Neighbourhood Consensus Networks

Ignacio Rocco[1,2] **Mircea Cimpoi**[3] **Relja Arandjelović**[4] **Akihiko Torii**[5] **Tomas Pajdla**[3] **Josef Sivic**[1,2,3]

[1]DI ENS, PSL Research University [2]INRIA [3]CIIRC, CTU in Prague [4]DeepMind [5]Tokyo Institute of Technology

Introduction

Task: find pixel-level visual correspondences

Challenges: strong illumination or appearance changes

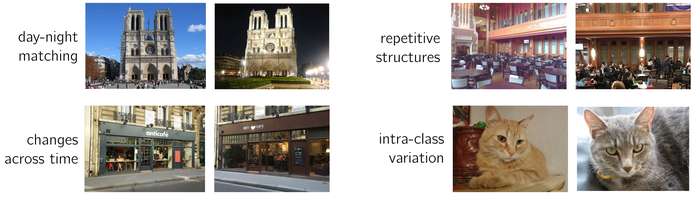

Contributions:

- Trainable method for feature extraction, matching and filtering
- Based on dense CNN descriptors and a novel 4D neighbourhood consensus CNN
- Trainable without correspondence annotation

Review - classical pipeline:

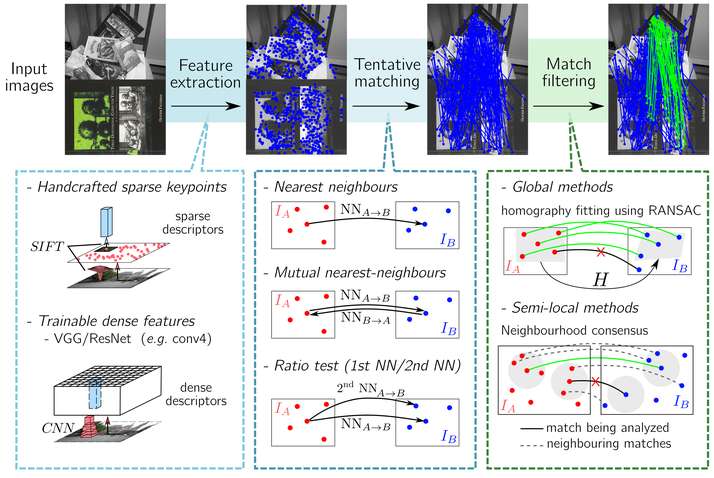

Proposed method

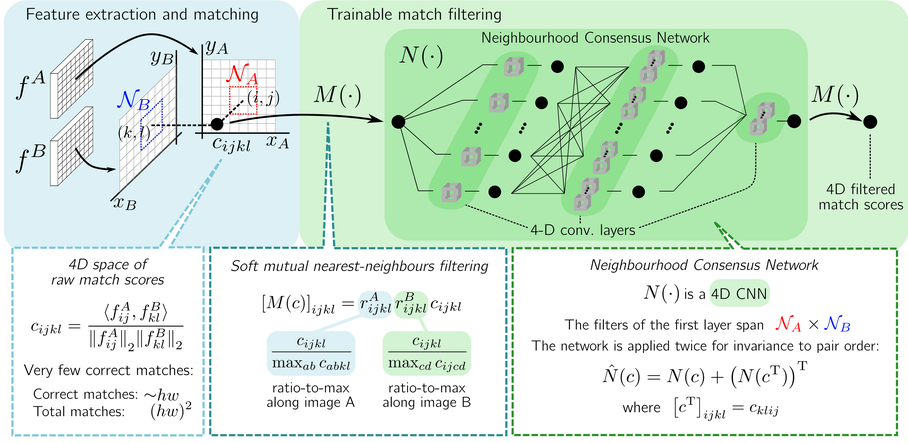

Extracting correspondences:

Matches can be extracted in both directions from the output c:

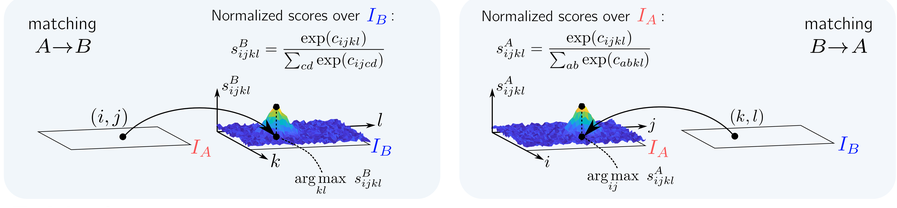

Training loss:

The network is trained with weak supervision:

$$\mathcal{L}(I^A, I^B) = -y\left(\bar{s}^A + \bar{s}^B\right)$$

mean score of matches $A \to B$

mean score of matches $B \to A$

- positive pairs ($y = 1$): **maximize** match score
- negative pairs ($y = -1$): **minimize** match score

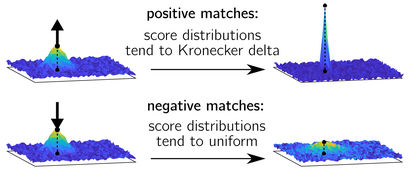

Experimental results

Category-level matching: PF-Pascal dataset [1]

- **Task:** match similar semantic parts
- **Metric:** percentage of correct keypoints (PCK)

Method	PCK
HOG+PF-LOM [1]	62.5
SCNet-AG+ [2]	72.2
CNNGeo [3]	71.9
WeakAlign [4]	75.8
NC-Net	**78.9**

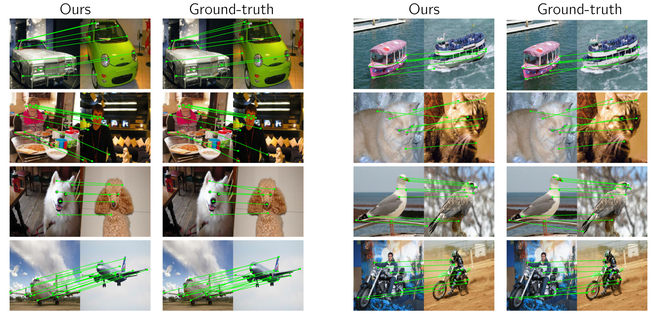

Instance-level matching: InLoc dataset [5]

- **Task:** retrieve 6-dof camera poses of query images
- **Metric:** percentage of correctly localized queries

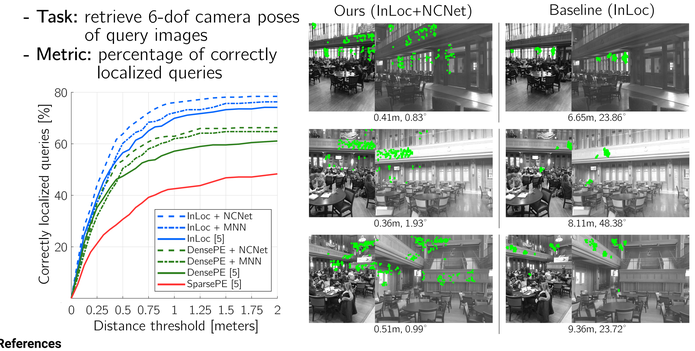

References

[1] B. Ham, M. Cho, C. Schmid, and J. Ponce. Proposal flow: Semantic correspondences from object proposals. *IEEE PAMI*, 2017.
[2] K. Han, R. S. Rezende, B. Ham, K.-Y. K. Wong, M. Cho, C. Schmid, and J. Ponce. SCNet: Learning Semantic Correspondence. In *Proc. ICCV*, 2017.
[3] I. Rocco, R. Arandjelović, and J. Sivic. Convolutional neural network architecture for geometric matching. In *Proc. CVPR*, 2017.
[4] I. Rocco, R. Arandjelović, and J. Sivic. End-to-end weakly-supervised semantic alignment. In *Proc. CVPR*, 2018.
[5] H. Taira, M. Okutomi, T. Sattler, M. Cimpoi, M. Pollefeys, J. Sivic, T. Pajdla, and A. Torii. InLoc: Indoor visual localization with dense matching and view synthesis. In *Proc. CVPR*, 2018.