Learning Cross-domain Information Transfer for Location Recognition and Clustering - Supplementary Material

Raghuraman Gopalan
Video and Multimedia Technologies Research Dept., AT&T Labs-Research, Middletown, NJ 07748 USA
raghuram@research.att.com

1. Karcher Mean Computation

1. Given a set of n points $\{f(x_i)\}$ on the manifold.
2. Let μ_0 be an initial estimate of the Karcher mean, usually obtained by picking one element of $\{f(x_i)\}$ at random. Set j = 0.
3. For each $i = 1, .., k$, compute the inverse exponential map (Algorithm 1 in main paper) ν_i of $f(x_i)$ about the current estimate of the mean, i.e. $\nu_i = \exp_{\mu_j}^{-1}(f(x_i))$.
4. Compute the average tangent vector $\bar{\nu} = \frac{1}{k}\sum_{i=1}^{k} \nu_i$.
5. If $\|\bar{\nu}\|$ is small, then stop. Else, move μ_j in the average tangent direction using $\mu_{j+1} = \exp_{\mu_j}(\epsilon\bar{\nu})$, where $\epsilon > 0$ is small step size, typically 0.5, and $\exp_{\mu_j}$ is the exponential map (Algorithm 2 in main paper) at μ_j.
6. Set $j = j + 1$ and return to Step 3. Continue till μ_j does not change anymore or till maximum iterations are exceeded.

Algorithm 1: Algorithm to compute the sample Karcher mean [2].

2. Experimental Results

Here we discuss experimental results that were not included in the main paper due to space constraints.

2.1. im2GPS Dataset

In Figure 1(a) and 1(b) we show the results on the 2K random and geographically uniform test sets of the im2GPS dataset [3] for all three values of the number of domains in the classification stage, $c' = 64, 128, 256$. As is the case with the default test set (Figure 5b in the main paper), we observe improvement in performance with finer concentration of the domains.

In Figure 2(a) and 2(b) we present results on Cases A1,A2,B1,B2 on the default test while varying c'. We see that while the approach proposed in Section 2 of the main paper is better than all these cases, Case A1 is better than Case A2 for all values of c', and $c' = 256$ gives the best performance over the other two values of c'. Similar trends are observed with Case B1 better than Case B2, and interestingly here, Case B1 with all three values of c' gives better performance than Case B2. This further emphasizes the need to perform PCA of the matrix M_i' to obtain M_i.

2.2. San Francisco Dataset

We then report results on the San Francisco dataset [1] PCI images for varying values of c', with and without using GPS, in Figure 3. We see that performance improves with finer domains, which is also the case with the PFI images (Figure 4). In Figure 5 we study cases A1,A2,B1 and B2 on the PCI images with and without GPS. As was the case with the im2GPS dataset, Case A1 performs better than Case A2, and Case B1 bests B2, while the approach presented in Section 2 of the main paper is better than all these cases.

References

[1] D. Chen, G. Baatz, K. Koser, S. Tsai, R. Vedantham, T. Pylvanainen, K. Roimela, X. Chen, J. Bach, M. Pollefeys, B. Girod, and R. Grzeszczuk. City-scale landmark identification on mobile devices. In *CVPR*, pages 737–744, 2011. 1

[2] Y. Chikuse. *Statistics on special manifolds*, volume 174 of *Lecture Notes in Statistics*. Springer Verlag, 2003. 1

[3] J. Hays and A. Efros. Im2gps: estimating geographic information from a single image. In *CVPR*, pages 1–8, 2008. 1

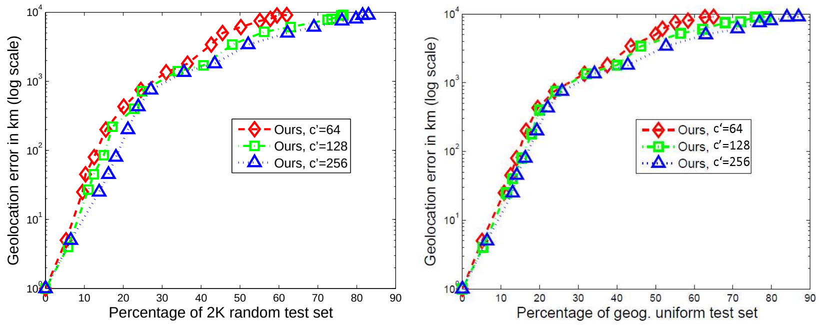

(a) (b)

Figure 1. Performance on the 2K random and geographically uniform test sets of the im2GPS dataset, across different c' values. Performance improves with the number of domains used in the classification stage.

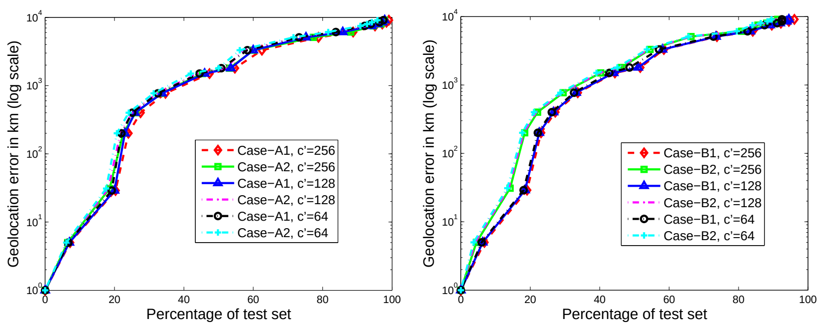

(a) (b)

Figure 2. Evaluating cases A1,A2,B1,B2 on the im2GPS default test set. Case A1 is better than A2, with B1 better than B2, and the performance improves with larger c' values.

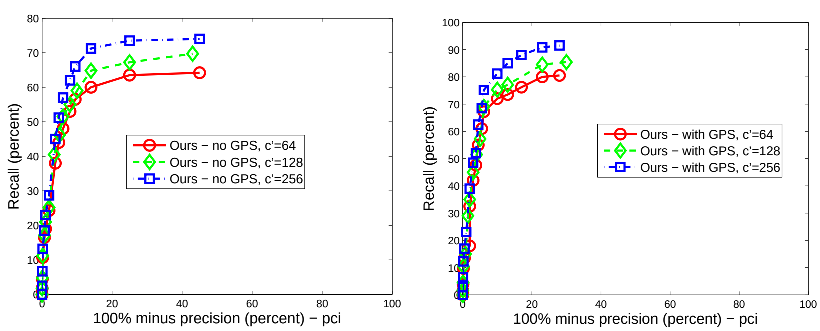

(a) (b)

Figure 3. Performance on San Francisco dataset PCI images with and without GPS, across different c' values. Performance improves with finer domains.

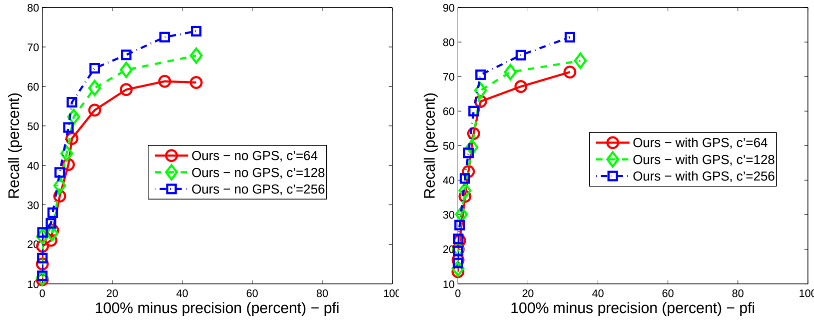

(a) (b)

Figure 4. Performance on San Francisco dataset PFI images with and without GPS, across different c' values. Performance improves with finer domains.

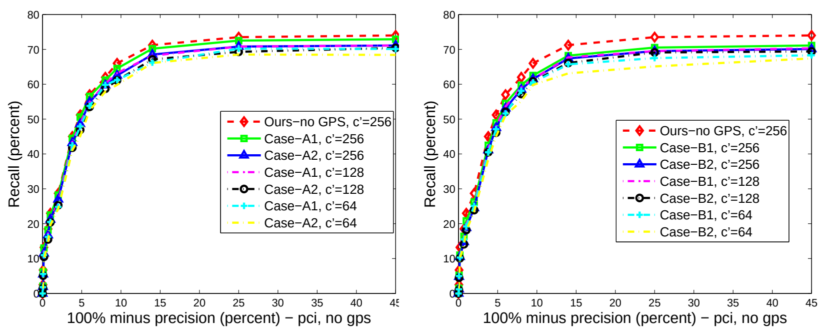

(a) (b)

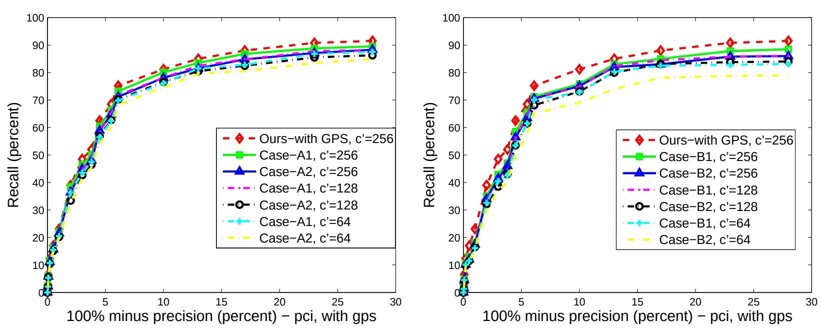

(c) (d)

Figure 5. Evaluating cases A1,A2,B1,B2 on the PCI images of San Francisco dataset, with and without GPS. The results parallel findings from the im2GPS dataset, which empirically reinforces some advantages of the design choices made in our approach for domain creation and for obtaining the embedded cross-domain representation.