Supplementary Material

Anonymous Author(s)
Affiliation
Address
`email`

Appendix: Parameter Sensitivity

Each of our non-linear extensions adds one additional hyper parameter to the LMNN problem. In section 6, we set these parameters by evaluation on a hold-out set. Here we explicitly examine their effect on the learned metric. For GB-LMNN, the new hyper-parameter is the regression tree depth. Figure 3(left) compares depths $4-7$ for several of the datasets evaluated in section 6. The figure depicts the ratio of kNN classification error for each depth setting to the kNN error of linear LMNN. GB-LMNN appears to be largely insensitive to tree depth within range.

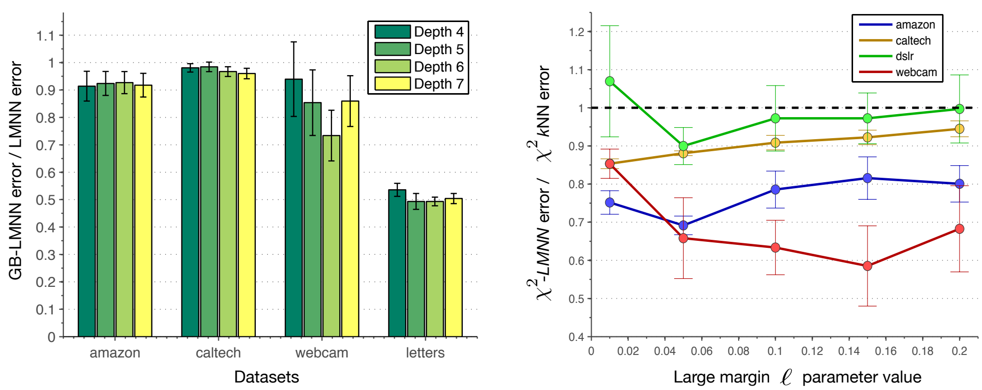

Figure 3: Parameter sensitivity measurements. *Left:* Varying tree depth for GB-LMNN. The measurement is the ratio between GB-LMNN error and LMNN error (lower is better). *Right:* Varying the large margin ℓ for χ^2-LMNN. The measurement is the ratio between χ^2-LMNN error and the χ^2 baseline error (lower is better).

For χ^2-LMNN, the additional hyper-parameter is the size of the large margin. Figure 3(right) examines several margin values: $0.01, 0.05, 0.10, 0.15$ and 0.20. The figure depicts the ratio of kNN classification error for each margin setting to the kNN error of the χ^2 distance baseline. For all but two settings, the transformation learned by χ^2-LMNN improves over the χ^2 baseline, generally by a large extent. However, the margin size parameter is clearly important to achieving the best performance. Fortunately, the parameter seems to be well-behaved and easily set by evaluation on a hold-out set.