K-Nearest Neighbors

Statistical Machine Learning (ENGG*6600*02)

School of Engineering, University of Guelph, ON, Canada

Course Instructor: Benyamin Ghojogh Summer 2023

k-Nearest Neighbors

k-Nearest Neighbors

- Consider the dataset $\{\mathbf{x}_i \in \mathbb{R}^d\}_{i=1}^n$.
- Some data points are close to each other and some are far from each other, in the *d*-dimensional Euclidean space. We can calculate the Euclidean distance between every pair of points in the space.
- Consider a point x_i. We can calculate the Euclidean distance of all points of dataset, except itself, from this point:

$$\|\mathbf{x}_1 - \mathbf{x}_i\|_2, \|\mathbf{x}_2 - \mathbf{x}_i\|_2, \dots, \|\mathbf{x}_n - \mathbf{x}_i\|_2.$$

- We sort these distances in ascending order and keep the k smallest distances. The corresponding k points with the smallest distances from x_i are the k-Nearest Neighbors (k-NN) of x_i. Note that k is a hyper-parameter positive integer for the number of neighbors.
- We do this procedure for all *n* points of dataset to have *k*-NN for all points.
- Consider a graph whose vertices are the indices of the points $\{1, 2, ..., n\}$ and an edge from vertex *j* to vertex *i* exists if x_j is one of the *k* nearest neighbors of x_i . Such a graph is called the *k*-**NN graph** of the dataset.

Adjacency Matrix

• Consider a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ whose (i, j)-th element is:

$$\mathbf{A}(i,j) := \begin{cases} 1 & \text{if } j \in k\text{-}\mathsf{NN}(\mathbf{x}_i) \\ 0 & \text{Otherwise.} \end{cases}$$
(1)

This matrix shows the indices of k-NN for all points. It is called the **adjacency matrix** or the k-NN matrix.

- Usually, the main diagonal of the adjacency matrix is zero or is ignored because we do not consider every point to be its own neighbor.
- If $k \ll n$, then the adjacency matrix is very sparse having lots of zeros.
- We can have the adjacency matrix between two datasets with size n_1 and n_2 . In this case, the size of matrix is $n_1 \times n_2$.
- Note that one may define the adjacency matrix to be A ∈ ℝ^{n×k} where every row lists the indices of the k-NN of x_i in the dataset.

k-NN for Machine Learning

k-NN for Machine Learning

- *k*-NN can be used for machine learning tasks.
- *k*-NN for **regression**:
 - For every input point x, we find the k-NN of x among the training data points.
 - The average or some statistics of the labels of the k-NN is used as the estimated label of the point x.
- *k*-NN for **classification**:
 - For every input point x, we find the k-NN of x among the training data points.
 - The majority of the class labels of the k-NN is used as the estimated class label of the point x.
- As explained above, *k*-NN for regression or classification does not have any training phase but it only has the test phase. The training phase of *k*-NN in libraries just takes the training data to be used later for *k*-NN calculation in the test phase.
- The *k*-NN classifier partitions the space into classes. The smaller the *k*, the more variance the decision boundary of classes has. Therefore, the small *k* leads to overfitting.
- In general, larger k in k-NN classifier or regressor results in less overfitting and better generalization. However, larger k has more computation. So, there is a trade-off between generalization and computation.

Kernel *k*-NN for Machine Learning

Distance in the Feature Space

 Let φ(.) be the pulling function from the input space to the feature space (reproducing kernel Hilbert space). Recall the kernel trick:

$$\boldsymbol{x}^{\top}\boldsymbol{x}\mapsto\phi(\boldsymbol{x})^{\top}\phi(\boldsymbol{x}). \tag{2}$$

• The squared Euclidean distance between points in the feature space (reproducing kernel Hilbert space) is [1]:

$$\|\phi(\mathbf{x}_i) - \phi(\mathbf{x}_j)\|_k^2 = k(\mathbf{x}_i, \mathbf{x}_i) + k(\mathbf{x}_j, \mathbf{x}_j) - 2k(\mathbf{x}_i, \mathbf{x}_j).$$
(3)

Proof:

$$\begin{split} \|\phi(\mathbf{x}_{i}) - \phi(\mathbf{x}_{j})\|_{k}^{2} &= \left(\phi(\mathbf{x}_{i}) - \phi(\mathbf{x}_{j})\right)^{\top} \left(\phi(\mathbf{x}_{i}) - \phi(\mathbf{x}_{j})\right) \\ &= \phi(\mathbf{x}_{i})^{\top} \phi(\mathbf{x}_{i}) + \phi(\mathbf{x}_{j})^{\top} \phi(\mathbf{x}_{j}) - \phi(\mathbf{x}_{i})^{\top} \phi(\mathbf{x}_{j}) - \phi(\mathbf{x}_{j})^{\top} \phi(\mathbf{x}_{i}) \\ &\stackrel{(a)}{=} \phi(\mathbf{x}_{i})^{\top} \phi(\mathbf{x}_{i}) + \phi(\mathbf{x}_{j})^{\top} \phi(\mathbf{x}_{j}) - 2\phi(\mathbf{x}_{i})^{\top} \phi(\mathbf{x}_{j}) \\ &\stackrel{(b)}{=} k(\mathbf{x}_{i}, \mathbf{x}_{i}) + k(\mathbf{x}_{i}, \mathbf{x}_{i}) - 2k(\mathbf{x}_{i}, \mathbf{x}_{i}), \end{split}$$

where (a) is because we can change the order of terms in inner product and (b) is because of the kernel trick.

Kernel k-NN for Machine Learning

- For calculation of the distance of points, we can pull the points to the feature space and then calculate their distances in that feature space. Then, *k*-NN can be calculated using the distances in the feature space. This algorithm is called **kernel** *k*-**NN**.
- Kernel k-NN can be used for machine learning tasks similar to how we use k-NN for machine learning.

Distance Metric Learning for Large Margin Nearest Neighbor Classification

Generalized Mahalanobis Distance

- *k*-Nearest Neighbor (*k*-NN) classification is highly impacted by the distance metric utilized for measuring the differences between data points.
- Euclidean distance does not weight the points and it values them equally.



 A general distance metric can be viewed as the Euclidean distance after projection of points onto a discriminative subspace [2]. This projection can be viewed as a linear transformation with a projection matrix denoted by *L*. We call this general metric the generalized Mahalanobis distance [2, 3, 4]:

$$\mathcal{D} := \|\mathbf{x}_i - \mathbf{x}_j\|_{\boldsymbol{M}}^2 := \|\boldsymbol{L}^\top (\mathbf{x}_i - \mathbf{x}_j)\|_2^2 = (\boldsymbol{L}^\top (\mathbf{x}_i - \mathbf{x}_j))^\top (\boldsymbol{L}^\top (\mathbf{x}_i - \mathbf{x}_j))$$

= $(\mathbf{x}_i - \mathbf{x}_j)^\top \boldsymbol{L} \boldsymbol{L}^\top (\mathbf{x}_i - \mathbf{x}_j) = (\mathbf{x}_i - \mathbf{x}_j)^\top \boldsymbol{M} (\mathbf{x}_i - \mathbf{x}_j),$ (4)

where $M := LL^{\top}$. The matrix M must be positive semi-definite, i.e. $M \succeq 0$, for the metric to satisfy convexity and the triangle inequality [5].

Large Margin Nearest Neighbor Classification

In order to improve the k-NN classification performance, we should decrease and increase the intra- and inter-class variances of data, respectively [6]. As can be seen in this figure, one way to achieve this goal is to pull the data points of the same class toward one another while pushing the points of different classes away.



- Let y_{il} be one (zero) if the data points x_i and x_l are (are not) from the same class. Moreover, let η_{ij} be one if x_j is amongst the k-nearest neighbors of x_i with the same class label; otherwise, it is zero.
- For tackling the goal of pushing together the points of a class and pulling different classes away, the following cost function can be minimized [7]:

$$\sum_{i,j} \eta_{ij} \| \boldsymbol{L}^{\top}(\boldsymbol{x}_{i} - \boldsymbol{x}_{j}) \|_{2}^{2} + c \sum_{i,j,l} \eta_{ij} (1 - y_{il}) \Big[1 + \| \boldsymbol{L}^{\top}(\boldsymbol{x}_{i} - \boldsymbol{x}_{j}) \|_{2}^{2} - \| \boldsymbol{L}^{\top}(\boldsymbol{x}_{i} - \boldsymbol{x}_{l}) \|_{2}^{2} \Big]_{+},$$
(5)

where $[.]_+ := \max(., 0)$ is the **standard Hinge loss**. The first term in Eq. (5) pushes the same-class points towards each other. The second term, on the other hand, is a **triplet loss** [8] which increases and decreases the inter- and intra-class variances, respectively.

Large Margin Nearest Neighbor Classification

• Eq. (5) was:

$$\sum_{i,j} \eta_{ij} \| \boldsymbol{L}^{\top}(\boldsymbol{x}_{i} - \boldsymbol{x}_{j}) \|_{2}^{2} + c \sum_{i,j,l} \eta_{ij} (1 - y_{il}) \Big[1 + \| \boldsymbol{L}^{\top}(\boldsymbol{x}_{i} - \boldsymbol{x}_{j}) \|_{2}^{2} - \| \boldsymbol{L}^{\top}(\boldsymbol{x}_{i} - \boldsymbol{x}_{l}) \|_{2}^{2} \Big]_{+},$$

 Inspired by support vector machines, the cost function (5) can be restated using slack variables:

$$\begin{array}{ll} \underset{\boldsymbol{M}, \, \xi_{ijl}}{\text{minimize}} & \mathcal{L} := \sum_{i,j} \eta_{ij} \, \|\boldsymbol{x}_i - \boldsymbol{x}_j\|_{\boldsymbol{M}}^2 + c \sum_{i,j} \eta_{ij} \left(1 - y_{il}\right) \xi_{ijl}, \quad \forall l \\ \text{subject to} & \|\boldsymbol{x}_i - \boldsymbol{x}_l\|_{\boldsymbol{M}}^2 - \|\boldsymbol{x}_i - \boldsymbol{x}_j\|_{\boldsymbol{M}}^2 \ge 1 - \xi_{ijl}, \\ & \xi_{ijl} \ge 0, \\ & \boldsymbol{M} \succeq \mathbf{0}, \end{array}$$

$$\begin{array}{l} \tag{6} \end{cases}$$

which is a SDP problem [9].

Large Margin Nearest Neighbor Classification

- We should solve this SDP problem using optimization toolboxes which use optimization algorithms such as the interior-point method.
- It is solved iteratively and usually slow to be solved. So, it is not used much in practice although its theory is solid.
- This method is called large margin metric learning for nearest neighbor classification [7, 10].
- I and my coauthors have a paper proposing **triplet mining** for this algorithm. See our paper "Acceleration of large margin metric learning for nearest neighbor classification using triplet mining and stratified sampling" [11].

References

- B. Schölkopf, "The kernel trick for distances," Advances in neural information processing systems, pp. 301–307, 2001.
- [2] B. Ghojogh, A. Ghodsi, F. Karray, and M. Crowley, "Spectral, probabilistic, and deep metric learning: Tutorial and survey," arXiv preprint arXiv:2201.09267, 2022.
- [3] B. Kulis et al., "Metric learning: A survey," Foundations and Trends (R) in Machine Learning, vol. 5, no. 4, pp. 287–364, 2013.
- [4] A. Globerson and S. T. Roweis, "Metric learning by collapsing classes," in Advances in Neural Information Processing Systems, pp. 451–458, 2006.
- [5] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.
- [6] B. Ghojogh, M. Sikaroudi, S. Shafiei, H. R. Tizhoosh, F. Karray, and M. Crowley, "Fisher discriminant triplet and contrastive losses for training siamese networks," in *IEEE International Joint Conference on Neural Networks (IJCNN)*, 2020.
- [7] K. Q. Weinberger, J. Blitzer, and L. K. Saul, "Distance metric learning for large margin nearest neighbor classification," in Advances in Neural Information Processing Systems, pp. 1473–1480, 2006.

References (cont.)

- [8] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 815–823, 2015.
- [9] L. Vandenberghe and S. Boyd, "Semidefinite programming," SIAM Review, vol. 38, no. 1, pp. 49–95, 1996.
- [10] K. Q. Weinberger and L. K. Saul, "Distance metric learning for large margin nearest neighbor classification," *Journal of Machine Learning Research*, vol. 10, no. Feb, pp. 207–244, 2009.
- [11] P. A. Poorheravi, B. Ghojogh, V. Gaudet, F. Karray, and M. Crowley, "Acceleration of large margin metric learning for nearest neighbor classification using triplet mining and stratified sampling," arXiv preprint arXiv:2009.14244, 2020.