^{t \to t+1} & \tilde{\mathbf{t}}_i^{t \to t+1} \\ \mathbf{0} & \mathbf{1} \end{pmatrix}$. As a convention, we convert all predicted transformations to the coordinate frame of the front camera $C_1$. Once all predictions are in the same coordinate frame, we constrain the translation vectors $\mathbf{t}$ and rotation matrices $\mathbf{R}$ to be similar across all cameras.

**Translation.** We constrain all predicted translation vectors to be similar to the prediction for the *front camera*, which generally performs best across all experiments. Defining the predicted front
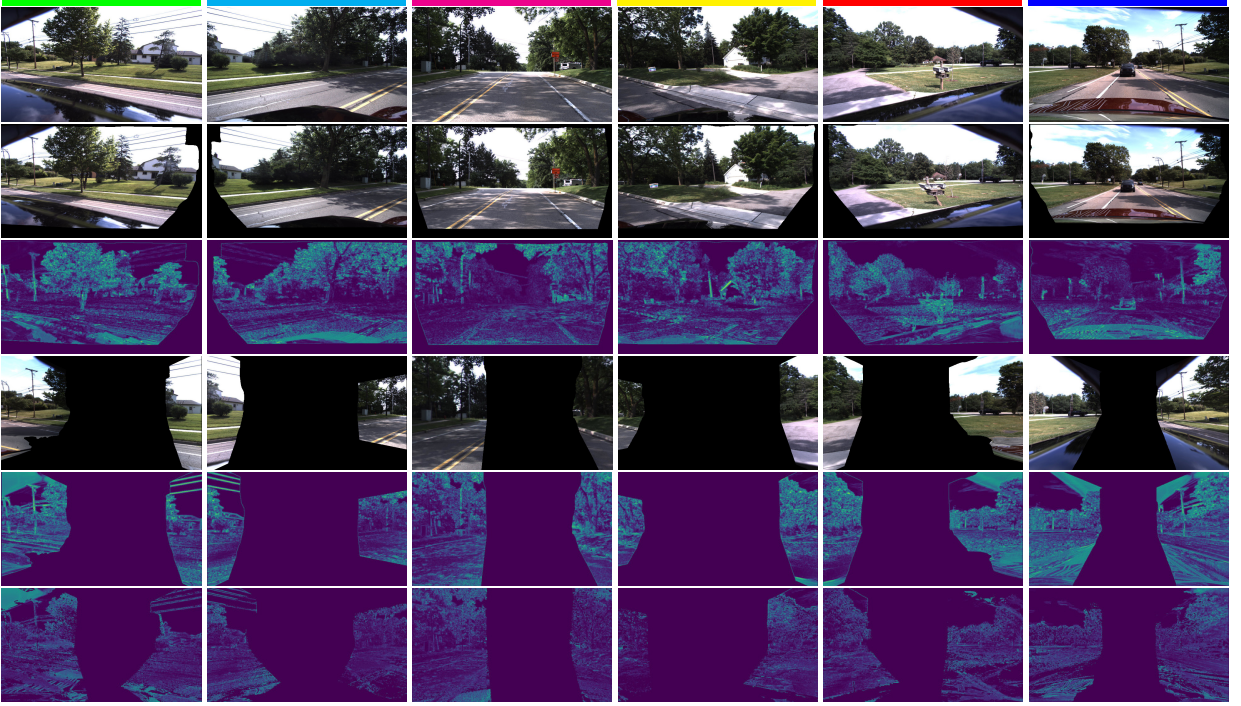
Figure 5.3: **Examples of spatial and temporal image warping** on the *DDAD* dataset (camera colors from Figure 6.1, clockwise). **First row:** Input RGB images. **Second and third rows:** Synthesized views from temporal contexts (Equation 2.3.1) and photometric losses. **Fourth and fifth rows:** Synthesized views from surrounding cameras (Equation 5.3.1), and photometric losses using only spatial contexts. **Sixth row** Photometric losses using our proposed spatio-temporal contexts. By also leveraging temporal contexts during cross-camera photometric warping, we are able to generate larger overlapping areas between images, as well as a smaller residual photometric error (darker colors) for optimization.

camera translation vector as $\hat{\mathbf{t}}_1^{t+1}$, for $N$ cameras the translation consistency loss is given by:

$$\mathbf{t}_{loss} = \sum_{j=2}^{N} \|\hat{\mathbf{t}}_1^{t+1} - \tilde{\mathbf{t}}_j^{t+1}\|^2 \tag{5.3.4}$$

**Rotation.** Similarly, we want to constrain other cameras to predict a rotation matrix similar to the front camera. To accomplish that, once the predictions are in the same coordinate frame we convert them to Euler angles $(\phi_i, \theta_i, \psi_i)$ and calculate the rotation consistency loss such that:

$$\mathbf{R}_{loss} = \sum_{j=2}^{N} \|\hat{\phi}_1 - \tilde{\phi}_j\|^2 + \|\hat{\theta}_1 - \tilde{\theta}_j\|^2 + \|\hat{\psi}_1 - \tilde{\psi}_j\|^2 \tag{5.3.5}$$

Similar to the trade-off between rotation losses and translation loss in the original PoseNet [30, 153], we trade off between the two constraints by defining $\mathcal{L}_{pcc} = \alpha_t \mathrm{t}_{loss} + \alpha_r \mathrm{R}_{loss}$, where $\alpha_t$ and $\alpha_r$ are weight coefficients.

### 5.3.3 The Importance of Masks

The photometric loss, as used for self-supervised monocular depth and ego-motion learning, has several assumptions that are not entirely justified in real-world scenarios. These include the static world assumption (violated by dynamic objects), brightness constancy (violated by luminosity changes), and dense overlap between frames (violated by large viewpoint changes). Although several works have been proposed to relax some of these assumptions [4], more often than not methods are developed to mask out those regions [88], to avoid spurious information from contaminating the final model.

In a multi-camera setting there are two scenarios that are particularly challenging for self-supervised depth and ego-motion learning: *non-overlapping areas*, due to large changes in viewpoint between cameras, and *self-occlusions*, due to camera positioning that results in the platform (i.e., ego-vehicle) partially covering the image. Here we describe how our proposed approach addresses each of these scenarios. The final masked photometric loss used during training takes the form:

$$\mathcal{L}_{\mathrm{mp}}(I_t, \hat{I}_t) = \mathcal{L}_p(I_t, \hat{I}_t) \odot \mathcal{M}_{\mathrm{no}} \odot \mathcal{M}_{\mathrm{so}} \tag{5.3.6}$$

where $\odot$ denotes element-wise multiplication, and $\mathcal{M}_{\mathrm{no}}$ and $\mathcal{M}_{\mathrm{so}}$ are binary masks respectively for non-overlapping and self-occluded areas, as described below.

**Non-Overlapping Areas**   We generate non-overlapping area masks by jointly warping with each context image a unitary tensor of the same spatial dimensions, using nearest-neighbor interpolation. The warped tensor is used to mask the photometric loss, thus avoiding gradient backpropagation in areas where there is no overlap between frames. Note that this unitary warping also uses network

predictions, and therefore is constantly updated at training time. This is similar to the motion masks described in Mahjourian et al. [137], however here we extend this concept to a spatio-temporal multi-camera setting. Figure 5.3 shows examples of spatial and temporal non-overlapping masks on the *DDAD* dataset, with a trained model. As expected, temporal contexts have a large amount of frame-to-frame overlap ($> 90\%$), even considering side-pointing cameras. Spatial context overlaps, on the other hand, are much smaller ($10-20\%$), due to radically different camera orientations.

**Self-Occlusions**    A common technique in self-supervised learning is the "auto-mask" procedure, which filters out pixels with a synthesized reprojection error higher than the original source image [88]. This mask particularly targets the "infinite depth problem", which occurs when scene elements move at the same speed as the vehicle, causing zero parallax and thus an infinitely-far away depth estimate for that region. However, this technique assumes brightness constancy, and the self-occlusions created by the robot (or car) body are often highly specular (Figure 5.4a), especially in the case of modern passenger vehicles.

Using a network with auto-masking enabled, specular self-occlusions create serious errors in the depth predictions, as shown in Figure 5.4c. We propose a simpler approach, consisting of creating manual masks for each camera (this needs only to be done once, assuming that the extrinsics remain constant). As shown in Figure 5.4d, and ablated in experiments, the introduction of these self-occlusion masks results in a substantial improvement in overall performance, to the point of enabling self-supervised depth and ego-motion learning under these conditions. Interestingly, self-occluded areas in the predicted depth maps are correctly "inpainted" to include the hidden ground plane. We posit that this is due to multi-camera training, in which a region unoccluded in one camera can be used to resolve self-occlusions in other viewpoints.

(a) Input RGB image.

(b) Self-occlusion mask.

(c) Without self-occlusion masking.

(d) With self-occlusion masking.
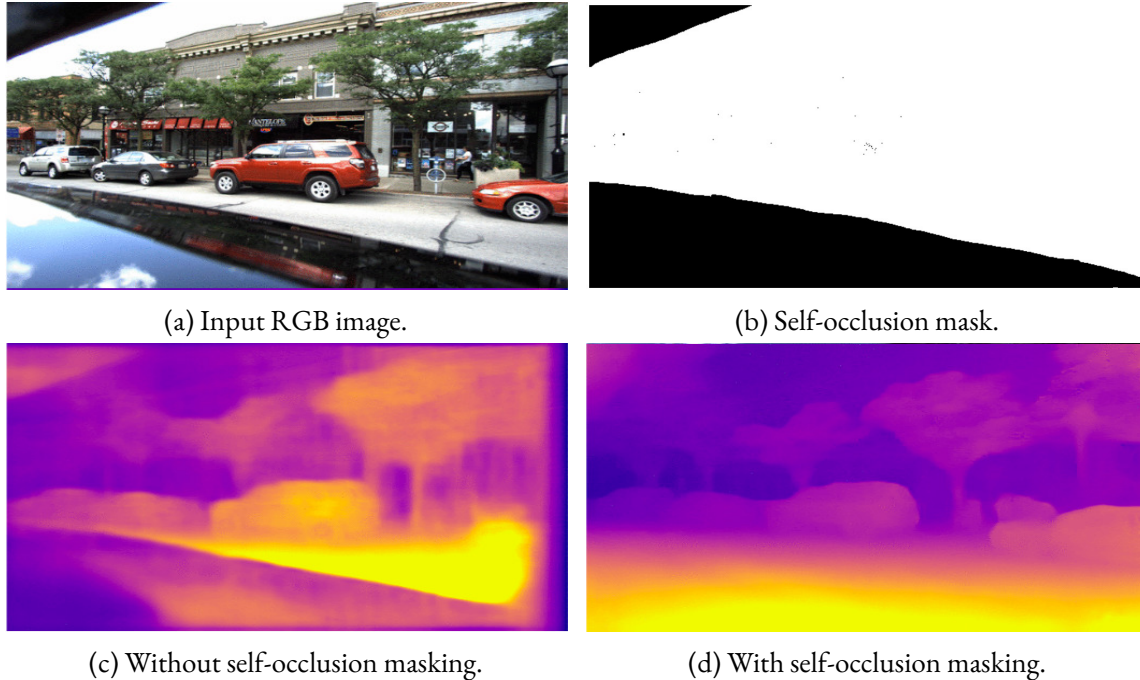
Figure 5.4: **Impact of self-occlusion masks on depth estimation** on the *DDAD* dataset. These masks remove self-occluded regions from the self-supervised photometric loss, enabling easier optimization (lower loss, cf. Figure 5.3) and better generalization (e.g., on the ground plane).

## 5.4 Experiments

### 5.4.1 Datasets

Traditionally, self-supervised depth and ego-motion learning uses monocular sequences [12, 88, 4, 9] or rectified stereo pairs [88] from forward-facing cameras in the KITTI [46] dataset. Recently, several datasets have been released with synchronized multi-camera sequences that cover the entire surrounding of the ego-vehicle [61, 9]. We focus on these datasets (DDAD and nuScenes) for our experiments, showing that our proposed approach, FSM, produces substantial improvements across all cameras. For more information on these datasets, please refer to Section 2.3.1.

- **KITTI** [46]. KITTI only contains forward-facing stereo pairs, so here it serves as a sanity check—we show that FSM accommodates the special case of high-overlapping rectified images, achieving competitive results with stereo methods.

- **DDAD** [9]. The DDAD dataset is the main focus of our experiments, since it contains six cameras with relatively small overlap and highly accurate dense ground-truth depth maps for evaluation. We show that, by jointly training FSM on all cameras, we considerably improve results and establish a new state of the art on this dataset by a large margin.

- **nuScenes** [61]. nuScenes is a challenging dataset for self-supervised depth estimation because of the relatively low resolution of the images, very small overlap between the cameras, high diversity of weather conditions and time of day, and unstructured environments. Thus, so far it has only been used as a dataset for depth evaluation [9]. We show that FSM is robust enough to overcome these challenges and substantially improve results relative to the baseline.

### 5.4.2 Training Protocol

Our models were implemented using PyTorch [119] and trained across eight V100 GPUs.[1] To highlight the flexibility of our proposed framework, all experiments used the same training hyper-parameters: Adam optimizer [82], with $\beta_1 = 0.9$ and $\beta_2 = 0.999$; batch size of 4 per GPU for single camera and 6 per GPU for multi-camera (all images from each sample in the same batch); learning rate of $2 \cdot 10^{-4}$ for 20 epochs; the previous $t - 1$ and subsequent $t + 1$ images are used as temporal context; color jittering and horizontal flipping as data augmentation; SSIM weight of $\alpha = 0.85$; and depth smoothness weight of $\lambda_d = 0.001$. We also used coefficients $\lambda_s = 0.1$ and $\lambda_t = 1.0$ to weight spatial and temporal losses respectively.

### 5.4.3 Multi-Camera Depth Evaluation Metrics

Our approach is *scale-aware*, due to its use of known extrinsics to generate metrically accurate predictions. In contrast, existing self-supervised monocular depth and ego-motion architectures learn up-to-scale models, resorting to *median-scaling* at test time in order to compare depth predictions against the (scaled) ground truth. Median scaling "borrows" the true scale of ground truth

---

[1] Training and inference code is available in the packnet repo.
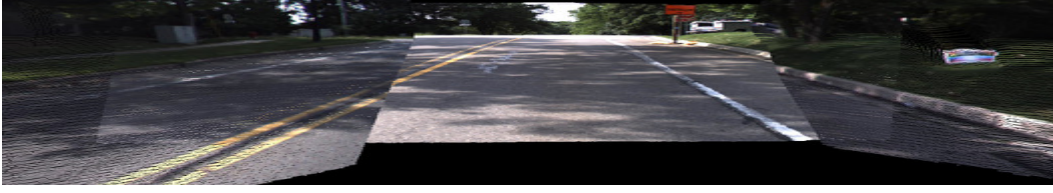
(a) Monocular photometric loss (Abs Rel 0.211/0.241)



(b) Spatio-temporal photometric loss (Abs Rel 0.201/0.207)

Figure 5.5: **Multi-camera pointcloud alignment** on DDAD using (a) the standard monocular photometric loss, and (b) our proposed spatio-temporal photometric constraints. We also report per-frame and shared median-scaling Abs Rel results (average of all cameras, see Table 5.3 for more details).

by multiplying each depth estimate by $\frac{\text{med}(D_{\text{gt}})}{\text{med}(D_{\text{pred}})}$, where med is the median operation and $D_{\text{gt}}$ and $D_{\text{pred}}$ are ground-truth and predicted depth maps. This scaling enables quantitative evaluation, but requires ground truth information at test time, limiting the utility of such methods in real-world applications.

Furthermore, median-scaling hides certain undesired behaviors, such as frame-to-frame variability [154]. This is exacerbated in the setting proposed in this chapter, where multiple images are used to generate a single, consistent pointcloud. If median-scaling is applied on a per-frame basis, the resulting depth maps will hide any scale discrepancies between camera predictions, which will not be reflected in the quantitative evaluation (Figure 5.5). Thus, instead of the standard median-scaling protocol, we propose to use a single scale factor $\gamma$ shared by all $N$ considered cameras defined as:

$$\gamma = \frac{\text{med}\left(\{D_{\text{gt}}^1, \cdots, D_{\text{gt}}^N\}\right)}{\text{med}\left(\{D_{\text{pred}}^1, \cdots, D_{\text{pred}}^N\}\right)} \tag{5.4.1}$$

This is similar to single median-scaling [88], in which the same factor is used to scale predictions for the entire dataset. In our setting, because multiple images are considered jointly, we instead produce a shared factor to scale all predictions at that timestep. This forces all predicted depth maps in the same timestep to have the same scale factor (and thus be relatively consistent), with any

| Method | Type | Abs Rel↓ | Sq Rel↓ | RMSE↓ | $\delta_{1.25}$ ↑ |
|---|---|---|---|---|---|
| UnDeepVO [155] | $S$ | 0.183 | 1.730 | 6.570 | - |
| Godard et al. [84] | $S$ | 0.148 | 1.344 | 5.927 | 0.803 |
| SuperDepth [156] | $S$ | 0.112 | 0.875 | 4.958 | 0.852 |
| Monodepth2† [88] | $M$ | 0.115 | 0.903 | 4.863 | 0.877 |
| Monodepth2 [88] | $S$ | 0.109 | 0.873 | 4.960 | 0.864 |
| Monodepth2 [88] | $M + S$ | **0.106** | 0.818 | 4.750 | **0.874** |
| **FSM** | $M + S$ | 0.108 | **0.737** | **4.615** | 0.872 |

Table 5.1: **Depth estimation results on the KITTI dataset,** relative to stereo (S) methods. Even though FSM (with monocular (M) and multi-camera (S) loss terms) relaxes several stereo assumptions, it remains competitive with published methods. The symbol † denotes per-frame median scaling.

| Method | Abs Rel↓ | Sq Rel↓ | RMSE↓ | $\delta_{1.25}$ ↑ |
|---|---|---|---|---|
| COLMAP (pseudo-depth) | 0.243 | 4.438 | 17.239 | 0.601 |
| Monodepth2† (R18) [88] | 0.213 | 4.975 | 18.051 | 0.761 |
| Monodepth2† (R50) [88] | 0.198 | 4.504 | 16.641 | 0.781 |
| PackNet† [9] | 0.162 | 3.917 | **13.452** | 0.823 |
| **FSM**† (w/o mask & spatial) | 0.184 | 4.049 | 17.109 | 0.735 |
| **FSM**† (w/o spatial) | 0.139 | 3.023 | 14.106 | 0.827 |
| **FSM** (w/ spatial) | 0.135 | 2.841 | 13.850 | 0.832 |
| **FSM** (w/ spatio-temporal) | **0.130** | **2.731** | 13.654 | **0.837** |

Table 5.2: **Quantitative depth evaluation of different methods on the DDAD [9] dataset,** for distances up to 200m on the forward-facing camera. The symbol † denotes per-frame median scaling.

deviation reflected in the calculated metrics. In practice, as our method is scale-aware, we report metrics both with (for comparison with baselines) and without median-scaling.

## 5.4.4   Stereo Methods

Though our proposed approach is intended for multi-camera rigs with arbitrarily small overlap, it can also be used without modification on stereo datasets, allowing us to learn depth directly instead of disparity. In fact, FSM can be seen as a generalization of stereo depth estimation, relaxing the assumption of fronto-parallel rectification. Nevertheless, in Table 5.1 we show that, despite not

| Method | Abs.Rel.↓ | | | | | | |
|---|---|---|---|---|---|---|---|
| | Front | F.Left | F.Right | B.Left | B.Right | Back | Avg. |
| Mono† - M | 0.184 | 0.366 | 0.448 | 0.417 | 0.426 | 0.438 | 0.380 |
| Mono† | 0.139 | 0.209 | 0.236 | 0.231 | 0.247 | 0.204 | 0.211 |
| **FSM†** | 0.131 | 0.203 | 0.226 | **0.223** | **0.240** | 0.188 | 0.202 |
| Mono‡ | 0.143 | 0.238 | 0.265 | 0.277 | 0.276 | 0.247 | 0.241 |
| **FSM‡** | 0.133 | 0.212 | 0.229 | 0.231 | 0.246 | 0.194 | 0.208 |
| **FSM - STC** | 0.133 | 0.219 | 0.246 | 0.252 | 0.259 | 0.197 | 0.218 |
| **FSM - PCC** | 0.131 | 0.206 | 0.228 | 0.238 | 0.248 | 0.188 | 0.207 |
| **FSM** | **0.130** | **0.201** | **0.224** | 0.229 | **0.240** | **0.186** | **0.201** |

(a) DDAD

| Method | Abs.Rel.↓ | | | | | | |
|---|---|---|---|---|---|---|---|
| | Front | F.Left | F.Right | B.Left | B.Right | Back | Avg. |
| Mono† | 0.214 | 0.304 | 0.388 | 0.314 | 0.438 | 0.304 | 0.327 |
| **FSM†** | 0.198 | 0.297 | **0.364** | 0.301 | **0.392** | 0.240 | 0.299 |
| Mono‡ | 0.251 | 0.403 | 0.546 | 0.429 | 0.616 | 0.321 | 0.428 |
| **FSM‡** | 0.200 | 0.337 | 0.448 | 0.354 | 0.521 | 0.267 | 0.355 |
| **FSM - STC** | 0.208 | 0.382 | 0.510 | 0.393 | 0.595 | 0.258 | 0.391 |
| **FSM - PCC** | 0.187 | 0.291 | 0.392 | 0.311 | 0.448 | 0.235 | 0.311 |
| **FSM** | **0.186** | **0.287** | 0.375 | **0.296** | 0.418 | **0.221** | **0.297** |

(b) nuScenes

Table 5.3: **Depth estimation results on multi-camera datasets**, using *FSM* relative to the single-camera photometric loss (*Mono*). The symbol † denotes per-frame median-scaling, and ‡ shared median-scaling (Eq. 5.4.1). *M* denotes the removal of masking (Sec. 5.3.3), *STC* the removal of spatio-temporal contexts (Sec. 5.3.1), and *PCC* the removal of pose consistency constraints (Sec. 5.3.2).
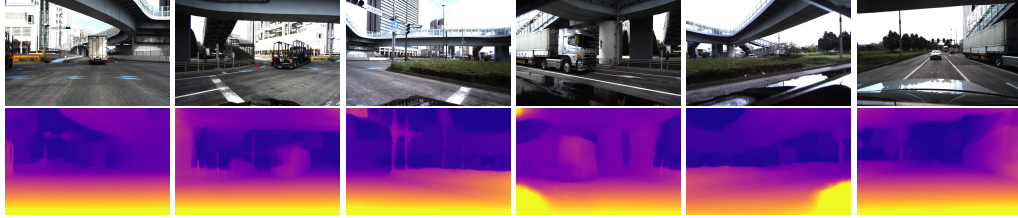
Figure 5.6: **Self-Supervised depth estimation results** using FSM on the *DDAD* dataset.
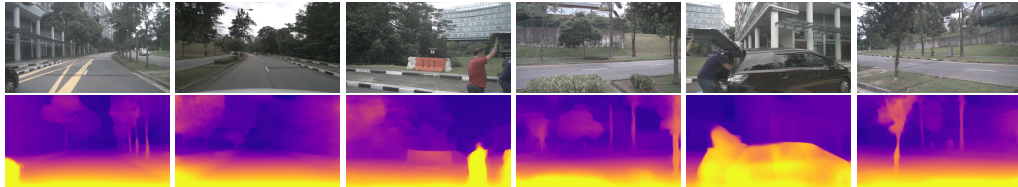


Figure 5.7: **Self-Supervised depth estimation results** using FSM on the *nuScenes* dataset.



(a) DDAD                                                (b) NuScenes
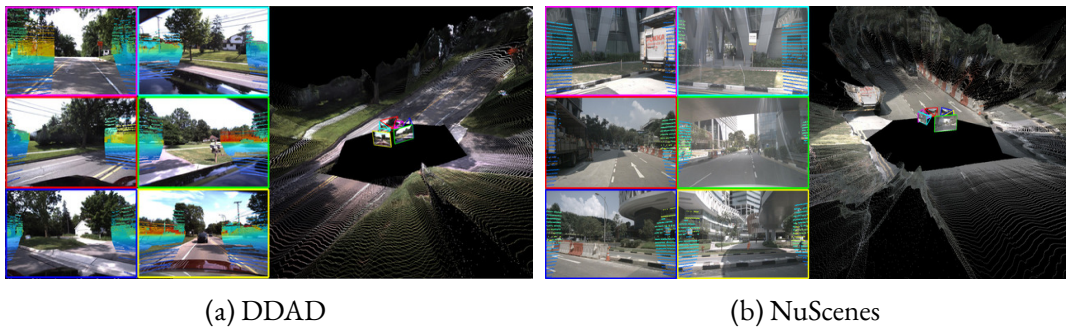
Figure 5.8: **Predicted pointclouds using FSM on the *DDAD* and *nuScenes* datasets**. For each dataset, the same network is used in all six images, predicted depth maps are lifted to 3D using camera intrinsics and extrinsics, and then combined *without any post-processing*. As a way to visualize camera overlapping, we also show projected LiDAR points from adjacent views overlaid on each RGB image (this information is not used at training or test time).

explicitly taking into consideration the stereo rectification on the KITTI dataset, our approach remains competitive with methods designed specifically for this setting.

## 5.4.5   Single-Camera Methods

Given that the majority of self-supervised monocular depth estimation papers focuses on single-image sequences with forward-facing cameras, we consider a variety of alternative baselines for our quantitative evaluation. In particular, we pick two state-of-the-art published methods: *monodepth2*

[88], that uses a simpler architecture with a series of modifications to the photometric loss calculation; and *PackNet* [9], that proposes an architecture especially designed for self-supervised depth learning. We use results of these methods for the front camera, as reported by Guizilini et al. [9], as baselines for our method.

We also take inspiration from the "Learning SfM from SfM" work [157, 158] and employ COLMAP [159], a state-of-the-art structure-from-motion system, on the unlabeled DDAD [9] training split to generate predicted depth maps, that are then used as pseudo-labels for supervised learning. Note that, while this approach is also self-supervised, it requires substantially more computation, since it processes all images from each sequence simultaneously to produce a single reconstruction of the scene.

The results of these experiments are summarized in Table 5.2. Our masking procedures for multi-camera training already significantly improve results from the previous state of the art [9]: from 0.162 to 0.139 absolute relative error (Abs Rel) on the front camera. By introducing spatial contexts (Equation 5.3.1), we gain a small but nontrivial boost in accuracy (especially for the nonforward facing cameras), but also learn *scale-aware* models by leveraging the camera extrinsics. Note that there is no limitation on the extrinsics transformation, only the assumption of some overlap between camera pairs (Figure 5.3). Finally, by introducing our proposed spatio-temporal contexts (Equation 5.3.2) and pose consistency constraints (Section 5.3.2) we further boost performance to 0.130, **achieving a new state of the art by a large margin**.

### 5.4.6 Multi-Camera Depth Estimation

We now evaluate FSM depth performance on all cameras of the *DDAD* and *nuScenes* datasets, ablating the impact of our contributions in the multi-camera setting (see Table 5.3).

**Photometric Masking**    As a baseline, we combine images from all cameras into a single dataset, without considering masking or cross-camera constraints. This is similar to Gordon et al. [4],

where multiple datasets from different cameras are pooled together to train a single model. As discussed in Section 5.3.3, the presence of self-occlusions on the DDAD dataset severely degrades depth performance when masking is not considered, reaching 0.380 Abs Rel (average of all cameras) versus 0.211 when self-occlusion masks are introduced (see Figure 5.4 for a qualitative comparison). Note that these results are still unscaled, and therefore median-scaling is required at test time for a proper quantitative evaluation.

**Spatio-Temporal Contexts**    The introduction of our proposed spatio-temporal contexts (STC), as described in Section 5.3.1, boosts performance on all cameras, from 0.211 to 0.202 (4.5%) on *DDAD* and 0.327 to 0.299 (9.1%) on *nuScenes*, by leveraging different levels of overlapping between views. This improvement becomes more apparent when considering our *shared median-scaling* evaluation protocol: 0.241 to 0.208 (16.1%) on *DDAD* and 0.428 to 0.355 (20.6%) on *nuScenes*. This is evidence that STC produces more consistent pointclouds across multiple cameras, as evidenced in Figure 5.5 and revealed by our proposed metric. Furthermore, the known extrinsics between cameras enables the generation of *scale-aware* models, with minimal degradation from their median-scaled counterparts: 0.207 versus 0.202 (2.2%) on *DDAD* and 0.311 vs 0.299 (4.0% on *nuScenes*). In this setting, we also evaluated the impact of STC relative to using only spatial and temporal contexts independently. As expected, there is a noticeable degradation in performance when spatio-temporal contexts are not considered: 0.206 to 0.218 (5.4%) on DDAD and 0.311 to 0.391 (25.8%) on nuScenes. We hypothesize that degradation on *nuScenes* is substantially higher due to an overall smaller overlapping between cameras (Figure 5.8), which will benefit more from STC as a way to improve cross-camera photometric losses for self-supervised training (Figure 5.3).

**Pose Consistency Constraints**    Finally, we include the pose consistency constraints (PCC) described in Section 5.3.2, as a way to enforce the learning of a similar rigid motion for all cameras. These additional constraints further boost performance, from 0.207 to 0.201 (2.9%) on *DDAD* and 0.311 to 0.297 (4.5%) on *nuScenes*. These improvements are more prominent on the side cameras, since per-frame pose estimation is harder in these cases due to larger relative motion and smaller overlap. Interestingly, the combination of all our contributions (masking, spatio-temporal contexts

and pose consistency constraints) lead to **scale-aware results that surpass median-scaled results**. This is evidence that FSM is capable of generating state-of-the-art metrically accurate models that are useful for downstream tasks.

## 5.5 Conclusion

Using cameras for 3D perception to complement or replace LiDAR scanners is an exciting new frontier for robotics. We have extended self-supervised learning of depth and ego-motion from monocular and stereo settings to the general multi-camera setting, predicting dense *scale-aware* point clouds around the ego-vehicle. We also introduced a series of key techniques that boost performance in this new setting, by leveraging known extrinsics between cameras: *spatial-temporal contexts*, *pose consistency constraints*, and studied the effects of *non-overlapping* and *self-occlusion* masking. In extensive experiments we demonstrated the capabilities of our methods and how they advance the state of the art.

# Chapter 6

# Learning Depth Fields

> The photograph itself doesn't interest me. I want only to capture a minute part of reality.
>
> Henri Cartier-Bresson.

In Chapter 5 we generalized monocular depth estimation to multi-camera rigs for training, but the method we proposed is still in the regime of single-frame *inference*: multi-view images are independently processed before accumulated after prediction. In order to incorporate information across frames, traditional methods rely on building multi-frame aggregation structures like cost volumes in order to enforce a multi-frame consistency constraint. Hybrid learning-based methods that rely on cost volumes and epipolar constraints depend on the correctness of such explicit geometric assumptions, e.g., accurate pose and calibration. Recently, a radically different approach–encoding geometric priors as *inputs* to general Transformer architectures–has been successfully used for tasks such as optical flow and stereo depth estimation. In this chapter, we introduce DeFiNe, a novel Transformer-based scene representation method for multi-view depth estimation. We also propose a series of 3D data augmentation techniques designed to inject key geometric inductive priors into our learned latent representation. Finally, we show that introducing view synthesis as an auxiliary task further improves depth estimation. As a result, DeFiNe achieves a new state of the art
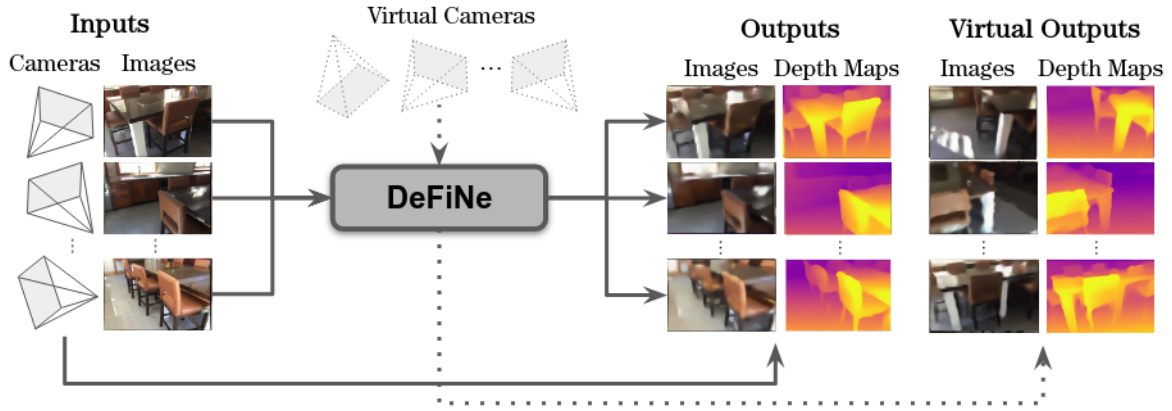
Figure 6.1: Our **DeFiNe** architecture achieves state of the art in multi-view depth estimation, while also enabling predictions from arbitrary viewpoints.

in stereo and video depth estimation without explicit geometric constraints. Our method achieves strong generalization performance, achieving state-of-the-art results in zero-shot domain transfer by a large margin, pointing to a new and robust *inductive bias* paradigm for geometric vision.

## 6.1 Introduction

Estimating 3D structure from a pair of images is a cornerstone problem of computer vision. Traditionally, this is treated as a correspondence problem, whereby one applies a homography to stereo-rectify the images, and then matches pixels (or patches) along epipolar lines to obtain disparity estimates. Given sufficiently accurate calibration, this disparity map can be converted into a per-pixel depth map. Contemporary approaches to stereo are specialized variants of classical methods, relying on correspondences to compute cost volumes, epipolar losses, bundle adjustment objectives, or projective multi-view constraints, among others. These are either baked directly into the model architecture or enforced as part of the loss function.

Applying the principles of classical vision in this way has given rise to architectures that achieve state-of-the-art results on tasks such as stereo depth estimation [10, 160], optical flow [161], and multi-view depth [58]. However, this success comes at a cost: each architecture is specialized and

purpose-built for a single task, and typically relies on an accurate underlying dataset-specific calibration. Though great strides have been made in alleviating the dependence on strong geometric assumptions [4, 98], two recent trends allow us to *decouple* the task from the architecture: (i) implicit representations and (ii) generalist networks. Our work draws upon both of these directions. Implicit representations of geometry and coordinate-based networks have recently achieved incredible popularity in the vision community. This direction is pioneered by work on neural radiance fields (NeRF) [36, 162], where a point- and ray-based parameterization along with a volume rendering objective allow simple MLP-based networks to achieve state-of-the-art view synthesis results. Follow-up works extend this coordinate-based representation to the pixel domain [163], allowing predicted views to be conditioned on image features.

The second trend in computer vision has been the use of generalist architectures, pioneered by Vision Transformers [164]. Emerging as an attention-based architecture for NLP, transformers have been used for a diverse set of tasks, including depth estimation [165, 166], optical flow [17], and image generation [167]. Transformers have also been applied to geometry-free view synthesis [168], demonstrating that attention can learn long-range correspondence between views for 2D-3D tasks. Scene Representation Transformers (SRT) [169] uses the transformer encoder-decoder model to learn scene representations for view synthesis from sparse, high-baseline data with no geometric constraints. However, owing to the $O(N^2)$ scaling of the self-attention module, experiments are limited to low-resolution images and require very long training periods (i.e., millions of iterations on a large-scale TPU architecture).

To alleviate the quadratic complexity of self-attention, the Perceiver architecture [16] disentangles the dimensionality of the latent representation from that of the inputs, by fixing the size of the latent representation. Perceiver IO [17] extends this architecture to allow for arbitrary outputs, with results in optical flow estimation that outperform traditional cost-volume based methods. Perceiver IO has also been recently used for stereo depth estimation [18], replacing traditional geometric *constraints* with input-level *inductive biases*. Building upon these works, we propose to learn a *geometric scene representation* for depth synthesis from novel viewpoints, including estimation, interpolation, and extrapolation. We expand the Perceiver IO framework to the scene representation setting,

taking sequences of images and predicting a consistent multi-view latent representation effective for downstream tasks. Taking advantage of the query-based nature of the Perceiver IO architecture, we propose a series of 3D augmentations that increase viewpoint density and diversity during training, thus encouraging (rather than enforcing) multi-view consistency. Furthermore, we show that the introduction of view synthesis as an auxiliary task, decoded from the same latent representation, improves depth estimation without additional ground-truth.

We test our model on the popular ScanNet benchmark [48], achieving state-of-the-art, real-time results for stereo depth estimation and competitive results for video depth estimation, without relying on memory- (and compute-) intensive operations such as cost volume aggregation and test-time optimization. We show that our 3D augmentations lead to significant improvements over baselines that are limited to the viewpoint diversity of training data. Furthermore, our zero-shot transfer results from ScanNet to 7-Scenes improve the state-of-the-art by a large margin, demonstrating that our method generalizes better than specialized architectures, which suffer from poor performance on out-of-domain data. Our contributions are summarized as follows:

- We use a generalist transformer-based architecture to learn a depth estimator from an arbitrary number of posed images. In this setting, we (i) propose a series of 3D augmentations that improve the geometric consistency of our learned latent representation; and (ii) show that **jointly learning view synthesis as an auxiliary task improves depth estimation**.

- Our DeFiNe not only achieve **state-of-the-art stereo depth estimation results** on the widely used ScanNet dataset, but also exhibit superior generalization properties with **state-of-the-art results on zero-shot transfer to 7-Scenes.**

- DeFiNe also **enables depth estimation from arbitrary viewpoints**. We evaluate this novel generalization capability in the context of *interpolation* (between timesteps), and *extrapolation* (future timesteps).

## 6.2   Related Work

**Multi-view Stereo**   Traditional multi-view stereo approaches have dominated even in the deep learning era. COLMAP [159] remains the standard framework for structure-from-motion, incorporating sophisticated bundle adjustment and keypoint refinement procedures, at the cost of speed. With the goal of producing closer to real-time estimates, multi-view stereo learning approaches adapt traditional cost volume-based approaches to stereo [10, 170] and multi-view [11, 150] depth estimation, often relying on known extrinsics to warp views into the frame of the reference camera. Recently, iterative refinement approaches that employ recurrent neural networks have made impressive strides in optical flow estimation [161]. Follow-on work applies this general recurrent correspondence architecture to stereo depth [160], scene-flow [171], and even full SLAM [172]. While their results are impressive, recurrent neural networks can be difficult to train, and test-time optimization increases inference time over a single forward pass.

Recently, transformer-based architectures [173] have replaced CNNs in many geometric estimation tasks. The Stereo Transformer [165] architecture replaces cost volumes with a correspondence approach inspired by sequence-to-sequence modeling. The Perceiver IO [17] architecture constitutes a large departure from cost volumes and geometric losses. For the task of optical flow, Perceiver IO feeds positionally encoded images through a Transformer [16], rather than using a cost volume for processing. Similar to our work, IIB [18] adapts the Perceiver IO architecture to generalized stereo estimation. Their main contribution is an epipolar parameterization as additional input-level inductive biases. We take an alternative approach and propose a series of geometric 3D data augmentation techniques designed to promote the learning of a *geometrically-consistent latent scene representation*, as well as using view synthesis as an auxiliary task. We extend Perceiver IO to the domain of scene representation learning. Our video-based representation (aided by our geometric augmentations) allows us to generalize to novel viewpoints, rather than be restricted to the stereo setting.

**Video Depth Estimation** Video and stereo depth estimation methods generally produce monocular depth estimates at test-time. ManyDepth [14] combines a monocular depth framework with multi-view stereo, aggregating predictions in a cost volume and thus enabling multi-frame inference at test-time. Recent methods accumulate predictions at train and test time, either with generalized stereo [174] or with sequence data [175]. DeepV2D [58] incorporates a cost-volume based multi-view stereo approach with an incremental pose estimator to iteratively improve depth and pose estimates at train and test time. Another line of work draws on the availability of monocular depth networks, which perform accurate but *multi-view inconsistent* estimates at test time [176]. In this setting, additional geometric constraints are enforced to finetune the network and improve multi-view consistency through epipolar constraints. Consistent Video Depth Estimation [176] refines COLMAP [159] results with a monocular depth network whose predictions are constrained to be multi-view consistent. Subsequent work jointly optimizes depth and pose for added robustness to challenging scenarios with poor calibration [177]. A recently proposed framework incorporates many of the architectural elements of prior work into a transformer-based architecture that takes video data as input for multi-view depth [178]. NeuralRecon [179] moves beyond depth-based architectures to learn TSDF volumes directly as a way to improve surface consistency.

**Novel View Synthesis** Since the emergence of neural radiance fields [36], implicit representations and volume rendering have emerged as the *de facto* standard for view synthesis. They parameterize viewing rays and points with positional encoding, and need to be re-trained on a scene-by-scene basis. Many recent improvements leverage depth supervision to improve view synthesis in a volume rendering framework [180, 181, 182, 183, 184]. Other works attempts to extend the NeRF approach beyond single scenes models by incorporating learned *features*, enabling few-shot volume rendering [163]. Feature-based methods have also treated view synthesis as a sequence-learning task, such as the scene representation transformer (SRT) architecture [169]. While SRT achieves impressive results, its transformer-based backbone restricts experiments to low-resolution images, and requires massive amounts of training data and computational resources.
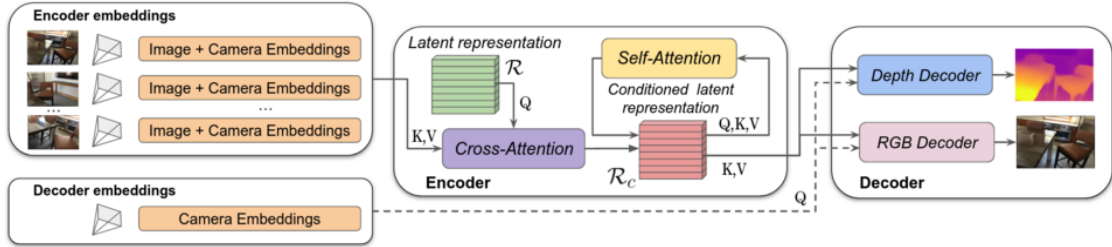
(a) Architecture overview.
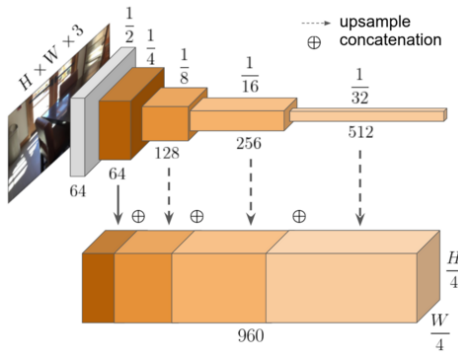


(b) Image embeddings.
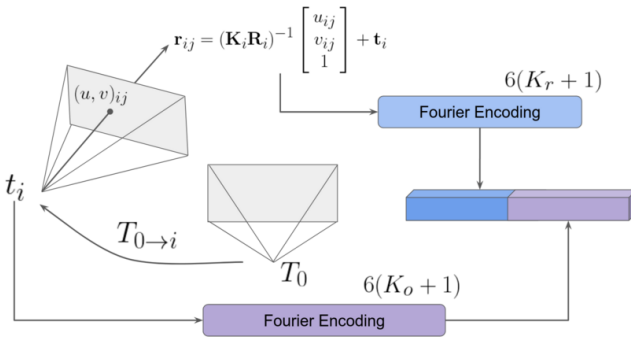
(c) Camera embeddings.

Figure 6.2: **Overview of our proposed DeFiNe architecture**, and the embeddings used to encode and decode information for depth and view synthesis.

## 6.3 Methodology

Our proposed DeFiNe architecture (Figure 6.2a) is designed with flexibility in mind, so data from different sources can be used as input and different output tasks can be estimated from the same latent space. We use Perceiver IO [16] as our general-purpose transformer backbone. During the encoding stage, our model takes RGB images from calibrated cameras, with known intrinsics and relative poses. The architecture processes this information according to the modality into different pixel-wise embeddings that serve as input to our Perceiver IO backbone. This encoded information can be queried using only camera embeddings, producing estimates from arbitrary viewpoints.

### 6.3.1 Perceiver IO

Perceiver IO [17] is a recent extension of the Perceiver [16] architecture. The Perceiver architecture alleviates one of the main weaknesses of transformer-based methods, namely the quadratic scaling

of self-attention with input size. This is achieved by using a fixed-size $N_l \times C_l$ latent representation $\mathcal{R}$, and learning to project high-dimensional $N_e \times C_e$ encoding embeddings onto this latent representation using cross-attention layers. Self-attention is performed in the lower-dimensional latent space, producing a *conditioned latent representation $\mathcal{R}_c$*, that can be queried using $N_d \times C_d$ decoding embeddings to generate estimates, again using cross-attention layers.

## 6.3.2   Input-Level Embeddings

**Image Embeddings (Figure 6.2b)**    Input $3 \times H \times W$ images are processed using a ResNet18 [185] encoder, producing a list of features maps at increasingly lower resolutions and higher dimensionality. Feature maps at $1/4$ the original resolution are concatenated with lower-resolution feature maps, after upsampling using bilinear interpolation. The resulting image embeddings are of shape $H/4 \times W/4 \times 960$, and are used in combination with the camera embeddings from each corresponding pixel (see below) to encode visual information.

**Camera Embeddings (Figure 6.2c)**    These embeddings capture multi-view scene geometry (e.g., camera intrinsics and extrinsics) as additional inputs during the learning process. Let $\mathbf{x}_{ij} = (u, v)$ be an image coordinate corresponding to pixel $i$ in camera $j$, with assumed known pinhole $3 \times 3$ intrinsics $\mathbf{K}_j$ and $4 \times 4$ transformation matrix $\mathbf{T}_j = \left[ \begin{smallmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{o} & 1 \end{smallmatrix} \right]$ relative to a canonical camera $\mathbf{T}_0$. Its origin $\mathbf{o}_j$ and direction $\mathbf{r}_{ij}$ are given by:

$$\mathbf{o}_j = -\mathbf{R}_j \mathbf{t}_j \quad , \quad \mathbf{r}_{ij} = \left( \mathbf{K}_j \mathbf{R}_j \right)^{-1} \left[ u_{ij}, v_{ij}, 1 \right]^T + \mathbf{t}_j \tag{6.3.1}$$

We ablate this variation in Table 6.1, showing that it leads to better performance for the task of depth estimation. These two vectors are then Fourier-encoded dimension-wise to produce higher-dimensional vectors, with a mapping of:

$$x \mapsto \left[ x, \ \sin(f_1 \pi x), \ \cos(f_1 \pi x), \ \ldots, \ \sin(f_K \pi x), \ \cos(f_K \pi x) \right]^\top \tag{6.3.2}$$

where $K$ is the number of Fourier frequencies used ($K_o$ for the origin and $K_r$ for the ray directions), equally spaced in the interval $[1, \frac{\mu}{2}]$. The resulting camera embedding is of dimensionality $2\big(3(K_o + 1) + 3(K_r + 1)\big) = 6\,(K_o + K_r + 2)$. During the encoding stage, camera embeddings are produced per-pixel assuming a camera with $1/4$ the original input resolution, resulting in a total of $\frac{HW}{16}$ vectors. During the decoding stage, embeddings from cameras with arbitrary calibration (i.e., intrinsics and extrinsics) can be queried to produce virtual estimates.

### 6.3.3 Geometric 3D Augmentations

Data augmentation is a core component of deep learning pipelines [186] that improves model robustness by applying transformations to the training data consistent with the data distribution in order to introduce desired equivariant properties. In computer vision and depth estimation in particular, standard data augmentation techniques are usually constrained to the 2D space and include color jittering, flipping, rotation, cropping, and resizing [13, 18]. Recent works have started looking into 3D augmentations [169], in an effort to increase robustness to errors in scene geometry, in terms of camera localization (i.e., extrinsics) and parameters (i.e., intrinsics). Conversely, we are interested in *encoding* scene geometry at the input-level, so our architecture can learn a multi-view-consistent geometric latent scene representation. Therefore, in this section we propose a series of 3D augmentations to increase the number and diversity of training views while maintaining the spatial relationship between cameras, thus enforcing desired equivariant properties within this setting.

**Virtual Camera Projection**    One of the key properties of our architecture is that it enables querying from arbitrary viewpoints, since only camera information (viewing rays) is required at the decoding stage. When generating predictions from these novel viewpoints, the network creates *virtual* information consistent with the implicit structure of the learned latent scene representation, conditioned on information from the encoded views. We evaluate this capability in Section 6.4.5, showing superior performance relative to the explicit projection of information from encoded views. Here, we propose to leverage this property at training time as well, generating additional supervision in the form of *virtual cameras* with corresponding ground-truth RGB images and depth maps

(a) Virtual Camera Projection.  (b) Canonical Jittering.
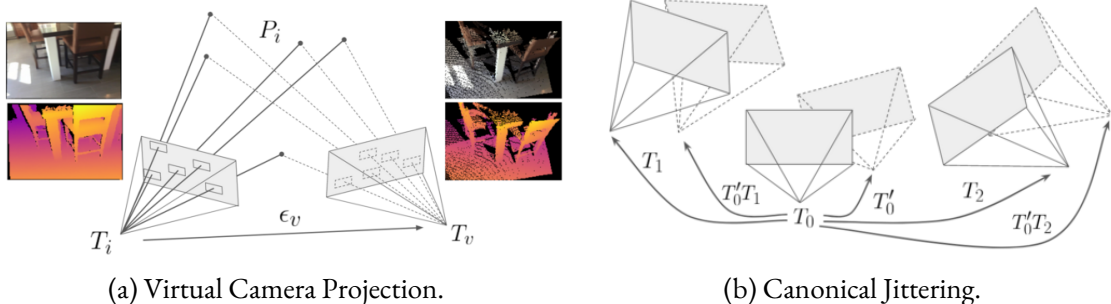
Figure 6.3: **Geometric augmentations**. (a) Information from camera $i$ is projected onto a virtual camera at $T_v$, creating additional supervision from other viewpoints. (b) Noise $T_0'$ is introduced to the canonical camera at $T_0$, and then propagated to other cameras to preserve relative scene geometry.

obtained by projecting available information onto these new viewpoints (Figure 6.3a). This novel augmentation technique forces the learned latent scene representation to be viewpoint-independent. Experiments show that this approach provides benefits in both the (a) stereo setting, with only two viewpoints; and (b) video setting, with similar viewpoints and a dominant camera direction.

From a practical perspective, virtual cameras are generated by adding translation noise $\boldsymbol{\epsilon}_v = [\epsilon_x, \epsilon_y, \epsilon_z]_v \sim \mathcal{N}(\mathbf{0}, \sigma_v \mathbf{I})$ to the pose of a camera $i$. The viewing angle is set to point towards the center $\mathbf{c}_i$ of the pointcloud $P_i$ generated by unprojecting information from the selected camera, which is also perturbed by $\boldsymbol{\epsilon}_c = [\epsilon_x, \epsilon_y, \epsilon_z]_c \sim \mathcal{N}(\mathbf{0}, \sigma_v \mathbf{I})$. When generating ground-truth information, we project the combined pointcloud of all available cameras onto these new viewpoints as a way to preserve full scene geometry. Furthermore, because the resulting RGB image and depth map will be sparse, we can improve efficiency by only querying at these specific locations.

**Canonical Jittering**   When operating in a multi-camera setting, the standard practice [150] is to select one camera to be the *reference* camera, and position all other cameras relative to it. One drawback of this convention is that one camera will always be at the same location (the origin of its own coordinate system) and therefore will produce the same camera embeddings, leading to overfitting. Intuitively, scene geometry should be invariant to the translation and rotation of the entire sensor suite. To enforce this property on our learned latent scene representation, we propose to inject some amount of noise to the canonical pose itself, so it is not located at the origin of

the coordinate system anymore. Note that this is different from methods that inject per-camera noise [187] with the goal of increasing robustness to localization errors. We only inject noise *once*, on the canonical camera, and propagate it to other cameras, so relative scene geometry is preserved within a translation and rotation offset (Figure 6.3b). However, this offset is reflected on the input-level embeddings produced by each camera, and thus forces the latent representation to be invariant to these transformations.

From a practical perspective, canonical jittering is achieved by randomly sampling translation $\boldsymbol{\epsilon}_t = [\epsilon_x, \epsilon_y, \epsilon_z]^\top \sim \mathcal{N}(\mathbf{0}, \sigma_t \mathbf{I})$ and rotation $\boldsymbol{\epsilon}_r = [\epsilon_\phi, \epsilon_\theta, \epsilon_\psi]^\top \sim \mathcal{N}(\mathbf{0}, \sigma_r \mathbf{I})$ errors from normal distributions with pre-determined standard deviations. Rotation errors are in Euler angles, and are converted to a $3 \times 3$ rotation matrix $\mathbf{R}_r$. These values are used to produce a jittered canonical transformation matrix $\mathbf{T}_0' = \begin{bmatrix} \mathbf{R}_r & \boldsymbol{\epsilon}_t \\ \mathbf{o} & 1 \end{bmatrix}$ that is then propagated to all other $N$ cameras, such that $T_i' = T_0' \cdot T_i, \forall i \in \{1, \ldots, N-1\}$.

**Canonical Randomization**   As an extension of canonical jittering, we also introduce canonical randomization to encourage generalization to different relative camera configurations, while still preserving scene geometry. Assuming $N$ cameras, we randomly select $o \in \{0, \ldots, N-1\}$ as the canonical index. Then $\forall i \in \{0, \ldots, N-1\}$, the relative transformation matrix $\mathbf{T}_i'$ given world-frame transformation matrix $\mathbf{T}_i$ is given by $\mathbf{T}_i' = \mathbf{T}_i \cdot \mathbf{T}_o^{-1}$.

### 6.3.4   Decoders

We use task-specific decoders, each consisting of one cross-attention layer between the $N_d \times C_d$ queries and the $N_l \times C_l$ conditioned latent representation $\mathcal{R}_c$, followed by a linear layer that creates an output of size $N_d \times C_o$, and a sigmoid activation function $\sigma(x) = \frac{1}{1+e^{-x}}$ to produce values in the interval $[0, 1]$. We set $C_o^d = 1$ for the depth estimation task, and $C_o^s = 3$ for view synthesis. Depth estimates are scaled to lie between a minimum $d_{min}$ and maximum $d_{max}$ range. Note that other decoders can be incorporated with DeFiNe without any modification to the underlying architecture, enabling the generation of multi-task estimates from arbitrary viewpoints.

**Losses**　We use an L1-log loss $\mathcal{L}_d = \|\log(d_{ij}) - \log(\hat{d}_{ij})\|_1$ to supervise depth estimation, where $\hat{d}_{ij}$ and $d_{ij}$ are depth estimates and ground-truth, respectively, for pixel $j$ at camera $i$. For view synthesis, we use an L2 loss $\mathcal{L}_s = \|I(\mathbf{p}_{ij}) - I(\hat{\mathbf{p}}_{ij})\|^2$, where $I(\hat{\mathbf{p}}_{ij})$ and $I(\mathbf{p}_{ij})$ are RGB estimates and ground-truth, respectively, for pixel $j$ at camera $i$. We use a weight coefficient $\lambda_s$ to balance these two losses, and another $\lambda_v$ to balance losses from available and virtual cameras. The final loss is of the form:

$$\mathcal{L} = \mathcal{L}_d + \lambda_s \mathcal{L}_s + \lambda_v \big( \mathcal{L}_{d,v} + \lambda_s(\mathcal{L}_{s,v}) \big) \tag{6.3.3}$$

Note that because our architecture enables querying at specific image coordinates, at training time we can improve efficiency by not computing estimates for pixels without ground truth (e.g., sparse depth maps or virtual cameras).

## 6.4　Experiments

### 6.4.1　Datasets

For more information on these datasets, please refer to Section 2.3.1.

**ScanNet [48]**　ScanNet is an RGB-D video dataset containing 2.5 million views from around 1500 scenes. It is the primary source of data for our experiments, and we explore two different settings: *stereo* and *video*. For the stereo experiments, we follow the same setting as Kusupati et al. [57]: we downsample scenes by a factor of 20, and use a custom split to create stereo pairs, resulting in 94212 training and 7517 test samples. For the video experiments, we follow the evaluation protocol of Teed et al. [58], with a total of 1405 scenes for training. For the test set, we use a custom split to select 2000 samples from 90 scenes not covered in the training set. Each training sample includes a target frame and a context of $[-3, 3]$ frames with stride 3. Each test sample includes a pair of frames, with a context of $[-3, 3]$ relative to the first frame of the pair with stride 3.

| Method | Abs.Rel↓ | RMSE↓ | $\delta_{1.25}$↑ |
|---|---|---|---|
| DPSNet [150] | 0.126 | 0.314 | — |
| NAS [57] | 0.107 | 0.281 | — |
| IIB [18] | 0.116 | 0.281 | 0.908 |
| **DeFiNe** $(128 \times 192)$ | 0.093 | 0.246 | 0.911 |
| **DeFiNe** $(240 \times 320)$ | **0.089** | **0.232** | **0.915** |

(a) Depth estimation results.
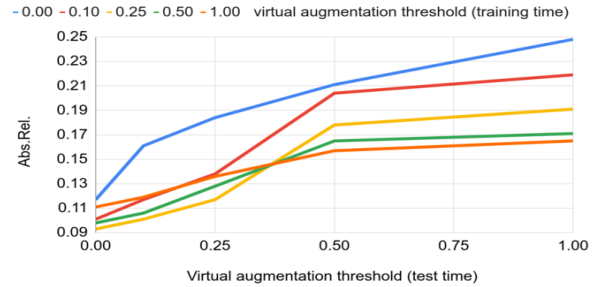


(b) Virtual depth estimation results.

Figure 6.4: **Depth estimation results on ScanNet-Stereo**. (a) We outperform contemporary methods by a large margin. (b) Depth estimation results on virtual cameras using different values for $\sigma_v$ at training and test time.

**7-Scenes [59]**   We also evaluate on the test split of 7-Scenes to measure zero-shot cross-dataset performance. This dataset, collected through KinectFusion [188], consists of $640 \times 480$ images in 7 settings, with a variable number of scenes in each setting. There are $500$–$1000$ images in each scene. We follow the evaluation protocol of Sun et al. [179], median-scaling predictions using ground-truth information before evaluation.

## 6.4.2   Stereo Depth Estimation

To test the benefits our proposed geometric 3D augmentation procedures, we first evaluate DeFiNe on the task of stereo depth estimation. Here, because each sample provides minimal information about the scene (i.e., only two frames), the introduction of additional virtual supervision should have the largest effect. We report our results in Figure 6.4a and visualize examples of reconstructed pointclouds in Figure 6.5. DeFiNe significantly outperforms other methods on this dataset, including IIB [18], which uses a similar Perceiver IO backbone and direct depth regression. Unlike IIB, which employs epipolar cues as additional input-level inductive biases, our approach focuses on generating additional virtual supervision in order to build a scene representation that encourages multi-view consistency, leading to a large $(20\%)$ relative improvement.

| | Variation | Lower is better ↓ | | | Higher is better ↑ | | |
|---|---|---|---|---|---|---|---|
| | | Abs. Rel | Sq. Rel | RMSE | $\delta_{1.25}$ | $\delta_{1.25^2}$ | $\delta_{1.25^3}$ |
| 1 | Depth-Only | 0.098 | 0.046 | 0.257 | 0.902 | 0.972 | 0.990 |
| 2 | w/ Conv. RGB encoder [17] | 0.114 | 0.058 | 0.294 | 0.866 | 0.961 | 0.982 |
| 3 | w/ 64-dim R18 RGB encoder | 0.104 | 0.049 | 0.270 | 0.883 | 0.966 | 0.985 |
| 4 | w/o camera information | 0.229 | 0.157 | 0.473 | 0.661 | 0.874 | 0.955 |
| 5 | w/o global rays encoding | 0.097 | 0.047 | 0.261 | 0.897 | 0.962 | 0.988 |
| 6 | w/ equal loss weights | 0.095 | 0.047 | 0.259 | 0.908 | 0.968 | 0.990 |
| 7 | w/ epipolar cues [18] | 0.094 | 0.048 | 0.254 | 0.905 | 0.972 | 0.990 |
| 8 | w/o Augmentations | 0.117 | 0.060 | 0.291 | 0.870 | 0.959 | 0.981 |
| 9 | w/o Virtual Cameras | 0.104 | 0.058 | 0.268 | 0.891 | 0.965 | 0.986 |
| 10 | w/o Canonical Jittering | 0.099 | 0.046 | 0.261 | 0.897 | 0.970 | 0.988 |
| 11 | w/o Canonical Randomization | 0.096 | 0.044 | 0.253 | 0.905 | 0.971 | 0.989 |
| | **DeFiNe** | **0.093** | **0.042** | **0.246** | **0.911** | **0.974** | **0.991** |

Table 6.1: **Ablation study for ScanNet-Stereo**, using different variations.

## 6.4.3 Ablation Study

We perform a detailed ablation study to evaluate the effectiveness of each component in our proposed architecture, with results shown in Table 6.1. Firstly, we evaluate performance when (1) learning only depth estimation, and see that the joint learning of view synthesis as an auxiliary task leads to significant improvements. The claim that depth estimation improves view synthesis has been noted before [181, 183], and attributed to the well-known fact that multi-view consistency facilitates the generation of images from novel viewpoints. However, our experiments also show the inverse: view synthesis improves depth estimation. Our hypothesis is that appearance is required to learn multi-view consistency, since it enables visual correlation across frames. By introducing view synthesis as an additional task, we are also encoding appearance-based information into our latent representation. This leads to improvements in depth estimation even though no explicit feature matching is performed at an architecture or loss level.

We also ablate different variations of our RGB encoder for the generation of image embeddings and show that (2) our proposed multi-level feature map concatenation (Figure 6.2b) leads to the best results relative to the standard single convolutional layer, or (3) using 1/4 resolution ResNet18 64-dimensional feature maps. Similarly, we also ablate some of our design choices, namely (4) the use
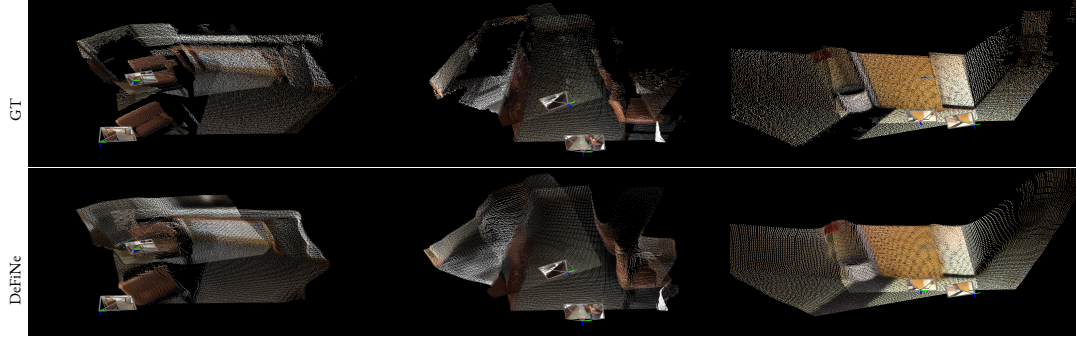
Figure 6.5: **Reconstructed two-view pointclouds**, from ScanNet-Stereo. DeFiNe pointclouds are generated using both depth maps and RGB images queried from our learned latent representation.
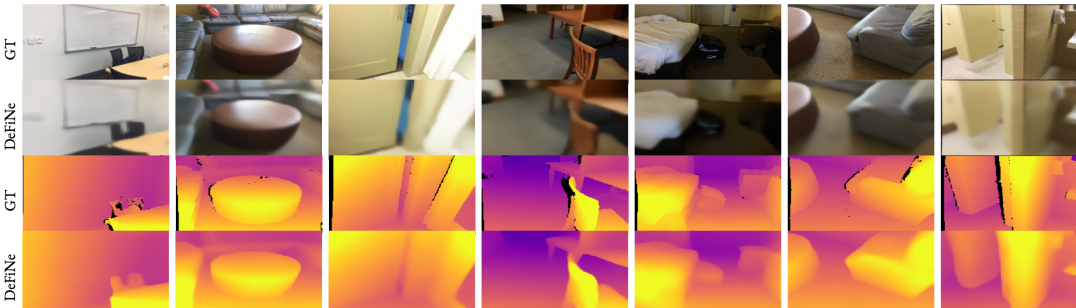


Figure 6.6: **Depth estimation and view synthesis results** on ScanNet. Although view synthesis is not our primary goal, it can be achieved with minimal modifications, and we show that it improves depth estimation performance.

of camera embeddings instead of positional encodings; (5) global viewing rays (Section 6.3.2) instead of traditional relative viewing rays; (6) the use of $\lambda_s = 1$ in the loss calculation (Equation 6.3.3) such that both depth and view synthesis tasks have equal weights; and (7) the use of epipolar cues as additional geometric embeddings, as proposed by IIB [18]. As expected, camera embeddings are crucial for multi-view consistency, and global viewing rays improve over the standard relative formulation. Interestingly, using a smaller $\lambda_s$ degrades depth estimation performance, providing further evidence that the joint learning of view synthesis is beneficial for multi-view consistency. We did not observe meaningful improvements when incorporating the epipolar cues from IIB [18], indicating that DeFiNe is capable of directly capturing these constraints at an input-level due to the increase in viewpoint diversity. Lastly, we ablate the impact of our various proposed geometric augmentations (Section 6.3.3), showing that they are key to our reported state-of-the-art performance.

Lastly, we evaluate depth estimation from virtual cameras, using different noise levels $\sigma_v$ at test time. Different models were also trained considering different noise levels, with results reported in Figure 6.4b. From these results we can see that the optimal virtual noise level, when evaluating at the target location, is $\sigma_v = 0.25\text{m}$ (yellow line), relative to the baseline without virtual noise (blue line). However, models trained with higher virtual noise (e.g., the orange line, with $\sigma_v = 1\text{m}$) are more robust to larger deviations from the target location.

### 6.4.4  Video Depth Estimation

To highlight the flexibility of our proposed architecture, we also experimented using video data from ScanNet following the training protocol of Tang et al. [189]. We evaluated performance on both ScanNet itself, using their evaluation protocol [189], as well as zero-shot transfer (without fine-tuning) to the 7-Scenes dataset. Table 6.2 reports quantitative results, while Figure 6.6 provides qualitative examples. On ScanNet, DeFiNe outperforms most published approaches, significantly improving over DeMoN [174], BA-Net [189], and CVD [176] both in terms of performance and speed. We are competitive with DeepV2D [58] in terms of performance, and roughly $14\times$ faster, owing to the fact that DeFiNe does not require bundle adjustment or any sort of test-time optimization. In fact, our inference time of 49 ms can be split into 44 ms for encoding and only 5 ms for decoding, enabling very efficient generation of depth maps after information has been encoded once. The only method that outperforms DeFiNe in terms of speed is NeuralRecon [179], which uses a sophisticated TSDF integration strategy. Performance-wise, we are also competitive with NeuralRecon, improving over their reported results in one of the three metrics (Sq. Rel).

Next, we evaluate zero-shot transfer from ScanNet to 7-Scenes, which is a popular test of generalization for video depth estimation. In this setting, DeFiNe significantly improves over all other methods, including DeepV2D (which fails to generalize) and NeuralRecon ($\sim 40\%$ improvement). We attribute this large gain to the highly intricate and specialized nature of these other architectures. In contrast, our method has no specialized module and instead learns a geometrically-consistent multi-view latent representation. In summary, we achieve competitive results on ScanNet while

| Method | Abs. Rel↓ | Sq. Rel↓ | RMSE↓ | Speed (ms)↓ |
|---|---|---|---|---|
| **ScanNet test split** [189] | | | | |
| DeMoN [174] | 0.231 | 0.520 | 0.761 | 110 |
| MiDas-v2 [190] | 0.208 | 0.318 | 0.742 | - |
| BA-Net [189] | 0.091 | 0.058 | 0.223 | 95 |
| CVD [176] | 0.073 | 0.037 | 0.217 | 2400 |
| DeepV2D [58] | 0.057 | **0.010** | <u>0.168</u> | 690 |
| NeuralRecon [179] | **0.047** | 0.024 | **0.164** | **30** |
| **DeFiNe** ($128 \times 192$) | 0.059 | 0.022 | 0.184 | <u>49</u> |
| **DeFiNe** ($240 \times 320$) | <u>0.056</u> | <u>0.019</u> | 0.176 | 78 |
| **Zero-shot transfer to 7-Scenes** [59] | | | | |
| DeMoN [174] | 0.389 | 0.420 | 0.855 | 110 |
| NeuralRGBD [191] | 0.176 | 0.112 | 0.441 | 202 |
| DPSNet [150] | 0.199 | 0.142 | 0.438 | 322 |
| DeepV2D [58] | 0.437 | 0.553 | 0.869 | 347 |
| CNMNet [192] | 0.161 | 0.083 | 0.361 | 80 |
| NeuralRecon [179] | 0.155 | 0.104 | 0.347 | **30** |
| EST [18] | <u>0.118</u> | <u>0.052</u> | <u>0.298</u> | 71 |
| **DeFiNe** ($128 \times 192$) | **0.100** | **0.039** | **0.253** | <u>49</u> |

Table 6.2: **Depth estimation results on ScanNet and 7-Scenes**. DeFiNe is competitive with other state-of-the-art methods on ScanNet, and outperforms all published methods in zero-shot transfer to 7-Scenes by a large margin.

significantly improving the state-of-the-art for video depth generalization, as evidenced by the large gap between DeFiNe and the best-performing methods on the 7-Scenes benchmark.

## 6.4.5 Depth from Novel Viewpoints

We previously discussed the strong performance that DeFiNe achieves on traditional depth estimation benchmarks, and showed how it improves out-of-domain generalization by a wide margin. Here, we explore another aspect of generalization that our architecture enables: viewpoint generalization. This is possible because, in addition to traditional depth estimation from RGB images, DeFiNe can also generate depth maps from arbitrary viewpoints since it only requires camera embeddings to decode estimates. We explore this capability in two different ways: *interpolation* and *extrapolation*. When interpolating, we encode frames at $\{t-5, \ldots, t+5\}$, and decode virtual depth

(a) Depth interpolation results.

(b) Depth extrapolation results.



(c) Depth extrapolation to future timesteps. Images and ground-truth depth maps are displayed only for comparison. DeFiNe can complete unseen portions of the scene in a geometrically-consistent way, generating dense depth maps from novel viewpoints.
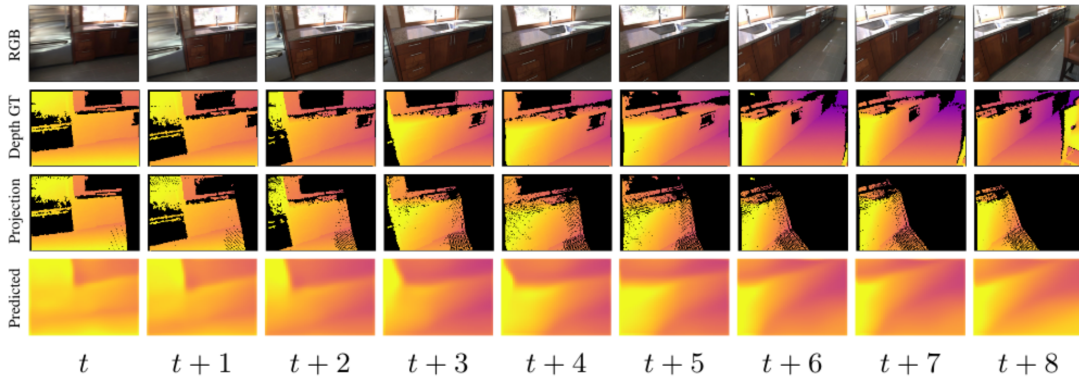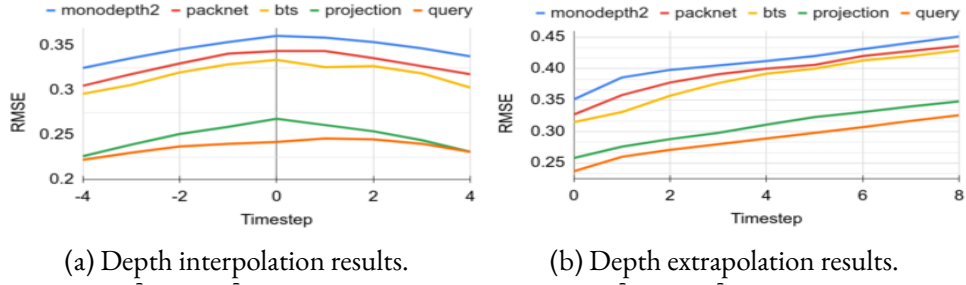
Figure 6.7: **Depth estimation results** from novel viewpoints.

maps at locations $\{t - 4, \ldots, t + 4\}$. When extrapolating, we encode frames at $\{t - 5, \ldots, t - 1\}$, and decode virtual depth maps at locations $\{t, \ldots, t + 8\}$. We use the same training and test splits as in our stereo experiments, with a downsample factor of 20 to encourage smaller overlaps between frames. As baselines for comparison, we consider the explicit projection of 3D information from encoded frames onto these new viewpoints. We evaluate both standard depth estimation networks [13, 9, 75] as well as DeFiNe itself, that can be used to either explicitly project information from encoded frames onto new viewpoints (projection), or query from the latent representation at that same location (query).

Figure 6.7 reports results in terms of root mean squared error (RMSE) considering only valid projected pixels. The first noticeable aspect is that our multi-frame depth estimation architecture significantly outperforms other single-frame baselines. However, and most importantly, results obtained by implicit querying consistently outperform those obtained via explicit projection. This

indicates that our model is able to improve upon available encoded information, via the learned latent representation. Furthermore, we also generate geometrically consistent estimates for areas without valid explicit projections (Figure 6.7c). As the camera tilts to the right, the floor is smoothly recreated in unobserved areas, as well as the partially observed bench. Interestingly, the chair at the extreme right was not recreated, which could be seen as a failure case. However, because the chair was never observed in the first place, it is reasonable for the model to assume the area is empty, and recreate it as a continuation of the floor.

## 6.5   Conclusion

We introduced DeFiNe, a generalist scene representation framework for multi-view depth estimation. Rather than explicitly enforcing geometric constraints at an architecture or loss level, we use geometric embeddings to condition network inputs, alongside visual information. To learn a geometrically-consistent latent representation, we propose a series of 3D augmentations designed to promote viewpoint, rotation, and translation equivariance. We also show that the introduction of view synthesis as an auxiliary task improves depth estimation without requiring additional ground-truth. We achieve state-of-the-art results on the popular ScanNet stereo benchmark, and competitive results on the ScanNet video benchmark with no iterative refinement or explicit geometry modeling. We also demonstrate strong generalization properties by achieving state-of-the-art results in zero-shot transfer from ScanNet to 7-Scenes. The general nature of our framework enables many exciting avenues for future work, including the estimation of additional tasks, extension to dynamic scenes, spatio-temporal representations, and uncertainty estimation.

# Chapter 7

# Conclusion

> It is an illusion that photos are made with the camera... they are made with the eye, heart and head.
>
> ———————————————————
> Henri Cartier-Bresson.

Monocular depth estimation has long been the realm of benchmark datasets: networks are trained and evaluated on calibrated and rectified data representing only a small subset of possible imaging geometries and scenes. In this thesis, we have expanded the domain of monocular depth networks from perspective views of streets and offices to underwater caves, micro-aerial vehicles, and exotic camera-mirror systems.

First, we extended the standard self-supervised learning framework from *perspective* cameras to a much more general class of parametric camera models (Chapter 3). This model achieved state-of-the-art depth results on a challenging fisheye benchmark in a completely self-calibrated fashion, learning the camera parameters jointly with depth and pose. The resulting camera parameters were comparable to traditional manual, target-based calibration.

Next, we relaxed the camera model assumptions even further, learning a fully non-parametric central camera model (Chapter 4) that can represent any central imaging geometry. Inspired by the

"raxel" model of Grossberg and Nayar [120], we devised a differentiable per-pixel model and applied it to settings where the perspective assumption often fails. Our experiments on a dashboard camera dataset (where the camera is behind a curved windshield), and on an underwater sequence (where the camera lies behind a glass dome) demonstrated the strength of a flexible model in settings where parametric models fail.

The next step was multi-camera systems with arbitrary overlap. In Chapter 5 we investigated how multi-camera constraints could allow monodepth estimators to be used to predict consistent, 360° point clouds, learning a single network across radically different viewpoints. This architecture achieved a large improvement on the DDAD benchmark, demonstrating the advantage of multi-camera training. The chapter focused on the calibrated setting, but it should be possible to jointly self-calibrate multi-camera rigs while learning multi-view consistent depth using some of the strategies introduced in prior chapters; we leave this interesting extension to future work.

The self-supervised architectures described above share a number of limitations: they do not explicitly model dynamic objects, they use a simple photometric loss that struggles on specular surfaces, and they are single-frame methods that do not take advantage of video input at test time. Recent advances using synthetic data [193], multi-frame inference [14, 194], and neural fields [195] have begun to alleviate some of these issues, but interesting future work remains.

**Transformers** In Chapter 1, we recalled a particularly poignant remark from a panelist at a CVPR workshop: he predicted that over time, traditional geometric constraints and solvers would fall by the wayside, and end-to-end learning would prevail. Our self-supervised depth architectures have the flavor of a purely end-to-end solution to geometric vision—our camera models are learned end-to-end without supervision, and our pose networks take input image pairs and predict transformations directly. Still, elements of traditional constraints remain—the networks are supervised using a depth-based view synthesis objective [12], which depends on the (calibrated or predicted) camera parameters at the *loss level*. Our FSM architecture explicitly incorporates multi-view consistency for both the depth and pose networks, relying on the accuracy of both the intrinsics and extrinsics.

For recent multi-frame monodepth extensions [14, 194] the dependence is even stronger: to aggregate information across multiple frames, we need to construct a cost volume that becomes specialized to camera information at train time. Any mis-calibration (or shift between train and test calibration) has a large impact on accuracy; the network inherently "overfits" to the training calibration in terms of both loss and architecture. The question remains: how do we fully *decouple* the loss and architecture from the camera?

Recent trends in vision point to a potential solution. PercieverIO [16], an extension of Vision Transformers [15], sets up a simple architecture for vision: positionally-encoded input is fed into a generic architecture that can map any input modality (e.g., speech, text, images, point clouds) to any output modality (e.g., class labels, semantic masks, optical flow, depth maps). A domain-specific *inductive bias* is provided at the input level, and the architecture can learn to use or ignore the information as it sees fit. The simple inductive bias of patch positions allows PercieverIO to outperform RAFT [161], a sophisticated recurrent 4D cost volume processing architecture, on the task of optical flow. Subsequent work adapted this architecture to stereo depth estimation [18]. An "epipolar embedding" representing the epipolar matching constraint is concatenated with positional encodings as input to a PercieverIO architecture, and the resulting depth estimator achieves competitive results with sophisticated cost-volume and bundle-adjustment-based stereo methods [18].

Our DeFiNe architecture introduced in Chapter 6 builds on the work of Yifan et al. [18] to train a multi-view consistent video depth estimator. We introduce a novel view query augmentation procedure, which allows us to generate many virtual views without access to RGB images (depth maps can be easily sampled at novel views), encouraging a multi-view consistent scene representation rather than enforcing consistency at the loss level. Applying this augmentation to the stereo task in [18] we find that the epipolar embedding is unnecessary. With only the standard viewing ray embedding, our view-augmented network significantly outperforms their architecture. In addition to the depth loss, we introduce an auxiliary view synthesis task to further improve our geometric scene representation. Taking inspiration from the light field view synthesis networks found in LFN [196] and Scene Representation Transformers [169], we decode RGB values in addition

to depth for query viewing rays. This form of "free" supervision further encourages geometric consistency between source and target views.

DeFiNe is competitive with state-of-the-art video depth estimators that use sophisticated cost volume aggregation and even bundle adjustment, but dramatically outperforms these methods on zero-shot generalization, demonstrating a major strength of decoupling the architecture from the camera parameters. Though the generalization results are compelling, a serious limitation remains—DeFiNe is supervised. Ideally, a camera-agnostic depth estimation architecture would combine the strengths of prior self-supervised depth networks with the new "input-level inductive bias" approach. One possible avenue for self-supervision is DeFiNe's light field view synthesis decoder. Removing supervised depth prediction, sparse depth maps can be extracted from a light field network [196], enabling a generic depth estimator with no ground-truth supervision.

We predict that novel input-level biases will continue to surpass bespoke geometric vision architectures on a variety of 3D tasks in the 2020s, proving that the CVPR workshop panelist's statement in 2019 was not quite as absurd as it initially seemed.

# Bibliography

[1] Vladyslav Usenko, Nikolaus Demmel, and Daniel Cremers. The double sphere camera model. In *Proceedings of International Conference on 3D Vision (3DV)*, pages 552–560, 2018.

[2] V. Usenko, N. Demmel, D. Schubert, J. Stueckler, and D. Cremers. Visual-inertial mapping with non-linear factor recovery. *IEEE Robotics and Automation Letters*, 5(2):422–429, 2020. doi: 10.1109/LRA.2019.2961227.

[3] Michael Burri, Janosch Nikolic, Pascal Gohl, Thomas Schneider, Joern Rehder, Sammy Omari, Markus W Achtelik, and Roland Siegwart. The euroc micro aerial vehicle datasets. *The International Journal of Robotics Research*, 35(10):1157–1163, 2016.

[4] Ariel Gordon, Hanhan Li, Rico Jonschkowski, and Anelia Angelova. Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8977–8986, 2019.

[5] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

[6] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018.

[7] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.

[8] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016.

[9] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2485–2494, 2020.

[10] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 66–75, 2017.

[11] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. MVSNet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 767–783, 2018.

[12] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[13] Clément Godard, Oisin Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging into self-supervised monocular depth prediction. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019.

[14] Jamie Watson, Oisin Mac Aodha, Victor Prisacariu, Gabriel Brostow, and Michael Firman. The temporal opportunist: Self-supervised multi-frame monocular depth. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1164–1174, 2021.

[15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image

recognition at scale. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.

[16] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 4451–4464, 2021.

[17] Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David Ding, Skanda Koppula, Daniel Zoran, Andrew Brock, Evan Shelhamer, et al. Perceiver IO: A general architecture for structured inputs & outputs. *arXiv preprint arXiv:2107.14795*, 2021.

[18] Wang Yifan, Carl Doersch, Relja Arandjelović, João Carreira, and Andrew Zisserman. Input-level inductive biases for 3d reconstruction. *arXiv preprint arXiv:2112.03243*, 2021.

[19] Vladyslav Usenko, Nikolaus Demmel, David Schubert, Jörg Stückler, and Daniel Cremers. Visual-inertial mapping with non-linear factor recovery. *IEEE Robotics and Automation Letters*, 5(2):422–429, 2019.

[20] Michael D Grossberg and Shree K Nayar. A general imaging model and a method for finding its parameters. In *Proceedings of the International Conference on Computer Vision (ICCV)*, volume 2, pages 108–115. IEEE, 2001.

[21] Dominique Raynaud. *Critical Edition of Ibn Al-Haytham's On the Shape of the Eclipse*. Springer, 2016.

[22] Stuart I Granshaw. Laussedat bicentenary: origins of photogrammetry, 2019.

[23] Noah Snavely, Steven M Seitz, and Richard Szeliski. Photo tourism: exploring photo collections in 3d. In *Special Interest Group on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pages 835–846. 2006.

[24] KB Atkinson. Deville and photographic surveying. *The Photogrammetric Record*, 15(86): 189–195, 1995.

[25] Timothy A Clarke and John G Fryer. The development of camera calibration methods and models. *The Photogrammetric Record*, 16(91):51–66, 1998.

[26] C Brown Duane. Close-range camera calibration. *Photogramm. Eng*, 37(8):855–866, 1971.

[27] Zhengyou Zhang. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(11):1330–1334, 2000.

[28] Edwin Olson. AprilTag: A robust and flexible visual fiducial system. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3400–3407, 2011.

[29] Richard I Hartley. In defense of the eight-point algorithm. *IEEE Transactions on pattern analysis and machine intelligence*, 19(6):580–593, 1997.

[30] Alex Kendall, Matthew Grimes, and Roberto Cipolla. PoseNet: A convolutional network for real-time 6-DOF camera relocalization. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 2938–2946, 2015.

[31] Torsten Sattler, Qunjie Zhou, Marc Pollefeys, and Laura Leal-Taixe. Understanding the limitations of cnn-based absolute camera pose regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3302–3312, 2019.

[32] Wang Zhao, Shaohui Liu, Yezhi Shu, and Yong-Jin Liu. Towards better generalization: Joint depth-pose learning without posenet. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9151–9161, 2020.

[33] Yuhua Chen, Cordelia Schmid, and Cristian Sminchisescu. Self-supervised learning with geometric constraints in monocular video: Connecting flow, depth, and camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7063–7072, 2019.

[34] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. 2000.

[35] Gengshan Yang and Deva Ramanan. Learning to segment rigid motions from two frames. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1266–1275, 2021.

[36] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 405–421, 2020.

[37] Srikumar Ramalingam, Peter Sturm, and Suresh K Lodha. Towards generic self-calibration of central cameras. In *Computer Vision and Image Understanding*, 2005.

[38] Srikumar Ramalingam, Suresh K Lodha, and Peter Sturm. A generic structure-from-motion framework. *Computer Vision and Image Understanding*, 103(3):218–228, 2006.

[39] Thomas Schöps, Viktor Larsson, Marc Pollefeys, and Torsten Sattler. Why having 10,000 parameters in your camera model is better than twelve. *arXiv preprint arXiv:1912.02908*, 2019.

[40] Yoonwoo Jeong, Seokjun Ahn, Christopher Choy, Animashree Anandkumar, Minsu Cho, and Jaesik Park. Self-calibrating neural radiance fields. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.

[41] Ashutosh Saxena, Jamie Schulte, Andrew Y Ng, et al. Depth estimation using monocular and stereo cues. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, volume 7, pages 2197–2203, 2007.

[42] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and F. Li. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.

[43] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva. Learning deep features for scene recognition using places database. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.

[44] T. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common objects in context. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2014.

[45] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-image depth perception in the wild. *Advances in neural information processing systems*, 29, 2016.

[46] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 3354–3361. IEEE, 2012.

[47] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from RGBD images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2012.

[48] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[49] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. M. Lopez. The SYNTHIA dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[50] Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-object tracking analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[51] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1746–1754, 2017.

[52] Abhishek Kadian, Joanne Truong, Aaron Gokaslan, Alexander Clegg, Erik Wijmans, Stefan Lee, Manolis Savva, Sonia Chernova, and Dhruv Batra. Sim2real predictivity: Does evaluation in simulation predict real-world performance? *IEEE Robotics and Automation Letters*, 5(4): 6670–6677, 2020.

[53] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The KITTI dataset. *International Journal of Robotics Research*, 2013.

[54] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.

[55] Miriam Schönbein, Tobias Strauß, and Andreas Geiger. Calibrating and centering quasi-central catadioptric cameras. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 4443–4450, 2014.

[56] Zichao Zhang, Henri Rebecq, Christian Forster, and Davide Scaramuzza. Benefit of large field-of-view cameras for visual odometry. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 801–808. IEEE, 2016.

[57] Uday Kusupati, Shuo Cheng, Rui Chen, and Hao Su. Normal assisted stereo depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2189–2199, 2020.

[58] Zachary Teed and Jia Deng. DeepV2D: Video to depth with differentiable structure from motion. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.

[59] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene coordinate regression forests for camera relocalization in RGB-D images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2930–2937, 2013.

[60] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. KinectFusion: Real-time dense surface mapping and tracking. In *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2011.

[61] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11621–11631, 2020.

[62] R. Garg, V. K. B.G., G. Carneiro, and I. Reid. Unsupervised CNN for single view depth estimation: Geometry to the rescue. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016.

[63] FARO™ S350 scanner. [https://www.faro.com/products/construction-bim-cim/faro-focus/](https://www.faro.com/products/construction-bim-cim/faro-focus/), 2019. URL [https://www.faro.com/products/construction-bim-cim/faro-focus/](https://www.faro.com/products/construction-bim-cim/faro-focus/). Accessed: 2019-03-20.

[64] M.A. Fischler and R.C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24 (6):381–395, June 1981.

[65] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor segmentation and support inference from RGBD images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2012.

[66] Ruohan Gao, Changan Chen, Ziad Al-Halah, Carl Schissler, and Kristen Grauman. Visualechoes: Spatial image representation learning through echolocation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 658–676. Springer, 2020.

[67] Julian Lienen, Eyke Hullermeier, Ralph Ewerth, and Nils Nommensen. Monocular depth estimation via listwise ranking using the plackett-luce model. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14595–14604, 2021.

[68] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3D scene shape from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 204–213, 2021.

[69] Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. Structure-guided ranking loss for single image depth prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 611–620, 2020.

[70] Soomin Kim, Yuchi Huo, and Sung-Eui Yoon. Single image reflection removal with physically-based training images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5164–5173, 2020.

[71] Ashutosh Saxena, Sung H Chung, Andrew Y Ng, et al. Learning depth from single monocular images. In *NIPS*, volume 18, pages 1–8, 2005.

[72] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 2650–2658, 2015.

[73] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[74] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *Proceedings of International Conference on 3D Vision (3DV)*, pages 239–248, 2016.

[75] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv:1907.10326*, 2019.

[76] Katrin Lasinger, René Ranftl, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *arXiv preprint arXiv:1907.01341*, 2019.

[77] Cesar Cadena, Yasir Latif, and Ian D Reid. Measuring the performance of single image depth estimation methods. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4150–4157. IEEE, 2016.

[78] Vitor Guizilini, Rui Hou, Jie Li, Rares Ambrus, and Adrien Gaidon. Semantically-guided representation learning for self-supervised monocular depth. *arXiv preprint arXiv:2002.12319*, 2020.

[79] Ibraheem Alhashim and Peter Wonka. High quality monocular depth estimation via transfer learning. *arXiv preprint arXiv:1812.11941*, 2018.

[80] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402, 2003.

[81] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[82] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[83] Ravi Garg, Vijay Kumar Bg, Gustavo Carneiro, and Ian Reid. Unsupervised CNN for single view depth estimation: Geometry to the rescue. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 740–756, 2016.

[84] Clément Godard, Oisin 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 (CVPR)*, pages 270–279, 2017.

[85] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. *arXiv:1506.02025*, 2015.

[86] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 2004.

[87] Chang Shu, Kun Yu, Zhixiang Duan, and Kuiyuan Yang. Feature-metric loss for self-supervised learning of depth and egomotion. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 572–588, 2020.

[88] Clément Godard, Oisin Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 3828–3838, 2019.

[89] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8445–8453, 2019.

[90] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. 2003.

[91] Sudheendra Vijayanarasimhan, Susanna Ricco, Cordelia Schmid, Rahul Sukthankar, and Katerina Fragkiadaki. Sfm-net: Learning of structure and motion from video. *arXiv preprint arXiv:1704.07804*, 2017.

[92] Davide Scaramuzza, Agostino Martinelli, and Roland Siegwart. A flexible technique for accurate omnidirectional camera calibration and structure from motion. In *Proceedings of the IEEE International Conference on Computer Vision Systems (ICVS)*, pages 45–45, 2006.

[93] Juho Kannala and Sami S Brandt. A generic camera model and calibration method for conventional, wide-angle, and fish-eye lenses. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(8):1335–1340, 2006.

[94] Bruno Caprile and Vincent Torre. Using vanishing points for camera calibration. *International Journal of Computer Vision*, 4(2):127–139, 1990.

[95] Marc Pollefeys and Luc Van Gool. A stratified approach to metric self-calibration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 407–412, 1997.

[96] Roberto Cipolla, Tom Drummond, and Duncan P Robertson. Camera calibration from vanishing points in image of architectural scenes. In *Proceedings of the British Machine Vision Conference (BMVC)*, pages 382–391, 1999.

[97] Oleksandr Bogdan, Viktor Eckstein, Francois Rameau, and Jean-Charles Bazin. DeepCalib: A deep learning approach for automatic intrinsic calibration of wide field-of-view cameras. In *Proceedings of the ACM SIGGRAPH European Conference on Visual Media Production*, 2018.

[98] Igor Vasiljevic, Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Wolfram Burgard, Greg Shakhnarovich, and Adrien Gaidon. Neural ray surfaces for self-supervised learning of depth and ego-motion. In *Proceedings of International Conference on 3D Vision (3DV)*, 2020.

[99] Gary Bradski and Adrian Kaehler. OpenCV. *Dr. Dobb's Journal of Software Tools*, 3, 2000.

[100] Joern Rehder, Janosch Nikolic, Thomas Schneider, Timo Hinzmann, and Roland Siegwart. Extending kalibr: Calibrating the extrinsics of multiple IMUs and of individual axes. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 4304–4311, 2016.

[101] J. Leonard, J. How, S. Teller, M. Berger, S. Campbell, G. Fiore, L. Fletcher, E. Frazzoli, A. Huang, S. Karaman, O. Koch, Y. Kuwata, D. Moore, E. Olson, S. Peters, J. Teo, R. Truax, M. Walter, D. Barrett, A. Epstein, K. Maheloni, K. Moyer, T. Jones, R. Buckley, M. Antone, R. Galejs, S. Krishnamurthy, and J. Williams. A perception-driven autonomous urban vehicle. *Journal of Field Robotics*, 25(10):727–774, October 2008.

[102] Chris Urmson, Joshua Anhalt, Drew Bagnell, Christopher Baker, Robert Bittner, MN Clark, John Dolan, Dave Duggins, Tugrul Galatali, Chris Geyer, et al. Autonomous driving in urban environments: Boss and the Urban Challenge. *Journal of Field Robotics*, 25(8):425–466, 2008.

[103] John G Fryer and Duane C Brown. Lens distortion for close-range photogrammetry. *Photogrammetric Engineering and Remote Sensing*, 52:51–58, 1986.

[104] Christopher Geyer and Kostas Daniilidis. A unifying theory for central panoramic systems and practical implications. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 445–461, 2000.

[105] Bogdan Khomutenko, Gaëtan Garcia, and Philippe Martinet. An enhanced unified camera model. *IEEE Robotics and Automation Letters*, 1(1):137–144, 2015.

[106] Scott Workman, Connor Greenwell, Menghua Zhai, Ryan Baltenberger, and Nathan Jacobs. DeepFocal: A method for direct focal length estimation. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, pages 1369–1373, 2015.

[107] Scott Workman, Menghua Zhai, and Nathan Jacobs. Horizon lines in the wild. *arXiv preprint arXiv:1604.02129*, 2016.

[108] Jiangpeng Rong, Shiyao Huang, Zeyu Shang, and Xianghua Ying. Radial lens distortion correction using convolutional neural networks trained with synthesized images. In *Proceedings of the Asian Conference on Computer Vision*, pages 35–49, 2016.

[109] Xiaoqing Yin, Xinchao Wang, Jun Yu, Maojun Zhang, Pascal Fua, and Dacheng Tao. FishEyeRectNet: A multi-context collaborative deep network for fisheye image rectification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 469–484, 2018.

[110] Manuel Lopez, Roger Mari, Pau Gargallo, Yubin Kuang, Javier Gonzalez-Jimenez, and Gloria Haro. Deep single image camera calibration with radial distortion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11817–11825, 2019.

[111] Yannick Hold-Geoffroy, Kalyan Sunkavalli, Jonathan Eisenmann, Matthew Fisher, Emiliano Gambaretto, Sunil Hadap, and Jean-François Lalonde. A perceptual measure for deep single image camera calibration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[112] Rui Zhu, Xingyi Yang, Yannick Hold-Geoffroy, Federico Perazzi, Jonathan Eisenmann, Kalyan Sunkavalli, and Manmohan Chandraker. Single view metrology in the wild. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 316–333, 2020.

[113] Chaoning Zhang, Francois Rameau, Junsik Kim, Dawit Mureja Argaw, Jean-Charles Bazin, and In So Kweon. DeepPTZ: Deep self-calibration for PTZ cameras. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1041–1049, 2020.

[114] Fabio Tosi, Filippo Aleotti, Pierluigi Zama Ramirez, Matteo Poggi, Samuele Salti, Luigi Di Stefano, and Stefano Mattoccia. Distilled semantics for comprehensive scene understanding from videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4654–4665, 2020.

[115] Yajie Zhao, Zeng Huang, Tianye Li, Weikai Chen, Chloe LeGendre, Xinglei Ren, Ari Shapiro, and Hao Li. Learning perspective undistortion of portraits. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 7849–7859, 2019.

[116] Shangrong Yang, Chunyu Lin, Kang Liao, Chunjie Zhang, and Yao Zhao. Progressively complementary network for fisheye image rectification using appearance flow. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6348–6357, 2021.

[117] Kang Liao, Chunyu Lin, and Yao Zhao. A deep ordinal distortion estimation approach for distortion rectification. *IEEE Transactions on Image Processing*, 30:3362–3375, 2021.

[118] Sudeep Pillai, Rares Ambrus, and Adrien Gaidon. Superdepth: Self-supervised, super-resolved monocular depth estimation. In *Robotics and Automation (ICRA), 2019 IEEE International Conference on*, 2018.

[119] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.

[120] Michael D Grossberg and Shree K Nayar. The raxel imaging model and ray-based calibration. *International Journal of Computer Vision*, 61(2):119–137, 2005.

[121] Jiading Fang, Igor Vasiljevic, Vitor Guizilini, Rares Ambrus, Greg Shakhnarovich, Adrien

Gaidon, and Matthew R Walter. Self-supervised camera self-calibration from video. *arXiv preprint arXiv:2112.03325*, 2021.

[122] Tali Treibitz, Yoav Schechner, Clayton Kunz, and Hanumant Singh. Flat refractive geometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(1):51–65, 2011.

[123] Varun Ravi Kumar, Sandesh Athni Hiremath, Stefan Milz, Christian Witt, Clement Pinnard, Senthil Yogamani, and Patrick Mader. Fisheyedistancenet: Self-supervised scale-aware distance estimation using monocular fisheye camera for autonomous driving. *arXiv preprint arXiv:1910.04076*, 2019.

[124] Srikumar Ramalingam and Peter Sturm. A unifying model for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(7):1309–1319, 2016.

[125] Srikumar Ramalingam, Peter Sturm, and Suresh K Lodha. Generic self-calibration of central cameras. *Computer Vision and Image Understanding*, 114(2):210–219, 2010.

[126] Filippo Bergamasco, Luca Cosmo, Andrea Gasparetto, Andrea Albarelli, and Andrea Torsello. Parameter-free lens distortion calibration of central cameras. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3847–3855, 2017.

[127] Pierre-André Brousseau and Sebastien Roy. Calibration of axial fisheye cameras through generic virtual central models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4040–4048, 2019.

[128] Robert Pless. Using many cameras as one. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pages II–587. IEEE, 2003.

[129] Peter Sturm, Srikumar Ramalingam, and Suresh Lodha. On calibration, structure from motion and multi-view geometry for generic camera models. In *Imaging Beyond the Pinhole Camera*, pages 87–105. Springer, 2006.

[130] Johannes Beck and Christoph Stiller. Generalized b-spline camera model. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 2137–2142. IEEE, 2018.

[131] Dennis Rosebrock and Friedrich M Wahl. Generic camera calibration and modeling using spline surfaces. In *2012 IEEE Intelligent Vehicles Symposium*, pages 51–56. IEEE, 2012.

[132] Sundar Vedula, Simon Baker, Peter Rander, Robert Collins, and Takeo Kanade. *Three-Dimensional Scene Flow*. Proceedings of the International Conference on Computer Vision (ICCV), 1999.

[133] Himangi Mittal, Brian Okorn, and David Held. Just go with the flow: Self-supervised scene flow estimation. *arXiv preprint arXiv:1912.00497*, 2019.

[134] Wenxuan Wu, Zhiyuan Wang, Zhuwen Li, Wei Liu, and Li Fuxin. Pointpwc-net: A coarse-to-fine network for supervised and self-supervised scene flow estimation on 3d point clouds. *arXiv preprint arXiv:1911.12408*, 2019.

[135] Miriam Schönbein and Andreas Geiger. Omnidirectional 3d reconstruction in augmented manhattan worlds. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 716–723. IEEE, 2014.

[136] Lipu Zhou, Jiamin Ye, Montiel Abello, Shengze Wang, and Michael Kaess. Unsupervised learning of monocular depth estimation with bundle adjustment, super-resolution and clip loss. *arXiv preprint arXiv:1812.03368*, 2018.

[137] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[138] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, 2018.

[139] Vincent Casser, Soeren Pirk, Reza Mahjourian, and Anelia Angelova. Unsupervised learning of depth and ego-motion: A structured approach. In *Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*, volume 2, page 7, 2019.

[140] François Chadebecq, Francisco Vasconcelos, George Dwyer, René Lacher, Sébastien Ourselin, Tom Vercauteren, and Danail Stoyanov. Refractive structure-from-motion through a flat refractive interface. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 5315–5323, 2017.

[141] Angelos Mallios, Eduard Vidal, Ricard Campos, and Marc Carreras. Underwater caves sonar data set. *The International Journal of Robotics Research*, 36(12):1247–1251, 2017.

[142] Myron Z Brown, Darius Burschka, and Gregory D Hager. Advances in computational stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(8):993–1008, 2003.

[143] Jure Zbontar, Yann LeCun, et al. Stereo matching by training a convolutional neural network to compare image patches. *Journal of Machine Learning Research*, 2016.

[144] Benjamin Attal, Selena Ling, Aaron Gokaslan, Christian Richardt, and James Tompkin. MatryODShka: Real-time 6DoF video view synthesis using multi-sphere images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 441–459, 2020.

[145] Tobias Bertel, Mingze Yuan, Reuben Lindroos, and Christian Richardt. OmniPhotos: casual 360° VR photography. *ACM Transactions on Graphics (TOG)*, 39(6):1–12, 2020.

[146] Hong-Xiang Chen, Kunhong Li, Zhiheng Fu, Mengyi Liu, Zonghao Chen, and Yulan Guo. Distortion-aware monocular depth estimation for omnidirectional images. *IEEE Signal Processing Letters*, 2021.

[147] Changhee Won, Jongbin Ryu, and Jongwoo Lim. OmniMVS: End-to-end learning for omnidirectional stereo matching. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 8987–8996, 2019.

[148] Varun Ravi Kumar, Marvin Klingner, Senthil Yogamani, Markus Bach, Stefan Milz, Tim Fingscheidt, and Patrick Mäder. Svdistnet: Self-supervised near-field distance estimation on surround view fisheye cameras. *ITS*, 2021.

[149] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. Deep-MVS: Learning multi-view stereopsis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[150] Sunghoon Im, Hae-Gon Jeon, Stephen Lin, and In So Kweon. DPSNet: End-to-end deep plane sweep stereo. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.

[151] Tejas Khot, Shubham Agrawal, Shubham Tulsiani, Christoph Mertz, Simon Lucey, and Martial Hebert. Learning unsupervised multi-view stereopsis via robust photometric consistency. *arXiv:1905.02706*, 2019.

[152] Yang Wang, Peng Wang, Zhenheng Yang, Chenxu Luo, Yi Yang, and Wei Xu. Unos: Unified unsupervised optical-flow and stereo-depth estimation by watching videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[153] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5974–5983, 2017.

[154] Jia-Wang Bian, Zhichao Li, Naiyan Wang, Huangying Zhan, Chunhua Shen, Ming-Ming Cheng, and Ian Reid. Unsupervised scale-consistent depth and ego-motion learning from monocular video. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[155] Ruihao Li, Sen Wang, Zhiqiang Long, and Dongbing Gu. Undeepvo: Monocular visual odometry through unsupervised deep learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 7286–7291, 2018.

[156] Sudeep Pillai, Rares Ambrus, and Adrien Gaidon. Superdepth: Self-supervised, super-resolved monocular depth estimation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2018.

[157] Maria Klodt and Andrea Vedaldi. Supervising the new with the old: Learning SFM from

SFM. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 698–713, 2018.

[158] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[159] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4104–4113, 2016.

[160] Lahav Lipson, Zachary Teed, and Jia Deng. RAFT-Stereo: Multilevel recurrent field transforms for stereo matching. In *Proceedings of International Conference on 3D Vision (3DV)*, pages 218–227, 2021.

[161] Zachary Teed and Jia Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.

[162] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. *arXiv preprint arXiv:2111.11426*, 2021.

[163] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

[164] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. 2020.

[165] Zhaoshuo Li, Xingtong Liu, Nathan Drenkow, Andy Ding, Francis X Creighton, Russell H Taylor, and Mathias Unberath. Revisiting stereo depth estimation from a sequence-to-sequence perspective with transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 6197–6206, 2021.

[166] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12179–12188, 2021.

[167] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12873–12883, 2021.

[168] Robin Rombach, Patrick Esser, and Björn Ommer. Geometry-free view synthesis: Transformers and no 3D priors. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 14356–14366, 2021.

[169] Mehdi SM Sajjadi, Henning Meyer, Etienne Pot, Urs Bergmann, Klaus Greff, Noha Radwan, Suhani Vora, Mario Lucic, Daniel Duckworth, Alexey Dosovitskiy, et al. Scene representation transformer: Geometry-free novel view synthesis through set-latent scene representations. *arXiv preprint arXiv:2111.13152*, 2021.

[170] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5410–5418, 2018.

[171] Zachary Teed and Jia Deng. Raft-3D: Scene flow using rigid-motion embeddings. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8375–8384, 2021.

[172] Zachary Teed and Jia Deng. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in Neural Information Processing Systems*, 34, 2021.

[173] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

[174] Benjamin Ummenhofer, Huizhong Zhou, Jonas Uhrig, Nikolaus Mayer, Eddy Ilg, Alexey Dosovitskiy, and Thomas Brox. Demon: Depth and motion network for learning monocular

stereo. In *IEEE Conference on computer vision and pattern recognition (CVPR)*, volume 5, page 6, 2017.

[175] Huizhong Zhou, Benjamin Ummenhofer, and Thomas Brox. DeepTAM: Deep tracking and mapping. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 822–838, 2018.

[176] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *ACM Transactions on Graphics (TOG)*, 39(4), 2020.

[177] Johannes Kopf, Xuejian Rong, and Jia-Bin Huang. Robust consistent video depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1611–1621, 2021.

[178] Xiaoxiao Long, Lingjie Liu, Wei Li, Christian Theobalt, and Wenping Wang. Multi-view depth estimation using epipolar spatio-temporal networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8258–8267, 2021.

[179] Jiaming Sun, Yiming Xie, Linghao Chen, Xiaowei Zhou, and Hujun Bao. NeuralRecon: Real-time coherent 3D reconstruction from monocular video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15598–15607, June 2021.

[180] Dejan Azinović, Ricardo Martin-Brualla, Dan B Goldman, Matthias Nießner, and Justus Thies. Neural RGB-D surface reconstruction. *arXiv preprint arXiv:2104.04532*, 2021.

[181] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised NeRF: Fewer views and faster training for free. *arXiv preprint arXiv:2107.02791*, 2021.

[182] Konstantinos Rematas, Andrew Liu, Pratul P Srinivasan, Jonathan T Barron, Andrea Tagliasacchi, Thomas Funkhouser, and Vittorio Ferrari. Urban radiance fields. *arXiv preprint arXiv:2111.14643*, 2021.

[183] Yi Wei, Shaohui Liu, Yongming Rao, Wang Zhao, Jiwen Lu, and Jie Zhou. NerfingMVS: Guided optimization of neural radiance fields for indoor multi-view stereo. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.

[184] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. NICE-SLAM: Neural implicit scalable encoding for SLAM. *arXiv preprint arXiv:2112.12130*, 2021.

[185] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[186] Connor Shorten and Taghi M Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):1–48, 2019.

[187] David Novotny, Diane Larlus, and Andrea Vedaldi. Learning [3d] object categories by looking around them. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 5218–5227, 2017.

[188] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. KinectFusion: Real-time dense surface mapping and tracking. In *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2011.

[189] Chengzhou Tang and Ping Tan. BA-Net: Dense bundle adjustment network. *arXiv preprint arXiv:1806.04807*, 2018.

[190] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

[191] Chao Liu, Jinwei Gu, Kihwan Kim, Srinivasa G Narasimhan, and Jan Kautz. Neural RGB $\rightarrow$ D sensing: Depth and uncertainty from a video camera. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10986–10995, 2019.

[192] Xiaoxiao Long, Lingjie Liu, Christian Theobalt, and Wenping Wang. Occlusion-aware depth estimation with adaptive normal constraints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 640–657, 2020.

[193] Vitor Guizilini, Kuan-Hui Lee, Rareş Ambruş, and Adrien Gaidon. Learning optical flow, depth, and scene flow without real-world labels. *IEEE Robotics and Automation Letters*, 7 (2):3491–3498, 2022.

[194] Vitor Guizilini, Rares Ambrus, Dian Chen, Sergey Zakharov, and Adrien Gaidon. Multi-frame self-supervised depth with transformers. *arXiv preprint arXiv:2204.07616*, 2022.

[195] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, W Yifan, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. Advances in neural rendering. In *Computer Graphics Forum*, volume 41, pages 703–735. Wiley Online Library, 2022.

[196] Vincent Sitzmann, Semon Rezchikov, Bill Freeman, Josh Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.