

NOCaL: Calibration-Free Semi-Supervised Learning of Odometry and Camera Ininsics

Ryan Griffiths, Jack Naylor, Donald G. Dansereau

Abstract—There are a multitude of emerging imaging technologies that could benefit robotics. However the need for bespoke models, calibration and low-level processing represents a key barrier to their adoption. In this work we present NOCaL, Neural Odometry and Calibration using Light fields, a semi-supervised learning architecture capable of interpreting previously unseen cameras without calibration. NOCaL learns to estimate camera parameters, relative pose, and scene appearance. It employs a scene-rendering hypernetwork pre-trained on a large number of existing cameras and scenes, and adapts to previously unseen cameras using a small supervised training set to enforce metric scale. We demonstrate NOCaL on rendered and captured imagery using conventional cameras, demonstrating calibration-free odometry and novel view synthesis. This work represents a key step toward automating the interpretation of general camera geometries and emerging imaging technologies. Code and datasets are available at <https://roboticimaging.org/Projects/NOCaL/>.

I. INTRODUCTION

Vision is a critical sense in modern robotics. Enormous advancements have been made in recent years to better utilise cameras for almost every task a robotic platform might be asked to do. Looking to nature, it is clear that the best way to see the world depends on what needs to be seen. In the same way, the best camera is specific to the application and domain in which a robot is operating. Whilst novel cameras are being developed to address shortcomings in existing modalities, this raises a key problem in deploying these sensors quickly and on new platforms: calibration and low-level interpretation.

Calibrating and interpreting new imaging devices is skilled and time-consuming work. Emerging devices like event cameras and light field cameras have taken years and even decades to adapt in robotics. Solutions generally involve the use of bespoke models, calibration procedures, and low-level interpretation and sensor fusion algorithms. An assumption is generally made of static camera characteristics, with device changes due to vibration, temperature, replacement or upgrading requiring re-calibration as these changes can affect performance [1], [2]. This makes both integrating new cameras and managing fleets of robots onerous and complex.

In this work, we propose a framework to automatically interpret previously unseen cameras by jointly learning to estimate novel views, odometry, and camera parameters – see Fig. 1. Our framework utilises recent advancements in neural rendering to provide self-supervision, leveraging the

The authors are with the Australian Centre For Robotics, School of Aerospace, Mechanical and Mechatronic Engineering, The University of Sydney. r.griffiths, jack.naylor, donald.dansereau@sydney.edu.au

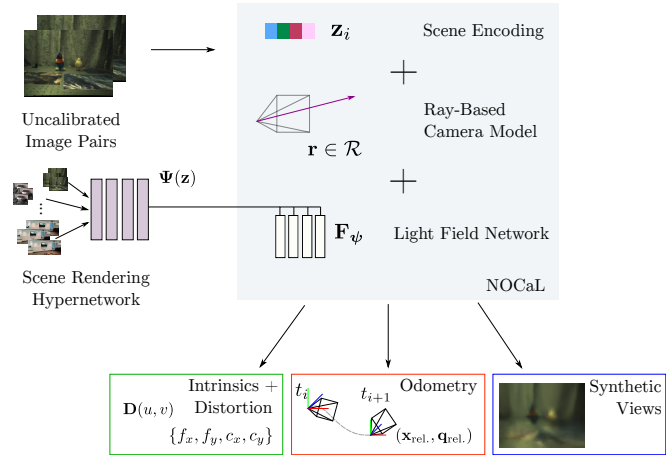


Fig. 1. NOCaL learns odometry and novel view synthesis from previously unseen cameras by leveraging a rendering hypernetwork pre-trained on existing cameras. Uncalibrated input image pairs are encoded in a latent space z_i that drives the hypernetwork $\Psi(z)$ to generate a light field rendering network F_ϕ . A differentiable ray-based camera model drives the renderer, enabling estimation of camera parameters, relative pose, and novel views. This work is a key step towards automatically interpreting more general camera geometries and emerging camera technologies.

large amount of imagery available from a newly introduced uncalibrated camera. To benefit from the availability of unlabelled training data from existing cameras, we employ a hypernetwork that learns to construct light field renderers, so that the hypernetwork can be trained on multiple scenes. This would not be possible with a fixed rendering network as this could only be trained on a single scene. Finally, to ground our odometry estimates in metric space, we employ a small labelled training set that we show is easy to collect where a complementary source of odometry is available.

Through unsupervised learning of camera parameters and odometry, our approach benefits from the large amount of unlabelled data available from existing cameras as well as a newly introduced camera. By introducing a small labelled training set we constrain the solution to a metric space with known scale, while there is always some ambiguity with monocular scale estimation, using this approach allows the network to use semantic information to be able to get reliable scale when working in a similar environment to that of the training data [3].

To demonstrate our approach, we employ cameras which are well described by a pinhole projection model with freeform ray-based distortion model. Allowing use of a broad range of monocular cameras without calibration. In future, we envision further relaxing this model to an entirely

freeform ray-based model, leading to a broader range of cameras including stereo, multi-aperture and light field, and fisheye, that are all well described by ray-based geometry.

We validate our approach on both captured and rendered images of indoor scenes, using cameras with different focal lengths and distortion parameters. We demonstrate our system accurately estimating camera intrinsics, distortion models, and relative pose, i.e. odometry.

To position our work, we compare against a fully supervised approach and an unsupervised approach that requires calibration. Perhaps surprisingly, our semi-supervised but uncalibrated approach outperforms both in accuracy of odometry, demonstrating the strength of combining small labelled datasets with readily available unlabelled data. We also include an ablation study that establishes the importance of the distortion model when using cameras that deviate substantially from an ideal pinhole model, comparing against variants of our method that lack distortion model and that estimate no camera parameters at all.

We anticipate this work to find broad applicability where recalibration is difficult and camera parameters can change, either due to optical shifts or replacing of hardware. Deployed systems on planetary missions, in harsh environments, and in domestic applications like robotic vacuums for example are typically difficult to recalibrate. Changes to on-board calibration can occur due to vibration and thermal effects, and replacing or upgrading cameras can be an expensive proposition especially where new camera models and/or recalibration are required. Our framework requires no prior knowledge of hardware or camera parameters, allowing such robots to perform accurate camera pose estimation without need for recalibration or manual intervention.

Limitations: Whilst our network is designed to work with a family of cameras, cameras not well-described by a pinhole projection and freeform distortion profile are unlikely to be well supported by our camera model. We anticipate generalising the camera model in future. Pose estimation requires substantial overlap between the input images, and so the approach also breaks down for fast motion / low-overlap input pairs.

II. RELATED WORK

State-of-the-art approaches to monocular visual odometry jointly learn scene depth and odometry of images using unsupervised learning [4]–[8]. This generally uses a warp function to predict images in a sequence based on depth and estimated pose. This warping usually requires accurately calibrated images which are often not available on robotic platforms that are deployed in harsh environments. More recent work by Fang et al. [2] and Gordon et al. [9] have also been able to jointly learn the camera model, which alleviates one of the main difficulties with this approach. Ultimately these approaches still use a warp function which will limit the types of camera geometries that can be learnt using this method. Warp based approaches also have no ability to deal with view dependent objects, such as non-Lambertian

surfaces. This is something that ray based methods, such as the one presented in this work, can accommodate.

Digumarti et al. [9] demonstrated the performance of warping using a novel 4D warp function by extending it to a new family of cameras: sparse light-field cameras. In using a warp function, these studies are limited to scenes with simple well-explained phenomena, for example these methods cannot handle view dependent phenomena (reflection, refraction).

Recent studies on neural novel view synthesis [10] have demonstrated state of the art results, with applications to many computer vision applications. While this has been adapted to robotics in [11]–[14], the fundamental limitation of such approaches exists in being only able to represent a specific scene, and within the spatial region captured by the input data. The computation required to ray-march is expensive, but provides a dense and continuous scene representation. Instant-NGP [15] addresses this limitation to an extent by employing a multi-resolution hash encoding, which drastically speeds up training and inference, but leads to more sparse geometry.

Light Field Networks (LFNs) [16] perform a single query of the network unlike prior works, enabling a reduction in training and inference time by several orders of magnitude compared to NeRF [10]. Recent studies [17]–[20], leverage LFNs to produce comparable results to that of NeRFs, with tradeoffs between visual fidelity and speed.

Back propagation of gradients through a neural field MLP provides an ability for networks to refine parameters such as pose. While bundle-adjusting radiance fields [21] refines poses during the formation of a radiance field, using a NeRF as a supervisory signal for absolute pose regression [22] shows advantages in accuracy around convex and extended scenes. NeRFs do not perform well on few-shot datasets and require dense image coverage of a scene to create high fidelity results, however once generated it is possible to perform accurate pose regression on a minimal dataset [23]. Of key note is that the refinement or determination of pose can be performed in addition to estimation of other parameters within the neural field, such as shape, reflectance functions or illumination [24].

Joint learning of camera intrinsics and neural fields show improved extrinsics estimation. Wang et al. [25] demonstrate an ability to jointly learn focal length and extrinsics, achieving similar results to traditional methods like COLMAP [26].

Jeong et al. [27] jointly learn a complex non-linear distortion camera model with the neural field in several stages, allowing the framework to learn a simple pinhole model followed by complex components including non-linear distortion parameters. In essence, this curriculum learning approach enables the network to obtain correct scene geometry without warping the scene to agree with camera distortion.

Hypernetworks [28] allow one network to produce weights for other networks that can perform additional tasks. The framework presented by von Oswald et al. [29] has the ability to retain a vast amount of memory for multitask learning using a hypernetwork. This benefits the individual networks




Fig. 2. The proposed NOCaL network architecture is made up of 4 separate sub-modules. The camera interpreter (green) encodes input frames into a scene-description latent space \mathbf{z} . The scene rendering hypernetwork (orange) uses the latent space to produce weights for the light field network (blue) which renders the scene captured by the input images. The camera model (yellow) has learnable parameters to be able to estimate the camera intrinsics used in capturing the input frames, and generates rays used by the light field network.

and allows for reuse, owing to the commonalities between the learning tasks.

Sitzmann et al. [16] also employs hypernetworks to leverage the ability to learn latent-based embeddings to produce a rendering network that can render the specific scene represented by latent embedding. Unlike prior works which focus on single scene, this network has the ability to render multiple scenes based on a single image.

III. METHOD

A. Network Architecture

The two main learnable parts of the proposed network are the encoding network and the hypernetwork. These two parts work together to be able to jointly learn pose and scene geometry from input images. Details for implementation are discussed in Section IV-B.

1) *Encoding Network*: The camera interpreter portion of our pipeline serves as an encoding network, which converts a pair of input frames from a specific camera to a pose and a latent space which represents the scene. This encoding network is built using a CNN structure with two separate heads at the end. The pose head outputs 3 values corresponding to translation and 6 values representing a continuous rotation space, of the form proposed by [30]. The scene head outputs a 256-dimensional latent space \mathbf{z}_i which contains the information required to be able to build a rendering of the scene that the two input frames comes from.

2) *Hypernetwork*: In this architecture the hypernetwork Ψ uses the latent space \mathbf{z}_i to output the weights for the rendering network ψ_i ,

$$\Psi(\mathbf{z}_i) = \psi_i. \quad (1)$$

Here the encoder has already distilled scene-specific information from the images, which ideally is independent of the camera geometry. The operation of the rendering network is on the level of light rays, requiring no camera model to generate new images. In this way, the hypernetwork is able to be used as a tool for multiple cameras, giving the rendering network an initialisation which may be used to train the extrinsics and intrinsics of unknown vision sensors.

3) *Neural Fields for Supervision*: We utilise the rendering from the light field network to train pose and camera intrinsics. There are a few key benefits with using view synthesis from a neural light field. Firstly, it is ray-based, providing a general model for all cameras, and resulting in novel views of sufficient visual fidelity ideal for use as supervision of odometry. Secondly, the implementation is fully differentiable which allows for an end to end system to be developed in which the input image into the neural field can be learnt. We utilise this second notion to learn camera parameters through a differentiable camera model.

The chosen light field network approximates a continuous scene in the form of an MLP $\mathbf{F}_\psi : (\mathbf{o}, \mathbf{d}) \mapsto \mathbf{c}$ with weights ψ . This formulation uniquely maps the ray direction \mathbf{d} through some origin \mathbf{o} using a Plücker coordinate encoding to a colour \mathbf{c} . As noted by Sitzmann et al. [16], whilst providing a compact and unique encoding, complex phenomena such as occlusion are not readily dealt with. The utilisation of a preceding frame to inform the hypernetwork and sequential camera motion enables the rendering network to avoid this shortcoming by evolving the scene representation over a trajectory. For this work we propose using a LFN over alternatives because of its speed advantages.




Fig. 3. Model of camera parameter estimation. The pixel coordinates (u, v) and the intrinsics K are first used to calculate the distorted points on the image plane (x_d, y_d) . A distortion network takes as input (u, v) and determines the distortion values on a per pixel basis and produces a Δx and Δy . These corrections are applied to the points on the image plane producing undistorted points (x, y) , which are converted to rays in world space by the homogenous transformation $T_{C \rightarrow W}^{-1}$.

B. Camera Modelling

Camera parameters are estimated in two parts: intrinsics, consisting of focal lengths (f_x, f_y) and principal points (c_x, c_y) , and a distortion model. Together these describe a large family of cameras, except those not well described by a pinhole projection.

1) *Focal Length Estimation*: The estimation of the focal length is performed through back-propagation of the rendering MLP. Setting the focal length as a tuneable parameters that can be optimised allows the network to change the physical model of the camera as it generates scene geometry. We found through experimentation that the camera model was fairly robust to intrinsics initialisation, however if initialised outside a sensible range than the focal would get stuck in a local minima. To ensure convergence of the focal length it is initialised as the width of the image in pixels following [25], this will typically place the focal length within a suitable error margin.

Given focal length is directly correlated to the scale of the geometry seen on the sensor, and we seek to jointly learn a representation of the scene and the camera values, small changes to focal length compared to geometry can trap the network in local minima. To this end, we use a higher learning rate to converge camera parameters prior to the network learning substantial scene geometry.

2) *Implicit Non-Linear Distortion*: To deal with generality of cameras, and to avoid the limitations of any single camera model, we model the non-linear distortion using an MLP, $\mathbf{D}(u, v) = (\Delta x, \Delta y)$. The MLP determines a *correction* to coordinates on the camera plane using pixel coordinates (u, v) . This pixel coordinate-based MLP effectively is modelling a differentiable, smooth and continuous distortion function of sufficient complexity to encompass most cameras covered by a pinhole model. Figure 3 shows a diagrammatic representation of the MLP based distortion function in conjunction with the camera model.

C. Semi-Supervised Learning

In previous unsupervised learning of odometry work the camera parameters were required to be known and images undistorted prior to entering the pipeline [4]. In this work we

proposed learning the camera parameters in addition to relative pose within the pipeline. This adds a substantial degree of complexity to the network, which is less constrained.

To reintroduce some constraints, a small amount of labelled data is used to semi-supervise the network. This allows us to directly impose a loss on the encoder, instead of having to backpropagate through the hypernetwork and allows us to confine the pose to metric terms. Previous unsupervised works had to scale the results after training to recover a metric scale [4]. By imposing the learning of the camera parameters, we enable direct recovery of relative pose by avoiding scale ambiguity.

The trade-off between needing to know the camera parameters and needing a small amount of labelled data is often a preferable one. Using a platform such as a robotic arm enables a set of ground truth poses to be acquired irrespective of camera installed on-board along a pre-defined trajectory. Novel cameras may require extensive processes to acquire accurate data for direct calibration, which may be costly to obtain in large volumes. Using a small amount of this data, it enables the network to extend automatically and generalise to new calibrations.

We demonstrate that this semi-supervised approach outperforms both a fully supervised and unsupervised approach. See Table. II for results.

D. Training Losses

Similar to other works in neural rendering, we employ a photometric loss term \mathcal{L}_{ph} as the primary loss function of our network between ground truth \mathbf{c} and predicted $\hat{\mathbf{c}}$ pixel colours. This is calculated for all rays $r \in \mathcal{R}$, where \mathcal{R} is the set of rays captured by an image,

$$\mathcal{L}_{\text{ph}} = \sum_{r \in \mathcal{R}} |\mathbf{c} - \hat{\mathbf{c}}|_2^2. \quad (2)$$

Where images have labelled poses during training, denoted as \mathcal{I}' , we enforce simple L_2 -norms between the translations $\mathbf{x} \in \mathbb{R}^3$ and rotation matrices $\mathbf{R} \in \text{SO}(3)$,

$$\mathcal{L}_{\text{trans}} = \sum_{\mathcal{I}'} |\mathbf{x} - \hat{\mathbf{x}}|_2^2, \quad (3)$$

$$\mathcal{L}_{\text{rot}} = \sum_{\mathcal{I}'} |\mathbf{R} - \hat{\mathbf{R}}|_2^2. \quad (4)$$

Finally, we encourage the latent space to have a mean of zero by assuming a Gaussian prior [16],

$$\mathcal{L}_{\text{enc}} = \sum_{\mathcal{I}} \text{mean}(\mathbf{z}). \quad (5)$$

The overall loss function hence encompasses a loss from rendering, any available pose supervision and an imposed constraint to the latent space

$$\mathcal{L} = \lambda_{\text{ph}} \mathcal{L}_{\text{ph}} + \lambda_{\text{trans}} \mathcal{L}_{\text{trans}} + \lambda_{\text{rot}} \mathcal{L}_{\text{rot}} + \lambda_{\text{enc}} \mathcal{L}_{\text{enc}}. \quad (6)$$

E. Curriculum Learning

Given the challenges of jointly estimating scene and camera parameters, we employ a curriculum learning approach to sequentially recover camera parameters within an evolving neural scene. Initially, the encoding and hypernetwork are trained with fixed initial camera intrinsics, providing a rough low-frequency representation of the scene. This rough representation provides sufficient geometry to begin supervising the camera model. Prior to the geometry being fully converged, we enable the tuning of focal length in a simple pinhole model, allowing for adjustment of scene scale. Finally, the full implicit distortion model is added, giving a metric and geometrically representative scene representation, and providing a camera calibration. Learning in this way lets the network avoid local minima as higher frequency scene information is learnt.

IV. RESULTS

A. Datasets

We demonstrate the results for the proposed system on both real and synthetic data, showing the system working for multiple cameras and scenes. The real world dataset used was part of the LearnLFOdo Dataset [9]. While this dataset is captured with a light field camera and this work is focused on monocular imaging, the dataset is the same as monocular when only the centre image is taken. This dataset contains 45 separate scenes and camera trajectories using the same camera captured by a UR5e robotic arm, providing accurate ground truth data to test the generalisation capabilities of NOCaL to new scenes.

The rendered dataset was trained separately to provide comparison between multiple cameras with defined intrinsics and distortion. This allowed for repeatable scene configurations and trajectories with multiple cameras of different distortion values.

B. Implementation Details

Both the hypernetwork and LFN are 6-layer MLPs with ReLU activations. The hypernetwork has 256 units per layer to allow for enhanced generalisation, while the LFN has 128 units. To encode the input image pairs a 7-layer CNN with kernel size of 3 is selected to maintain resolution of features. The pose head is a 3-layer CNN of 1x1 convolutions, while the scene head is a 3-layer fully connected network of 256 unit wide layers. Finally, the distortion model is an 8-unit wide MLP with 4-layers. We select the following hyperparameters, as laid out in Eqn. 6, $\lambda_{\text{ph}} = 100$, $\lambda_{\text{trans}} = 30$, $\lambda_{\text{rot}} = 20$, $\lambda_{\text{enc}} = 1 \times 10^{-6}$.

Four separate learning rates were used to train the separate parts of the model, one for each of the separate sub-modules; the intrinsics, the distortion, the hyper rendering network, and the encoding network. As the network jointly optimises the separate sections if one trains faster than another, it can lead to modules getting stuck in a local minima. Weighting the separate parts through the learning rates is critical to ensure the framework as a whole converges appropriately. The chosen learning rates were, 5×10^{-5} , 8×10^{-5} , 5×10^{-6} ,




Fig. 4. View synthesis from the learnt LFN on a test scene. A pair of input images at times i , $i+1$ are shown on the left, with the corresponding predicted views on the right closely matching in visual appearance. The temporal difference frames shows the motion through time, it can be seen that the predicted and measured motion are very similar. Animations of this motion will be made available on the project website.

5×10^{-1} for the encoder, hypernetwork, distortion model and intrinsics respectively. Adam optimisers were used for all separate learning rates. All layers were initialised with an Xavier uniform distribution [31].

C. Scene Reconstruction

As shown in Fig. 4, the framework is able to produce novel views of scenes it has not been trained on. Given a pair of inputs with some motion between frames, a pair of predicted frames can be retrieved from the network. As these views are used as the supervisory signal for the rest of the network, the temporal difference or motion between the predicted frames should reflect the same motion between the input frames. This is the case in Fig. 4, which indicates that the network can be supervised with this signal. The reconstructions have lost some of the high frequency scene content, however the reconstruction quality is not critical, but rather how well the renderings can supervise the motion between the frames.

D. Camera Modeling

NOCaL is able to recover accurate camera intrinsics, close to ground truth and those attained by traditional methods. Table I demonstrates an ability to recover comparable results to COLMAP [26] in the absence of distortion, validating the case of an ideal pinhole camera by a significant reduction in error. We compare results based on the mean radial shift per pixel $\Delta r = \sqrt{\Delta x^2 + \Delta y^2}$. We sample the radial distortion function used by COLMAP on a grid to obtain a comparable result owing to how the distortion network is formulated.

We note that while our method is able to significantly reduce the error in the case of large distortion, the fixed camera model used by COLMAP provides a better approximation to the continuous radial distortion profile. As we are also jointly learning focal length, some error in focal length is taken up within the distortion network leading to additional error. Additional constraints could be imposed on the implicit distortion model in the future, though we note this did not affect achieved odometry performance.

TABLE I
EVALUATING CAMERA PARAMETER ESTIMATION

Camera	Initial		NOCaL (ours)		COLMAP [26]	
	f [px]	Error $\overline{\Delta r}$	f [px]	Error $\overline{\Delta r}$	f [px]	Error $\overline{\Delta r}$
Focal: 600px, no distortion	640.0	0.0225	601.1	0.0008	599.4	0.0003
Focal: 600px, large distortion	640.0	0.0440	595.3	0.0186	602.1	0.0010

TABLE II
EVALUATING ODOMETRY PERFORMANCE ON CAPTURED AND RENDERED IMAGERY

Method	Labelled Images	Unlabelled Images	Translation Error [m]			Rotation Error [degrees]		
			Mean	STD	RMSE	Mean	STD	RMSE
Odometry accuracy on captured indoor imagery								
Fully supervised	800	0	0.025	0.009	0.027	1.553	1.847	2.414
Unlabelled calibrated [4]	0	8000	0.029	0.016	0.033	1.522	0.969	1.808
NOCaL (ours)	800	7200	0.020	0.008	0.022	0.412	0.295	0.505
Ablation study using rendered indoor imagery with camera distortion								
Ours no intrinsics or distortion	100	900	0.157	0.060	0.168	8.026	9.180	12.194
Ours no distortion	100	900	0.147	0.053	0.156	4.790	2.209	5.275
Ours full	100	900	0.145	0.054	0.154	4.024	1.971	4.481

E. Odometry Results

Odometry results for NOCaL are shown in Table II. We compared NOCaL to two other odometry methods: a fully supervised approach with labelled imagery, and an unsupervised approach based on [4] that requires the camera to be calibrated and imagery rectified. The unsupervised method was provided with similar numbers of unlabelled images, around 8000, representative of the availability of unlabelled imagery in practical scenarios.

While the unsupervised approach did not require any labelled data, it is scale ambiguous, with the results needing to be correctly scaled before an error can be calculated. NOCaL does not require such scaling as the labelled data during training establish scale based on semantic information in the image. Furthermore our framework does not require camera calibration or rectification.

Perhaps surprisingly, our zero-calibration approach outperformed both fully supervised and calibrated unsupervised methods. This is partially explained by the amount of training data available to each method: NOCaL and the fully supervised approach were provided with 800 labelled images, but NOCaL also has the benefit of additional unlabelled imagery. In the case of the unsupervised calibrated method, we hypothesise that our freeform distortion model did a better job of describing the camera distortion compared to the parametric model employed in typical camera calibration.

We performed an ablation study to measure the effectiveness of our camera model – the results are shown at the bottom of Table II. We tested the full NOCaL, NOCaL without a distortion model, and a version that does not adjust the camera intrinsics or distortion model. This study employed rendered imagery simulating a camera with realistic and known radial distortion. Intrinsics were initialised close to but not exactly matching correct values, in line with a typical imaging scenario. The study shows both the distortion model and intrinsic refinement play an important role in NOCaLs strong odometry performance. The presented error metrics were calculated using [32].

F. Training and Inference Time

Typical training time for NOCaL on the LFodo dataset [9] was approximately 1.5 hours. Training time was tempered by use of the LFN for rendering and use of down-sampled training images. As the rendering and the odometry can be uncoupled, inference is performed on the odometry network alone, yielding an inference time of 16.9 ms, compatible with real-time applications. Performing inference of the full framework takes 33.9 ms. The networks were trained and timed on an NVIDIA RTX 3060 12GB GPU.

V. CONCLUSIONS

We presented NOCaL, a framework that jointly learns odometry, camera parameters and visual appearance in an end to end fashion. This framework generalises to a large range of cameras that can be modelled with a pinhole projective model and freeform distortion map. We demonstrated our method learning odometry on previously unseen and uncalibrated cameras, and an ablation study established the importance of learning camera parameters for this task. Our approach requires only a small amount of labelled data, allowing it to provide metric results with scale, and benefits from the availability of large amounts of unlabelled data from both a newly introduced camera and the vast quantities of existing unlabelled data from existing cameras.

Future work will entail extending the camera model to work for more general camera geometries, including fish-eye lenses and multi-aperture cameras. Deployment on robotic platforms undergoing extended operation with online updates to the camera model represents a logical next step in validating the approach. Finally, we aim to extend the work to support time of flight, event, and more general computational imaging devices, working towards truly autonomous integration of emerging camera technologies.

ACKNOWLEDGEMENTS

This research was supported in part by funding from Ford Motor Company.

We would like to thank the reviewers for their time and valuable feedback, which helped us improve this work.

REFERENCES

- [1] M. Smith and E. Cope, "The effects of temperature variation on single-lens-reflex digital camera calibration parameters," *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 38, no. Part 5, 2010.
- [2] J. Fang, I. Vasiljevic, V. Guizilini, R. Ambrus, G. Shakhnarovich, A. Gaidon, and M. R. Walter, "Self-supervised camera self-calibration from video," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 8468–8475.
- [3] J. Jiao, Y. Cao, Y. Song, and R. Lau, "Look deeper into depth: Monocular depth estimation with semantic booster and attention-driven loss," in *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [4] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised Learning of Depth and Ego-motion from Video," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1851–1858.
- [5] L. Han, Y. Lin, G. Du, and S. Lian, "DeepVIO: Self-supervised Deep Learning of Monocular Visual Inertial Odometry using 3D Geometric Constraints," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 6906–6913.
- [6] Z. Yin and J. Shi, "GeoNet: Unsupervised Learning of Dense Depth, Optical Flow and Camera Pose," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1983–1992.
- [7] H. Zhan, R. Garg, C. S. Weerasekera, K. Li, H. Agarwal, and I. Reid, "Unsupervised Learning of Monocular Depth Estimation and Visual Odometry with Deep Feature Reconstruction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 340–349.
- [8] C. Godard, O. Mac Aodha, and G. J. Brostow, "Unsupervised Monocular Depth Estimation with Left-Right Consistency," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 270–279.
- [9] S. T. Digumarti, J. Daniel, A. Ravendran, R. Griffiths, and D. G. Dansereau, "Unsupervised Learning of Depth Estimation and Visual Odometry for Sparse Light Field Cameras," in *Intelligent Robots and Systems (IROS)*. IEEE, 2021.
- [10] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [11] M. Adamkiewicz, T. Chen, A. Caccavale, R. Gardner, P. Culbertson, J. Bohg, and M. Schwager, "Vision-only Robot Navigation in a Neural Radiance World," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4606–4613, 2022.
- [12] E. Sucar, S. Liu, J. Ortiz, and A. Davison, "iMAP: Implicit Mapping and Positioning in Real-Time," in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.
- [13] J. Ichnowski*, Y. Avigal*, J. Kerr, and K. Goldberg, "Dex-NeRF: Using a Neural Radiance field to Grasp Transparent Objects," in *Conference on Robot Learning (CoRL)*, 2020.
- [14] L. Yen-Chen, P. Florence, J. T. Barron, T.-Y. Lin, A. Rodriguez, and P. Isola, "NeRF-Supervision: Learning Dense Object Descriptors from Neural Radiance Fields," in *IEEE Conference on Robotics and Automation (ICRA)*, 2022.
- [15] T. Müller, A. Evans, C. Schied, and A. Keller, "Instant Neural Graphics Primitives with a Multiresolution Hash Encoding," *ACM Trans. Graph.*, vol. 41, no. 4, pp. 102:1–102:15, July 2022. [Online]. Available: <https://doi.org/10.1145/3528223.3530127>
- [16] V. Sitzmann, S. Rezkikov, W. T. Freeman, J. B. Tenenbaum, and F. Durand, "Light Field Networks: Neural Scene Representations with Single-Evaluation Rendering," in *Proc. NeurIPS*, 2021.
- [17] B. Attal, J.-B. Huang, M. Zollhöfer, J. Kopf, and C. Kim, "Learning Neural Light Fields with Ray-Space Embedding Networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [18] Z. Li, L. Song, C. Liu, J. Yuan, and Y. Xu, "NeuLF: Efficient Novel View Synthesis with Neural 4D Light Field," in *Eurographics Symposium on Rendering*, A. Ghosh and L.-Y. Wei, Eds. The Eurographics Association, 2022.
- [19] M. Suhail, C. Esteves, L. Sigal, and A. Makadia, "Light Field Neural Rendering," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 8269–8279.
- [20] P. Wang, Y. Liu, G. Lin, J. Gu, L. Liu, T. Komura, and W. Wang, "Progressively-connected Light Field Network for Efficient View Synthesis," *arXiv preprint arXiv:2207.04465*, 2022.
- [21] C.-H. Lin, W.-C. Ma, A. Torralba, and S. Lucey, "BARF: Bundle-Adjusting Neural Radiance Fields," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5741–5751.
- [22] S. Chen, Z. Wang, and V. Prisacariu, "Direct-PoseNet: Absolute Pose Regression with Photometric Consistency," in *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, pp. 1175–1185.
- [23] A. Moreau, N. Piasco, D. Tsishkou, B. Stanculescu, and A. de La Fortelle, "LENS: Localization enhanced by NeRF synthesis," in *Conference on Robot Learning*. PMLR, 2022, pp. 1347–1356.
- [24] M. Boss, A. Engelhardt, A. Kar, Y. Li, D. Sun, J. T. Barron, H. P. Lensch, and V. Jampani, "SAMURAI: Shape And Material from Unconstrained Real-world Arbitrary Image collections," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [25] Z. Wang, S. Wu, W. Xie, M. Chen, and V. A. Prisacariu, "NeRF--: Neural Radiance Fields Without Known Camera Parameters," *arXiv preprint arXiv:2102.07064*, 2021.
- [26] J. L. Schönberger and J.-M. Frahm, "Structure-from-Motion Revisited," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [27] Y. Jeong, S. Ahn, C. Choy, A. Anandkumar, M. Cho, and J. Park, "Self-Calibrating Neural Radiance Fields," in *Proceedings of the IEEE International Conference on Computer Vision*, 2021, pp. 5846–5854.
- [28] D. Ha, A. M. Dai, and Q. V. Le, "Hypernetworks," in *International Conference on Learning Representations*, 2017. [Online]. Available: <https://openreview.net/forum?id=rkpACe11x>
- [29] J. von Oswald, C. Henning, B. F. Grewe, and J. Sacramento, "Continual learning with hypernetworks," in *International Conference on Learning Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=SJgwNerKvB>
- [30] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, "On the Continuity of Rotation Representations in Neural Networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5745–5753.
- [31] X. Glorot and Y. Bengio, "Understanding the Difficulty of Training Deep Feedforward Neural Networks," in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics. JMLR Workshop and Conference Proceedings*, 2010, pp. 249–256.
- [32] M. Grupp, "evo: Python package for the Evaluation of Odometry and SLAM." <https://github.com/MichaelGrupp/evo>, 2017.