

Solving the Blind Perspective-n-Point Problem End-To-End With Robust Differentiable Geometric Optimization

Dylan Campbell*, Liu Liu*, and Stephen Gould

Australian National University, Australian Centre for Robotic Vision
{dylan.campbell, liu.liu, stephen.gould}@anu.edu.au

Abstract. Blind Perspective-n-Point (PnP) is the problem of estimating the position and orientation of a camera relative to a scene, given 2D image points and 3D scene points, without prior knowledge of the 2D–3D correspondences. Solving for pose and correspondences simultaneously is extremely challenging since the search space is very large. Fortunately it is a coupled problem: the pose can be found easily given the correspondences and vice versa. Existing approaches assume that noisy correspondences are provided, that a good pose prior is available, or that the problem size is small. We instead propose the first fully end-to-end trainable network for solving the blind PnP problem efficiently and globally, that is, without the need for pose priors. We make use of recent results in differentiating optimization problems to incorporate geometric model fitting into an end-to-end learning framework, including Sinkhorn, RANSAC and PnP algorithms. Our proposed approach significantly outperforms other methods on synthetic and real data.

Keywords: Camera pose estimation · PnP · Implicit differentiation

1 Introduction

The blind Perspective-n-Point (PnP) problem [35] aims to estimate the camera pose from which a set of 2D points were viewed, relative to an unordered 3D point-set. Specifically, the task is to find the rotation and translation that aligns a set of 2D bearing vectors with a set of 3D points, without knowledge of the true 2D–3D correspondences. The camera intrinsic parameters are assumed to be known, which allows 2D points to be expressed as bearing vectors. While a fundamental technique for many computer vision and robotic applications, including augmented reality and visual localization, it remains a challenging problem that has not as yet been satisfactorily solved.

The standard (non-blind) PnP problem [23, 32], where 2D–3D correspondences are known, is significantly less complex. It has a closed-form solution for three points [30] and, for a larger number of points, can be embedded in a

* Equal contribution




Fig. 1. Network architecture for our end-to-end blind PnP solver. We combine standard neural layers with declarative layers in a bi-level optimization framework to instantiate the traditional camera pose estimation pipeline (feature extraction, feature matching and optimization) in a single neural network. The input is a set of 2D and 3D point coordinates, from which point-wise features are extracted using standard network layers. Feature matching is then performed by computing the pairwise distance between the 2D and 3D point features, and using the Sinkhorn algorithm [42] to obtain a joint probability matrix. Finally, a probability-weighted blind PnP objective function is optimized from a RANSAC initialization to estimate the camera rotation and translation. The key contribution of this work is showing how this optimization procedure may be incorporated into an end-to-end learnable network by the use of declarative layers.

RANSAC framework [23] to reduce its sensitivity to outliers. However, it is inherently difficult to establish 2D–3D correspondences between modalities. As a result, PnP solvers are typically restricted to applications where both the 2D and 3D data contain visual information, such as structure-from-motion datasets [33]. Even for these, appearance may change seasonally, diurnally, and with weather, and so using geometric rather than visual features may improve generalizability.

Solving the PnP problem without correspondences is much more challenging, because the search space of correspondences and camera poses is very large, the objective function has many local optima, and outliers are prevalent. As a result, it was traditionally the domain of robust geometric algorithms that overcame the search space and non-convexity problems by requiring good pose priors [17, 35] or time-consuming global optimization [9, 11, 12]. Since these techniques were typically iterative, randomized and non-differentiable, this problem has not been amenable to a deep learning solution. Moreover, the geometry of the problem is difficult for a network to learn. However, there is significant opportunity in using a neural network for this problem, since it can effectively recognize patterns in the geometric data and thus reduce the search space and influence of outliers.

Fortunately, the framework of deep declarative networks [25] has recently been proposed, which provides a solution to the problem of including standard neural layers and geometric optimization layers inside the same end-to-end learnable network. This paper applies many of the ideas associated with deep declarative networks by formulating our deep blind PnP solver as a bi-level optimization problem. In this way, we aim to benefit from the pattern recognition capabilities of standard neural networks and the physical models and optimization algorithms used in traditional geometric approaches to the PnP problem.

We focus on the *optimization* part of the traditional camera pose estimation pipeline (feature extraction, matching and optimization) shown in Figure 1, that is, the remit of the PnP solver itself. To this end, we use an existing network architecture for feature extraction [47] and matching [34]. However, our key insight is that camera pose optimization algorithms, including robust global search techniques such as RANSAC [23] and state-of-the-art nonlinear PnP solvers, can be seamlessly integrated into an end-to-end deep learning framework. Our contributions are: 1) the first fully end-to-end trainable network for solving the blind PnP problem efficiently and globally; 2) the novel deployment of geometric model fitting algorithms as declarative layers inside the network; 3) the novel embedding of non-differentiable robust estimation techniques into the network; and 4) state-of-the-art performance on synthetic and real datasets.

2 Related Work

The majority of camera pose estimation methods assume that a set of 2D–3D correspondences is available and thus a PnP solver [23, 32] can be used. Hence, much of the effort in improving the visual localization pipeline focuses on robustly establishing correspondences [38, 40] or removing outliers [19, 44, 47]. These approaches are not appropriate for situations where correspondences cannot be easily obtained, such as when the 3D point-set has no associated visual information. In contrast, we address this problem by deferring the correspondence estimation task until the PnP stage of the pipeline, jointly estimating pose and correspondences. An alternative approach, which does not require explicit correspondences, is learning-based pose regression [28, 27, 45]. However, Sattler et al. [39] show that this essentially solves an image retrieval task rather than reasoning about 3D structure. Also, the camera is not localized with respect to an explicit 3D map, instead representing the scene implicitly. DSAC and extensions [7, 8] can localize with respect to a 3D model, but require many training images from the test scene. In contrast, we eschew visual information to learn generalizable geometric features, and never see the test scene during training.

To solve the blind PnP problem [35] of localizing the camera relative to an explicit unseen 3D model when correspondences cannot be obtained, some approaches seek a local optimum and assume that a good pose prior is available [17, 4]. For example, David et al. [17] propose an algorithm that alternates between solving for pose and correspondences, using the Sinkhorn algorithm [42]. To mitigate the pose prior requirement, global search strategies have been proposed, including multiple random starts [17] and probabilistic pose priors [35]. RANSAC [23] can also be used, but becomes intractable for moderately-sized problems. To obviate the need for pose priors and guarantee that a global optimum is found, globally-optimal approaches [11, 9, 12] use branch-and-bound to systematically reduce the search space. For example, Campbell et al. [12] globally optimize a robust distance between mixture distributions to solve the blind PnP problem. However, these optimal methods are time-consuming and limited to a moderate number of points, unlike our (orders of magnitude) faster approach.

Deep PnP solvers are proposed in existing work [16] (standard PnP) and concurrent work [34] (blind PnP). Due to difficulties inherent in eigendecomposition and outlier filtering, neither approach is end-to-end, despite being highly effective at learning 2D–3D correspondences. The former shows how to avoid unstable eigendecomposition gradients by applying a loss before the pose parameters are estimated, while the latter shows that high-quality 2D–3D correspondence matrices can be learned using optimal transport via the Sinkhorn algorithm [42]. Metric learning can also be used to learn matchable features, as shown for 2D–2D and 3D–3D matching [20]. However, estimating pose from these features or correspondence matrices requires a non-differentiable selection step, such as nearest neighbor search, to reduce the set of correspondences to a tractable size. Different to these approaches, we propose a fully end-to-end trainable blind PnP solver. We directly use an existing ResNet-based [26] feature extraction architecture [47] and a Sinkhorn-based [42] feature matching technique [34] in our network, since our focus is on the joint optimization of all parameters in the camera pose estimation pipeline. Our contribution is orthogonal to these works.

The declarative framework that allows us to incorporate geometric optimization algorithms into a deep network is described in Gould et al. [25]. They present theoretical results and analyses on how to differentiate constrained optimization problems via implicit differentiation. Differentiable convex problems have also been studied recently, including quadratic programs [3] and cone programs [2, 1]. In computer vision, the technique has been applied to video classification [21, 22], action recognition [14], visual attribute ranking [37], few-shot learning for visual recognition [31], and non-blind PnP in concurrent work [13]. In this work, we show that we can embed geometric model fitting algorithms and non-differentiable robust estimation techniques into a network as declarative layers to solve the blind PnP problem end-to-end.

3 An End-To-End Blind PnP Solver

In this section we present our end-to-end trainable network for solving the blind PnP problem, which we name BPnPNet. We start by formally defining the problem, then provide background on deep declarative networks and show how critical components of a blind PnP solver can be implemented as declarative layers. We then describe our network architecture, loss functions, and learning strategy.

3.1 Problem Formulation

Let $\mathbf{p} \in \mathbb{R}^3$ denote a 3D point and $\mathbf{f} \in \mathbb{R}^3$ denote a unit bearing vector corresponding to a 2D point in the image plane of a calibrated camera. That is, $\|\mathbf{f}\| = 1$ and $\mathbf{f} \propto \mathbf{K}^{-1}[u, v, 1]^T$, where \mathbf{K} is the matrix of intrinsic camera parameters and (u, v) are the 2D image coordinates. Given a set of bearing vectors $\mathcal{F} = \{\mathbf{f}_i\}_{i=1}^m$ and 3D points $\mathcal{P} = \{\mathbf{p}_i\}_{i=1}^n$, the objective of blind PnP is to find the rotation $\mathbf{R} \in SO(3)$ and translation $\mathbf{t} \in \mathbb{R}^3$ that transforms \mathcal{P} to the coordinate system of \mathcal{F} with the greatest number of one-to-one inlier correspondences, defined by an angular inlier threshold $\theta \in (0, \pi)$. The optimization problem is

$$\begin{aligned}
& \underset{\mathbf{R}, \mathbf{t}, \mathbf{C}}{\text{maximize}} && \sum_{i=1}^m \sum_{j=1}^n \mathbf{C}_{ij} (2 \llbracket \angle(\mathbf{f}_i, \mathbf{R}\mathbf{p}_j + \mathbf{t}) \leq \theta \rrbracket - 1) \\
& \text{subject to} && \mathbf{R} \in SO(3), \mathbf{t} \in \mathbb{R}^3 \\
& && \mathbf{C} \in \mathbb{B}^{m \times n}, \mathbf{C}\mathbf{1}^n \in \mathbb{B}^m, \mathbf{C}^\top \mathbf{1}^m \in \mathbb{B}^n
\end{aligned} \tag{1}$$

where \mathbf{C} is a Boolean one-to-one correspondence matrix with at most one non-zero element in each row and column, \mathbf{C}_{ij} is the element at row i and column j , $\mathbb{B} = \{0, 1\}$ is the Boolean domain, $\llbracket \cdot \rrbracket$ is an Iverson bracket, and $\angle(\mathbf{x}, \mathbf{y}) = \arccos(\|\mathbf{x}\|^{-1} \|\mathbf{y}\|^{-1} \mathbf{x}^\top \mathbf{y})$ is a function that returns the angle in $[0, \pi]$ between the vector arguments. This inlier maximization formulation optimizes a robust angular reprojection error. The joint optimization problem can be simplified if either the correspondences or the camera pose is known. If the correspondence matrix \mathbf{C} is known, we have the standard PnP problem. The camera pose can be estimated by minimizing the angular reprojection error

$$\frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n \mathbf{C}_{ij} \angle(\mathbf{f}_i, \mathbf{R}\mathbf{p}_j + \mathbf{t}). \tag{2}$$

If rotation \mathbf{R} and translation \mathbf{t} are known, then \mathbf{C} can be computed as

$$\tilde{\mathbf{C}}_{ij} = \llbracket \angle(\mathbf{f}_i, \mathbf{R}\mathbf{p}_j + \mathbf{t}) \leq \theta \rrbracket \tag{3}$$

followed by the Hungarian algorithm to enforce the one-to-one constraint. If a good pose or correspondence matrix initialization is available, these steps can be alternated to find a good estimate of \mathbf{R} , \mathbf{t} and \mathbf{C} . This strategy, analogous to the Iterative Closest Point algorithm [6], is taken by the SoftPOSIT algorithm [17]. In contrast, our approach applies this alternation implicitly during training.

3.2 Bi-Level Optimization

The conceptual framework that underpins this method is the deep declarative network [25], which interprets training a network as a bi-level optimization problem. According to this view, a network can be composed of multiple imperative and declarative layers. An imperative layer explicitly defines a function for transforming the input to the output, e.g., a convolution layer. In contrast, a declarative layer is implicitly defined in terms of the desired output, formulated as a constrained mathematical optimization problem. Declarative layers are more flexible and general than standard layers, since they admit constraints on the output and decouple the gradient computation from the algorithm used to solve the optimization problem. Crucially, the technique of implicit differentiation enables the back-propagation of gradients through a declarative layer without having to traverse the forward processing function, as shown in Figure 2.

A bi-level optimization problem [43, 5] for end-to-end learning has an upper-level problem solved subject to constraints imposed by a lower-level problem:

$$\begin{aligned}
& \text{minimize} && L(\mathbf{x}, \mathbf{y}) \\
& \text{subject to} && \mathbf{y} \in \arg \min_{\mathbf{u} \in \mathcal{C}} f(\mathbf{x}, \mathbf{u})
\end{aligned} \tag{4}$$




Fig. 2. Comparison of imperative and declarative layers. An imperative layer (left) transforms the input \mathbf{x} to the output \mathbf{y} using explicit forward functions, parameterized by network weights \mathbf{w} . A declarative layer (right) computes the output \mathbf{y} as a minimizer of an objective function, parameterized by the input \mathbf{x} and network weights \mathbf{w} . During learning, the gradient $DL(\mathbf{y})$ of the global loss function with respect to the output is propagated backwards using the chain rule. While the backward function of an imperative layer is tightly coupled with every step of the forward function, effectively unrolling any algorithm applied, the backward function of a declarative layer computes the gradient of the entire layer in one step. The individual forward processing nodes can be recursive or non-differentiable, provided that the objective function optimized by the final forward node is (sub)differentiable in \mathbf{x} .

where L is a global loss function, f is an objective function, \mathcal{C} is an arbitrary constraint set, and the loss is minimized over all network weights. To solve the bi-level optimization problem (4) by gradient descent, we require the derivative

$$DL(\mathbf{x}, \mathbf{y}) = D_X L(\mathbf{x}, \mathbf{y}) + D_Y L(\mathbf{x}, \mathbf{y}) D\mathbf{y}(\mathbf{x}) \quad (5)$$

in order to back-propagate gradients. The key challenge is to compute $D\mathbf{y}(\mathbf{x})$, for which we use implicit differentiation. We use the notation Df for the derivative of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, an $m \times n$ matrix with entries $(Df(\mathbf{x}))_{ij} = \partial f_i / \partial \mathbf{x}_j$, and we denote the partial derivative over the formal variable X , with all other variables fixed, as $D_X f(\mathbf{x}, \mathbf{y})$. We also use $D_{XY}^2 f$ as shorthand for $D_X (D_Y f)^\top$. For completeness, we collect the two results from Gould et al. [25] that are used in this paper. The unconstrained case applies Dini’s implicit function theorem [18, p19] to the first-order optimality condition $D_Y f(\mathbf{x}, \mathbf{y}) = \mathbf{0}$.

Lemma 1. Consider a function $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ and let $\mathbf{y}(\mathbf{x}) \in \arg \min_{\mathbf{u}} f(\mathbf{x}, \mathbf{u})$. Assume $\mathbf{y}(\mathbf{x})$ exists and that f is second-order differentiable in the neighborhood of $\mathbf{u} = \mathbf{y}(\mathbf{x})$. Set $\mathbf{H} = D_{YY}^2 f(\mathbf{x}, \mathbf{y}(\mathbf{x})) \in \mathbb{R}^{m \times m}$ and $\mathbf{B} = D_{XY}^2 f(\mathbf{x}, \mathbf{y}(\mathbf{x})) \in \mathbb{R}^{m \times n}$. Then for non-singular \mathbf{H} the derivative of \mathbf{y} with respect to \mathbf{x} is

$$D\mathbf{y}(\mathbf{x}) = -\mathbf{H}^{-1} \mathbf{B}. \quad (6)$$

We also require the linear equality constraints case.

Lemma 2. Consider a function $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ and let $\mathbf{A} \in \mathbb{R}^{p \times m}$ and $\mathbf{d} \in \mathbb{R}^p$ with $\text{rank}(\mathbf{A}) = p$ define a set of p under-constrained linear equations $\mathbf{A}\mathbf{u} = \mathbf{d}$. Also let $\mathbf{y}(\mathbf{x}) \in \arg \min_{\mathbf{u}} f(\mathbf{x}, \mathbf{u})$ subject to $\mathbf{A}\mathbf{u} = \mathbf{d}$. Assume that $\mathbf{y}(\mathbf{x})$ exists and that $f(\mathbf{x}, \mathbf{u})$ is second-order differentiable in the neighborhood of $\mathbf{u} = \mathbf{y}(\mathbf{x})$. Set $\mathbf{H} = D_{YY}^2 f(\mathbf{x}, \mathbf{y}(\mathbf{x}))$ and $\mathbf{B} = D_{XY}^2 f(\mathbf{x}, \mathbf{y}(\mathbf{x}))$. Then

$$D\mathbf{y}(\mathbf{x}) = (\mathbf{H}^{-1} \mathbf{A}^\top (\mathbf{A} \mathbf{H}^{-1} \mathbf{A}^\top)^{-1} \mathbf{A} \mathbf{H}^{-1} - \mathbf{H}^{-1}) \mathbf{B}. \quad (7)$$

3.3 Declarative Layers for Blind PnP

In this section, we will demonstrate how the theory of bi-level optimization may be applied to the blind PnP problem.

Weighted Blind PnP Layer: This declarative layer operates on the $m \times n$ product set of 2D–3D correspondences, optimizing the lower-level objective function

$$f(\mathbf{P}, \mathbf{r}, \mathbf{t}) = \sum_{i=1}^m \sum_{j=1}^n \mathbf{P}_{ij} \left(1 - \mathbf{f}_i^\top \frac{\mathbf{R}_r \mathbf{p}_j + \mathbf{t}}{\|\mathbf{R}_r \mathbf{p}_j + \mathbf{t}\|} \right) \quad (8)$$

over $\mathbf{r} \in \mathbb{R}^3$ and $\mathbf{t} \in \mathbb{R}^3$, where \mathbf{P} is a fixed joint correspondence probability matrix with $\sum \mathbf{P}_{ij} = 1$, a relaxation of the Boolean correspondence matrix \mathbf{C} , and \mathbf{r} is the angle-axis representation of the rotation \mathbf{R}_r such that $\mathbf{R}_r = \exp[\mathbf{r}]_\times$ for the skew symmetric operator $[\cdot]_\times$. The Rodrigues’ rotation formula provides an efficient closed-form solution to this exponential map. We use the angle-axis representation to automatically satisfy the constraints on \mathbf{R} , that is, $\mathbf{R} \in SO(3)$. We minimize this nonlinear function using the native PyTorch L-BFGS optimizer [10] to find $(\mathbf{r}^*, \mathbf{t}^*)$ for a given joint probability matrix \mathbf{P} .

Given optimal $(\mathbf{r}^*, \mathbf{t}^*)$, corresponding to \mathbf{y} in Lemma 1, we can compute the derivatives $D\mathbf{r}^*(\mathbf{P})$ and $D\mathbf{t}^*(\mathbf{P})$ using (6). Observe that the gradient computation is agnostic to the choice of optimization algorithm; we do not back-propagate through the L-BFGS iterations that were used to determine $(\mathbf{r}^*, \mathbf{t}^*)$. Instead, since the objective function is twice-differentiable, we only require that a (locally) optimal solution be found in order to compute the gradient in one step. While an analytic solution for the gradient can be obtained, it is quite unwieldy. In lieu of this, we use automatic differentiation to compute the necessary Jacobian and Hessian matrices. To be clear, automatic differentiation is applied to the specification of the objective function, not the algorithmic steps used to optimize it, which is distinctly different from standard usage in deep learning.

Importantly, the $m \times n$ product set of correspondences is too large for existing differentiable PnP solvers, such as DLT [16] and DSAC [7]. For $m = n = 1000$, 99.9% of the mn possible correspondences are outliers. Even with the weights \mathbf{P} , outliers will dominate the solver. We achieve robustness with top- k RANSAC (see below) and nonlinear optimization. In contrast, the non-robust linear estimate of DLT is unusably poor, and has severe numerical issues [16]. DSAC also fails in this case, because the probability of selecting an inlier hypothesis at random from the product set is vanishingly small. Unlike our declarative layer, DSAC cannot differentially select hypotheses from the top- k subset.

RANSAC: While Lemma 1 guarantees a local descent direction given the local minimizer \mathbf{y} , it is unlikely to be useful for the learning problem (or indeed the inference problem) if \mathbf{y} is a bad estimate. Hence it is helpful for \mathbf{y} to be, on average, a good local optimum—preferably the global optimum. Since the blind PnP objective function is non-convex with many local minima, a standard technique is to apply robust randomized global search such as RANSAC [23]. The declarative framework gives us the opportunity to incorporate this non-differentiable algorithm into an end-to-end learning network. We select a subset

of $k = 1.5 \min\{m, n\}$ correspondences from the mn possibilities, choosing those with the highest joint probability in \mathbf{P} . We then run RANSAC with the P3P algorithm [24] to find the inlier set, followed by the EPnP algorithm [32] on all inliers to refine the estimate. This robust estimate of the camera pose parameters is used to initialize the nonlinear weighted PnP optimization algorithm.

Since the final processing node in the declarative layer optimizes a twice-differentiable objective function, the non-differentiability of any intermediate computation is irrelevant to the gradient calculation. Note that this procedure for robustly estimating the camera pose parameters has no analytic solution and involves a non-differentiable algorithm. It would not be possible to use standard techniques such as explicit or automatic differentiation to obtain the gradient.

Sinkhorn Layer: We also define a declarative layer for feature matching in order to estimate the joint correspondence probability of the 2D–3D point pairs from a cost matrix $\mathbf{M} \in \mathbb{R}_+^{m \times n}$. This is achieved by encapsulating the Sinkhorn algorithm [42] in a declarative layer, as has been previously demonstrated in the literature [37]. The layer optimizes the lower-level objective function [15]

$$f(\mathbf{M}, \mathbf{P}) = \sum_{i=1}^m \sum_{j=1}^n (\mathbf{M}_{ij} \mathbf{P}_{ij} + \mu \mathbf{P}_{ij} (\log \mathbf{P}_{ij} - 1)) \quad (9)$$

with respect to $\mathbf{P} \in U(\mathbf{r}, \mathbf{c})$, where the transport polytope

$$U(\mathbf{r}, \mathbf{c}) = \{\mathbf{P} \in \mathbb{R}_+^{m \times n} \mid \mathbf{P}\mathbf{1}^n = \mathbf{r}, \mathbf{P}^\top \mathbf{1}^m = \mathbf{c}\} \quad (10)$$

is defined for the prior probability vectors $\mathbf{r} \in \mathbb{R}_+^m$ and $\mathbf{c} \in \mathbb{R}_+^n$ with $\sum \mathbf{r} = 1$ and $\sum \mathbf{c} = 1$, which represent the probability that any given 2D or 3D point has a valid match. In this work, we use uniform priors $\mathbf{r} = \frac{1}{m} \mathbf{1}$ and $\mathbf{c} = \frac{1}{n} \mathbf{1}$.

We run the highly-efficient Sinkhorn algorithm which optimizes this objective function, an entropy-regularized Wasserstein distance, in $O(m^2)$ [15]. This is considerably more efficient than the Hungarian algorithm, which exactly optimizes the Wasserstein distance in $O(m^3)$, while also converging to the Wasserstein distance as $\mu \rightarrow 0$. Given optimal \mathbf{P}^* , corresponding to \mathbf{y} in Lemma 2, we can compute the derivative $D\mathbf{P}^*(\mathbf{M})$ using (7). Unlike the PnP layer, we compute the derivative analytically to ensure memory efficiency. To do so, we need to form the matrices \mathbf{A} , \mathbf{B} and \mathbf{H} and perform the necessary inversions. We defer the details to the supplementary material.

The benefits of enclosing this algorithm in a declarative layer include being able to run the algorithm to convergence, rather than fixing the number of iterations, and obviating the need for unrolling the algorithmic steps and maintaining the requisite computation graph, which saves a significant amount of memory. Our implementation is much more memory efficient than that of Santa Cruz et al. [37], reducing $O(m^2 n^2)$ memory requirements to $O(mn)$. This allows much larger problems than were considered previously, such as $m = n = 1000$ used in this work. This was achieved by exploiting the block structure of the matrices \mathbf{A} and \mathbf{H} rather than storing them in full, and by computing the vector–Jacobian product rather than the Jacobian itself.

3.4 Network Architecture

Our network architecture is shown in Figure 1. First, we extract discriminative features from the 2D and 3D point-sets, aiming to recognise patterns in the data that are useful for establishing correspondences. Next, we estimate the correspondence probability for every 2D–3D pair by computing the pairwise distance between features and solving an optimal transport problem. Last, we optimize a weighted blind PnP objective function to obtain the (locally) optimal camera pose, given the data and estimated correspondence probability matrix.

Feature Extraction: To extract discriminative features from the 2D and 3D point-sets, we directly use the point feature extraction model from Yi et al. [47]. This model is a 12-layer ResNet [26], where each layer consists of a perceptron with 128 neurons per point, context normalization, batch normalization and a ReLU nonlinearity, with weights shared between points. Context normalization is the mechanism for sharing information between points, by normalizing with respect to the mean $\boldsymbol{\mu}^l$ and standard deviation σ^l of the feature vectors \mathbf{z}_i^l of every point at the l^{th} layer, and is given by $\text{CN}(\mathbf{z}_i^l) = (\mathbf{z}_i^l - \boldsymbol{\mu}^l)/\sigma^l$.

Before passing the data to the feature extraction networks, we do some initial processing. We convert the homogeneous bearing 3-vectors into inhomogeneous 2-vectors by dividing through by the z coordinate, since this requires fewer network parameters. We also apply a learned 3×3 transformation matrix to the 3D points using the input transform from PointNet [36], to align the points to a canonical orientation. Finally, after obtaining the pointwise feature vectors, we apply L_2 normalization, which is helpful for the ensuing optimization procedure.

Hence the 2D feature extractor encodes a mapping Φ with parameters ϕ from the 2D bearing vector set \mathcal{F} to the feature vector set $\mathcal{Z}_{\mathcal{F}}$, given by $\mathcal{Z}_{\mathcal{F}} = \Phi_{\phi}(\mathcal{F})$ with $\mathcal{Z}_{\mathcal{F}} = \{\mathbf{z}_{\mathbf{f}_i}\}_{i=1}^m$ and $\mathbf{z}_{\mathbf{f}_i} \in \mathbb{R}^{128}$. The 3D feature extractor encodes a similar mapping, given by $\mathcal{Z}_{\mathcal{P}} = \Psi_{\psi}(\mathcal{P})$ with $\mathcal{Z}_{\mathcal{P}} = \{\mathbf{z}_{\mathbf{p}_i}\}_{i=1}^n$ and $\mathbf{z}_{\mathbf{p}_i} \in \mathbb{R}^{128}$.

Correspondence Probability: To estimate the probability that a 2D–3D point pair is an inlier correspondence, we compute the pairwise distances between the feature vector sets $\mathcal{Z}_{\mathcal{F}}$ and $\mathcal{Z}_{\mathcal{P}}$ and then solve an optimal transport problem, as was shown to be effective in concurrent work [34]. The elements of the pairwise L_2 distance matrix $\mathbf{M} \in \mathbb{R}_+^{m \times n}$ are computed as $\mathbf{M}_{ij} = \|\mathbf{z}_{\mathbf{f}_i} - \mathbf{z}_{\mathbf{p}_j}\|_2$. We then solve the regularized transport problem [15] using the Sinkhorn algorithm [42] to obtain a joint probability matrix \mathbf{P} . The advantage of this approach is that it considers the entirety of \mathbf{M} when estimating the probability \mathbf{P}_{ij} , in order to resolve correspondence ambiguities. Finding a jointly optimal solution is critical if the learning goal is to approach the *one-to-one* correspondence matrix \mathbf{C} up to scale. This optimization problem was outlined in Section 3.3, where we showed that the gradient computation can be decoupled from the Sinkhorn algorithm using implicit differentiation.

Blind PnP Optimization: Given the joint correspondence probability matrix \mathbf{P} , we can now optimize the weighted nonlinear blind PnP objective function (8) to obtain the optimal camera pose (\mathbf{R}, \mathbf{t}) for that set of correspondence probabilities. The optimization problem was outlined in Section 3.3, where we

showed how to find a locally-optimal camera pose using the L-BFGS algorithm, how to ensure it is a good local optimum (on average) using RANSAC, and how to back-propagate through the layer. Hence we have a fully-differentiable way to generate camera pose parameters that are likely to be near the global optimum of the non-convex objective function. At test time, we can either take the network output or the RANSAC output computed within the network; both are evaluated in the experiments.

3.5 Learning From Pose-Labelled Data

Loss Functions: We use two loss functions, one for each component of the coupled problem. The first is a correspondence loss L_c to bring the estimated correspondence matrix \mathbf{P} closer to the ground-truth. The second is a pose loss L_p to encourage the network to generate correspondence matrices that are amenable to our PnP solver. Note that these are distinct, albeit complementary, aims. While a perfect correspondence matrix would generate an accurate camera pose, this is not achievable in practice. Instead, there is a family of correspondence matrices for a fixed suboptimal value of L_c , which will differ considerably in their suitability for the weighted PnP solver.

The correspondence loss L_c arises directly from the problem formulation (1) given the ground-truth rotation \mathbf{R}_{gt} and translation \mathbf{t}_{gt} , yielding

$$L_c = \sum_i^m \sum_j^n \mathbf{P}_{ij} (1 - 2 \llbracket \angle(\mathbf{f}_i, \mathbf{R}_{\text{gt}} \mathbf{p}_j + \mathbf{t}_{\text{gt}}) \leq \theta \rrbracket \rrbracket) \quad (11)$$

where $\llbracket \cdot \rrbracket$ is an Iverson bracket (indicator function), and θ is the angular inlier threshold. The loss is bounded, since $\sum \mathbf{P}_{ij} = 1$ and so $L_c \in [-1, 1]$, and has the interpretation of maximizing the probability of the inlier correspondences and minimizing the probability of the outlier correspondences, since \mathbf{P} is a joint probability matrix. If the ground-truth correspondence matrix \mathbf{C} is available, this can be used instead of the indicator function. It is also possible to use the (less robust) reprojection error (2), however we found no advantage to this.

The camera pose loss L_p uses standard error measures on rotations and translations, and is given by

$$L_p = L_r + L_t \quad (12)$$

$$L_r = \angle(\mathbf{R}, \mathbf{R}_{\text{gt}}) = \arccos \frac{1}{2} (\text{trace } \mathbf{R}_{\text{gt}}^T \mathbf{R} - 1) \quad (13)$$

$$L_t = \|\mathbf{t} - \mathbf{t}_{\text{gt}}\|_2 \quad (14)$$

This loss is not bounded, since $L_r \in [0, \pi]$ and $L_t \in [0, \infty)$. The argument of \arccos is clamped to between $\pm(1 - \epsilon)$ for $\epsilon = 10^{-7}$ to prevent an infinite gradient at 0° and 180° . Finally, the total loss is given by

$$L = L_c + \gamma_p L_p \quad (15)$$

where γ_p is a hyperparameter that controls the relative influence of L_p .

Learning Strategy: We train the network implemented in PyTorch using the Adam optimizer [29] with a learning rate of 10^{-5} and otherwise default parameters. We use a batch size of 16 and train for 120 epochs (to convergence) with the correspondence loss only ($\gamma_p = 0$), followed by 20–80 epochs with the pose loss as well ($\gamma_p = 1$). This reflects the intuition that the pose loss is more meaningful once the correspondence probability matrix \mathbf{P} has useful information, having reduced the correspondence search space.

Implementation Details: For the Sinkhorn algorithm, the entropy parameter μ was set to 0.1; for RANSAC, the inlier reprojection error was set to 0.01 and the maximum number of iterations was set to 1000; and for the L-BFGS solver, the line search function was set to strong Wolfe, the maximum number of iterations varied with the number of points to standardize the batch runtime, and the gradient norms were clipped to 100. Ground-truth correspondences were used in training instead of specifying inlier threshold θ . All experiments were run on a single Titan V GPU, and the PyTorch code, including modular Sinkhorn and weighted PnP layers, will be released.

4 Results

Our blind PnP network, named BPnPNet, is evaluated with respect to the baseline algorithms SoftPOSIT [17], RANSAC [23], and GOSMA [12] on synthetic and real data. These are state-of-the-art representative examples of a local blind PnP solver (SoftPOSIT), a global solver (RANSAC), and a globally-optimal solver (GOSMA). For RANSAC, we randomly sample 2D–3D correspondences and use a minimal P3P solver [30].¹ For SoftPOSIT, we provide an initialization using ground-truth pose information, since it is a local solver and therefore requires a good pose prior. The algorithms were stopped early if their runtime for a single point-set pair exceeded 30s, returning the best pose found so far. This ensured that evaluation time was bounded at four days per algorithm on the datasets tested. Globally-optimal algorithms often exceed this limit, but it is infeasible to evaluate them to convergence on large datasets.

We use the synthetic ModelNet40 dataset [46] and the real-world MegaDepth dataset [33] for evaluation. The former is a CAD mesh model dataset, while the latter is a multi-view photo dataset with COLMAP [41] reconstructions providing the 2D and 3D point-sets. MegaDepth has highly diverse scenes, camera poses, and point distributions. We report quartiles for rotation error (in degrees), translation error, and reprojection error (in degrees), according to (13), (14), and (2) respectively. We denote the first, second (median) and third quartiles as Q1, Q2 and Q3. We also report average runtime for inference (in seconds) and recall at a particular error threshold (as a percentage), that is, the percentage of poses with an error less than that threshold.

¹ The probability of choosing a minimal set of 4 true 2D–3D correspondences from the size mn set of all correspondences without replacement is $\prod_{i=0}^3 \frac{m-i}{(m-i)(n-i)} \approx 10^{-12}$ for $m = n = 1000$ and no outliers. The number of RANSAC iterations required to achieve 90% confidence is thus $\log(1 - 0.9) / \log(1 - 10^{-12}) \approx 2.3 \times 10^{12}$.

Table 1. Results on the ModelNet40 [46] test set. We report quartiles for rotation error ($^\circ$), translation error and reprojection error ($^\circ$), and the mean runtime T (s). Note that Ours L_cR is a standard RANSAC baseline: deep 2D–3D feature matching followed by P3P-RANSAC. [†]Algorithms were run for a maximum of 30s.

Method	Rotation Error			Translation Error			Reproj. Error			T
	Q1	Q2	Q3	Q1	Q2	Q3	Q1	Q2	Q3	\bar{x}
SoftPOSIT [17]	16.1	21.8	28.0	0.33	0.49	0.72	2.82	3.98	5.21	27 [†]
RANSAC [23]	90.8	139	165	0.43	1.15	3.08	4.22	5.87	8.06	30 [†]
GOSMA [12]	10.1	22.1	52.0	0.25	0.46	0.75	1.04	1.62	3.11	30 [†]
Ours L_c	6.08	11.3	18.3	0.34	0.52	0.81	0.56	0.86	1.31	0.1
Ours L_cL_p	4.88	9.66	16.0	0.04	0.08	0.15	0.36	0.61	1.03	0.1
Ours L_cR	5.49	11.7	20.0	0.04	0.09	0.20	0.37	0.70	1.25	0.1
Ours L_cL_pR	3.33	8.09	15.8	0.04	0.08	0.16	0.28	0.52	1.01	0.1

4.1 Synthetic Data Experiments

In this section, we evaluate our network on the synthetic ModelNet40 dataset [46] and conduct ablation studies. To generate the synthetic data from the mesh models, we uniformly sampled 1000 3D points from each model and generated virtual cameras by drawing Euler rotation angles uniformly from $[0, \pi/4]$ and translations from $[-0.5, 0.5]$, with an offset of 4.5 along the z axis. The points were projected to a 640×480 virtual image with a focal length of 800 and normal noise with $\sigma = 2$ pixels was applied to the 2D points. In this way, we generated training and testing sets of 40000 and 2468 2D–3D point-set pairs respectively, each from the standard train and test splits of ModelNet40. SoftPOSIT was initialized using the mean ground-truth Euler angles and translation, corresponding to a median initial rotation error of 21.5° and translation error of 0.49.

The results on the ModelNet40 test set are shown in Table 1 and Figure 3. They demonstrate that our network obtains significantly better camera pose results than state-of-the-art local, global and globally-optimal algorithms. The results also include our ablation study, where we compare our model’s camera pose output (Ours L_cL_p) with a variant of our model that is learnt without the pose loss (Ours L_c). In all cases, the pose loss improves the results significantly, especially the translation errors. In particular, the recall for rotation errors less than 15° and translation errors less than 0.5 is 72%, an improvement of 25% over using the L_c loss only. We also compare our model’s output with the RANSAC pose computed within our network, denoted by R in the results and $(\mathbf{R}_0, \mathbf{t}_0)$ in Figure 1. The robust RANSAC estimate tends to be more accurate than the model’s final output since it is more resistant to errors in the correspondence probability matrix, and so should be used in any applications. Our method is also at least two orders of magnitude faster than the other methods, taking 0.12s on average. Note that Ours L_cR is an example of the standard RANSAC baseline of deep 2D–3D feature matching followed by P3P-RANSAC, without end-to-end training. Finally, a small reduction of only $2^\circ/0.01$ on the median statistics is observed with weaker pose-only supervision (see supplementary material).




Fig. 3. Recall on the ModelNet40 (top) and MegaDepth (bottom) test sets, with respect to an error threshold τ . R denotes using the RANSAC estimate.




Fig. 4. Qualitative results for the MegaDepth dataset: 3D point-sets projected onto the image using the camera pose found by GOSMA [12] (top) and our method (bottom).

4.2 Real Data Experiments

Here we evaluate our network on the MegaDepth dataset [33]. To generate the splits, we randomly selected landmarks and obtained train and test sets of 40828 and 10795 2D–3D point-set pairs respectively. The landmarks do not overlap across splits; we test how well the network *generalizes* to unseen locations. There are between 5–15000 points per set, reflecting the variability of real structure-from-motion data. SoftPOSIT was initialized using the ground-truth, perturbed by up to 10° about a random axis and up to 0.5m in a random direction.

Table 2. Results on the MegaDepth [33] test set. We report quartiles for rotation error ($^\circ$), translation error and reprojection error ($^\circ$), and the mean runtime T (s). [†]Algorithms were run for a maximum of 30s.

Method	Rotation Error			Translation Error			Reproj. Error			T
	Q1	Q2	Q3	Q1	Q2	Q3	Q1	Q2	Q3	\bar{x}
SoftPOSIT [17]	1.81	21.4	165	0.24	1.53	6.10	0.92	7.85	24.1	18 [†]
RANSAC [23]	66.6	122	155	6.80	15.2	28.2	4.45	8.77	13.3	30 [†]
GOSMA [12]	8.69	86.8	145	1.07	5.67	9.34	1.30	13.7	37.1	30 [†]
Ours L_c	1.91	4.47	11.4	0.52	1.05	2.34	0.54	1.12	2.81	0.2
Ours L_cL_p	1.32	3.31	8.84	0.21	0.46	1.08	0.21	0.53	1.64	0.2
Ours L_cR	0.44	1.55	7.70	0.05	0.18	0.80	0.06	0.16	1.27	0.2
Ours L_cL_pR	0.34	1.00	4.88	0.04	0.12	0.53	0.06	0.12	0.74	0.2

The results on the MegaDepth test set are shown in Table 2 and Figure 3, with qualitative results given in Figure 4. Our approach outperforms the other algorithms by a significant margin, notably doing better than a local optimization algorithm initialized very close to the ground-truth. GOSMA performs poorly on this dataset, despite being optimal, since it rarely converges within the 30s evaluation limit (often taking minutes). As with the synthetic data, the pose loss improves the results, especially the translation errors. In particular, the recall for rotation errors less than 10° and translation errors less than 1 is 73%, 25% better than without the loss. With the RANSAC estimate, the recall further improves to 82%. Additional results, including outlier analysis, failure cases, and another feature matching approach, are provided in the supplementary material.

5 Conclusion

In this paper, we have proposed the first fully end-to-end trainable network for solving the blind PnP problem. The key insight is that we can back-propagate through a geometric optimization algorithm using the technique of implicit differentiation. This allows us to compute a gradient even when the declarative layer involves non-differentiable RANSAC search and L-BFGS optimization of a nonlinear geometric objective. For such a layer, unrolling the algorithmic steps and computing the gradient with automatic differentiation is not possible, and would not be advisable even if it were due to the memory and computational requirements. Furthermore, we show that our method outperforms state-of-the-art geometric blind PnP solvers by a considerable margin when pose-labelled training data is available. Promisingly, our declarative approach admits the possibility of an unsupervised reprojection error loss, which may be used to fine-tune our pre-trained model to test scene data without the need for ground-truth labels.

Acknowledgements

This work was conducted by the Australian Research Council Centre of Excellence for Robotic Vision (CE140100016), funded by the Australian Government.

References

1. Agrawal, A., Amos, B., Barratt, S., Boyd, S., Diamond, S., Kolter, Z.: Differentiable convex optimization layers. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 32 (NIPS 2019)*. pp. 9562–9574. Curran Associates, Inc. (2019)
2. Agrawal, A., Barratt, S., Boyd, S., Busseti, E., Moursi, W.: Differentiating through a cone program. *Journal of Applied and Numerical Optimization* **1**(2), 107–115 (2019)
3. Amos, B., Kolter, J.Z.: OptNet: Differentiable optimization as a layer in neural networks. In: Precup, D., Teh, Y.W. (eds.) *Proceedings of the 34th International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 70, pp. 136–145. PMLR, International Convention Centre, Sydney, Australia (06–11 Aug 2017)
4. Baka, N., Metz, C., Schultz, C.J., van Geuns, R.J., Niessen, W.J., van Walsum, T.: Oriented Gaussian mixture models for nonrigid 2D/3D coronary artery registration. *IEEE Transactions on Medical Imaging* **33**(5), 1023–1034 (May 2014). <https://doi.org/10.1109/TMI.2014.2300117>
5. Bard, J.F.: *Practical Bilevel Optimization: Algorithms and Applications*. Kluwer Academic Press (1998)
6. Besl, P.J., McKay, N.D.: A method for registration of 3-D shapes. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)* **14**(2), 239–256 (1992)
7. Brachmann, E., Krull, A., Nowozin, S., Shotton, J., Michel, F., Gumhold, S., Rother, C.: DSAC – Differentiable RANSAC for camera localization. In: *Proceedings of the 2017 Conference on Computer Vision and Pattern Recognition*. pp. 2492–2500. IEEE Computer Society (Jul 2017). <https://doi.org/10.1109/CVPR.2017.267>
8. Brachmann, E., Rother, C.: Learning less is more - 6D camera localization via 3D surface regression. In: *Proceedings of the 2018 Conference on Computer Vision and Pattern Recognition*. pp. 4654–4662. IEEE Computer Society (2018)
9. Brown, M., Windridge, D., Guillemaut, J.Y.: Globally optimal 2D-3D registration from points or lines without correspondences. In: *Proceedings of the 2015 International Conference on Computer Vision*. pp. 2111–2119 (Dec 2015)
10. Byrd, R.H., Lu, P., Nocedal, J., Zhu, C.: A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing* **16**(5), 1190–1208 (1995)
11. Campbell, D., Petersson, L., Kneip, L., Li, H.: Globally-optimal inlier set maximisation for camera pose and correspondence estimation. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)* **42**(2), 328–342 (Feb 2020). <https://doi.org/10.1109/TPAMI.2018.2848650>
12. Campbell, D., Petersson, L., Kneip, L., Li, H., Gould, S.: The alignment of the spheres: Globally-optimal spherical mixture alignment for camera pose estimation. In: *Proceedings of the 2019 Conference on Computer Vision and Pattern Recognition*. pp. 11796–11806. IEEE Computer Society (2019)
13. Chen, B., Parra, A., Cao, J., Li, N., Chin, T.J.: End-to-end learnable geometric vision by backpropagating PnP optimization. In: *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society (Jun 2020)
14. Cherian, A., Fernando, B., Harandi, M., Gould, S.: Generalized rank pooling for action recognition. In: *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society (2017)

15. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. In: Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 2292–2300. Curran Associates Inc. (2013)
16. Dang, Z., Moo Yi, K., Hu, Y., Wang, F., Fua, P., Salzmann, M.: Eigendecomposition-free training of deep networks with zero eigenvalue-based losses. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *Proc. of the European Conference on Computer Vision (ECCV)*. pp. 768–783. Springer (2018)
17. David, P., Dementhon, D., Duraiswami, R., Samet, H.: SoftPOSIT: simultaneous pose and correspondence determination. *International Journal of Computer Vision (IJCV)* **59**(3), 259–284 (2004)
18. Dontchev, A.L., Rockafellar, R.T.: *Implicit Functions and Solution Mappings: A View from Variational Analysis*. Springer-Verlag, 2nd edn. (2014)
19. Enqvist, O., Kahl, F.: Robust optimal pose estimation. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) *Proceedings of the 2008 European Conference on Computer Vision*. pp. 141–153. Springer (Oct 2008)
20. Fathy, M.E., Tran, Q.H., Zeeshan Zia, M., Vernaza, P., Chandraker, M.: Hierarchical metric learning and matching for 2D and 3D geometric correspondences. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *Proc. of the European Conference on Computer Vision (ECCV)*. pp. 803–819. Springer (2018)
21. Fernando, B., Gould, S.: Learning end-to-end video classification with rank-pooling. In: Balcan, M.F., Weinberger, K.Q. (eds.) *Proc. of the International Conference on Machine Learning (ICML)*. PMLR (2016)
22. Fernando, B., Gould, S.: Discriminatively learned hierarchical rank pooling networks. In: *International Journal of Computer Vision (IJCV)*. vol. 124, pp. 335–355 (Sep 2017)
23. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* **24**(6), 381–395 (1981)
24. Gao, X.S., Hou, X.R., Tang, J., Cheng, H.F.: Complete solution classification for the perspective-three-point problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **25**(8), 930–943 (Aug 2003)
25. Gould, S., Hartley, R., Campbell, D.: Deep declarative networks: A new hope. Tech. rep., Australian National University (arXiv:1909.04866) (Sep 2019)
26. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society (2016)
27. Kendall, A., Cipolla, R.: Geometric loss functions for camera pose regression with deep learning. In: *Proceedings of the 2017 Conference on Computer Vision and Pattern Recognition*. pp. 6555–6564. IEEE Computer Society (Jul 2017). <https://doi.org/10.1109/CVPR.2017.694>
28. Kendall, A., Grimes, M., Cipolla, R.: PoseNet: A convolutional network for real-time 6-DOF camera relocalization. In: *Proceedings of the 2015 International Conference on Computer Vision*. pp. 2938–2946. IEEE Computer Society (Dec 2015). <https://doi.org/10.1109/ICCV.2015.336>
29. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Bengio, Y., LeCun, Y. (eds.) *Proc. of the International Conference on Learning Representations (ICLR)* (2015)
30. Kneip, L., Scaramuzza, D., Siegwart, R.: A novel parametrization of the perspective-three-point problem for a direct computation of absolute camera po-

- sition and orientation. In: Proceedings of the 2011 Conference on Computer Vision and Pattern Recognition. pp. 2969–2976. IEEE, IEEE Computer Society (Jun 2011)
31. Lee, K., Maji, S., Ravichandran, A., Soatto, S.: Meta-learning with differentiable convex optimization. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE Computer Society (2019)
 32. Lepetit, V., Moreno-Noguer, F., Fua, P.: EPnP: An accurate $O(n)$ solution to the PnP problem. *International Journal of Computer Vision* **81**(2), 155–166 (2009)
 33. Li, Z., Snavely, N.: MegaDepth: Learning single-view depth prediction from internet photos. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2041–2050. IEEE Computer Society (2018)
 34. Liu, L., Campbell, D., Li, H., Zhou, D., Song, X., Yang, R.: Learning 2D–3D correspondences to solve the blind perspective-n-point problem. Tech. rep., Australian National University (arXiv:2003.06752) (Mar 2019)
 35. Moreno-Noguer, F., Lepetit, V., Fua, P.: Pose priors for simultaneously solving alignment and correspondence. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) Proceedings of the 2008 European Conference on Computer Vision. pp. 405–418. Springer, Springer (Oct 2008)
 36. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: PointNet: deep learning on point sets for 3D classification and segmentation. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 652–660. IEEE, IEEE Computer Society, Honolulu, USA (Jul 2017)
 37. Santa Cruz, R., Fernando, B., Cherian, A., Gould, S.: Visual permutation learning. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)* **41**(12), 3100–3114 (2019)
 38. Sattler, T., Leibe, B., Kobbelt, L.: Efficient effective prioritized matching for large-scale image-based localization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **39**(9), 1744–1756 (Sep 2017). <https://doi.org/10.1109/TPAMI.2016.2611662>
 39. Sattler, T., Zhou, Q., Pollefeys, M., Leal-Taixe, L.: Understanding the limitations of CNN-based absolute camera pose regression. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE Computer Society (Jun 2019)
 40. Schönberger, J.L., Pollefeys, M., Geiger, A., Sattler, T.: Semantic visual localization. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6896–6906. IEEE Computer Society (2018)
 41. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, IEEE Computer Society (2016)
 42. Sinkhorn, R.: Diagonal equivalence to matrices with prescribed row and column sums. *The American Mathematical Monthly* **74**(4), 402–405 (1967)
 43. von Stackelberg, H., Bazin, D., Urch, L., Hill, R.R.: Market structure and equilibrium. Springer (2011)
 44. Svärm, L., Enqvist, O., Kahl, F., Oskarsson, M.: City-scale localization for cameras with known vertical direction. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **39**(7), 1455–1461 (2016)
 45. Walch, F., Hazirbas, C., Leal-Taixe, L., Sattler, T., Hilsenbeck, S., Cremers, D.: Image-based localization using lstms for structured feature correlation. In: Proc. of the International Conference on Computer Vision (ICCV). pp. 627–637. IEEE Computer Society (2017)

46. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3D ShapeNets: A deep representation for volumetric shapes. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1912–1920. IEEE Computer Society (2015)
47. Yi, K.M., Trulls, E., Ono, Y., Lepetit, V., Salzmann, M., Fua, P.: Learning to find good correspondences. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2666–2674. IEEE Computer Society (2018)