

CHEAT SHEET: ALGORITHMS FOR SUPERVISED- AND UNSUPERVISED LEARNING ¹

ALGORITHM	DESCRIPTION	MODEL	OBJECTIVE	TRAINING	REGULARISATION	COMPLEXITY	NON-LINEAR	ONLINE LEARNING
<i>k</i>-nearest neighbour	The label of a new point \hat{x} is classified with the most frequent label $\hat{\ell}$ of the k nearest training instances.	$\hat{\ell} = \arg \max_C \sum_{i: x_i \in N_k(\mathbf{x}, \hat{x})} \delta(t_i, C)$	No optimisation needed.	Use cross-validation to learn the appropriate k ; otherwise no training, classification based on existing points.	k acts as to regularise the classifier: as $k \rightarrow N$ the boundary becomes smoother.	$\mathcal{O}(NM)$ space complexity, since all training instances and all their features need to be kept in memory.	Natively finds non-linear boundaries.	To be added.
Naive Bayes	Learn $p(C_k x)$ by modelling $p(x C_k)$ and $p(C_k)$, using Bayes' rule to infer the class conditional probability. Assumes each feature independent of all others, ergo 'Naive.'	$y(\mathbf{x}) = \arg \max_k p(C_k x)$ $= \arg \max_k p(x C_k) \times p(C_k)$ $= \arg \max_k \prod_{i=1}^D p(x_i C_k) \times p(C_k)$ $= \arg \max_k \sum_{i=1}^D \log p(x_i C_k) + \log p(C_k)$	No optimisation needed.	Multivariate likelihood $p(x C_k) = \sum_{i=1}^D \log p(x_i C_k)$ $p_{\text{MLE}}(x_i = v C_k) = \frac{\sum_{j=1}^N \delta(t_j = C_k \wedge x_{ji} = v)}{\sum_{j=1}^N \delta(t_j = C_k)}$ Multinomial likelihood $p(x C_k) = \prod_{i=1}^D p(\text{word}_i C_k)^{x_i}$ $p_{\text{MLE}}(\text{word}_i = v C_k) = \frac{\sum_{j=1}^N \delta(t_j = C_k) \times x_{ji}}{\sum_{j=1}^N \sum_{d=1}^D \delta(t_j = C_k) \times x_{dj}}$... where: <ul style="list-style-type: none"> x_{ji} is the count of word i in test example j; x_{di} is the count of feature d in test example j. Gaussian likelihood $p(x C_k) = \prod_{i=1}^D \mathcal{N}(v; \mu_{ik}, \sigma_{ik})$	Use a Dirichlet prior on the parameters to obtain a MAP estimate. Multivariate likelihood $p_{\text{MAP}}(x_i = v C_k) = \frac{(\beta_i - 1) + \sum_{j=1}^N \delta(t_j = C_k \wedge x_{ji} = v)}{ \alpha_i (\beta_i - 1) + \sum_{j=1}^N \delta(t_j = C_k)}$ Multinomial likelihood $p_{\text{MAP}}(\text{word}_i = v C_k) = \frac{(\alpha_i - 1) + \sum_{j=1}^N \delta(t_j = C_k) \times x_{ji}}{\sum_{j=1}^N \sum_{d=1}^D (\delta(t_j = C_k) \times x_{dj}) - D + \sum_{d=1}^D \alpha_d}$	$\mathcal{O}(NM)$, each training instance must be visited and each of its features counted.	Can only learn linear boundaries for multivariate/multinomial attributes. With Gaussian attributes, quadratic boundaries can be learned with uni-modal distributions.	To be added.
Log-linear	Estimate $p(C_k x)$ directly, by assuming a maximum entropy distribution and optimising an objective function over the conditional entropy distribution.	$y(x) = \arg \max_k p(C_k x)$ $= \arg \max_k \sum_m \lambda_m \phi_m(x, C_k)$... where: $p(C_k x) = \frac{1}{Z_\lambda(x)} e^{\sum_m \lambda_m \phi_m(x, C_k)}$ $Z_\lambda(x) = \sum_k e^{\sum_m \lambda_m \phi_m(x, C_k)}$	Minimise the negative log-likelihood: $\mathcal{L}_{\text{MLE}}(\lambda, \mathcal{D}) = \prod_{(x,t) \in \mathcal{D}} p(t x) = - \sum_{(x,t) \in \mathcal{D}} \log p(t x)$ $= \sum_{(x,t) \in \mathcal{D}} \left(\log Z_\lambda(x) - \sum_m \lambda_m \phi_m(x, t) \right)$ $= \sum_{(x,t) \in \mathcal{D}} \left(\log \sum_k e^{\sum_m \lambda_m \phi_m(x, C_k)} - \sum_m \lambda_m \phi_m(x, t) \right)$	Gradient descent (or gradient ascent if maximising objective): $\lambda^{n+1} = \lambda^n - \eta \Delta \mathcal{L}$... where η is the step parameter. $\Delta \mathcal{L}_{\text{MLE}}(\lambda, \mathcal{D}) = \sum_{(x,t) \in \mathcal{D}} \mathbb{E}[\phi(x, \cdot)] - \phi(x, t)$ $\Delta \mathcal{L}_{\text{MAP}}(\lambda, \mathcal{D}, \sigma) = \frac{\lambda}{\sigma^2} + \sum_{(x,t) \in \mathcal{D}} \mathbb{E}[\phi(x, \cdot)] - \sum_{(x,t) \in \mathcal{D}} \phi(x, t)$... where $\sum_{(x,t) \in \mathcal{D}} \phi(x, t)$ are the empirical counts. For each class C_k : $\sum_{(x,t) \in \mathcal{D}} \mathbb{E}[\phi(x, \cdot)] = \sum_{(x,t) \in \mathcal{D}} \phi(x, \cdot) p(C_k x)$	Penalise large values for the λ parameters, by introducing a prior distribution over them (typically a Gaussian). Objective function $\mathcal{L}_{\text{MAP}}(\lambda, \mathcal{D}, \sigma) = \arg \min_\lambda \left(-\log p(\lambda) - \sum_{(x,t) \in \mathcal{D}} \log p(t x) \right)$ $= \arg \min_\lambda \left(-\log e^{-\frac{(\lambda-\lambda)^2}{2\sigma^2}} - \sum_{(x,t) \in \mathcal{D}} \log p(t x) \right)$ $= \arg \min_\lambda \left(\frac{\sum_m \lambda_m^2}{2\sigma^2} - \sum_{(x,t) \in \mathcal{D}} \log p(t x) \right)$	$\mathcal{O}(INMK)$, since each training instance must be visited and each combination of class and features must be calculated for the appropriate feature mapping.	Reformulate the class conditional distribution in terms of a kernel $K(x, x')$, and use a non-linear kernel (for example $K(x, x') = (1 + \mathbf{w}^T x)^2$). By the Representer Theorem: $p(C_k x) = \frac{1}{Z_\lambda(x)} e^{\lambda^T \phi(x, C_k)}$ $= \frac{1}{Z_\lambda(x)} e^{\sum_{n=1}^N \sum_{i=1}^K \alpha_{nk} \phi(x_n, C_i)^T \phi(x, C_k)}$ $= \frac{1}{Z_\lambda(x)} e^{\sum_{n=1}^N \sum_{i=1}^K \alpha_{nk} K((x_n, C_i), (x, C_k))}$ $= \frac{1}{Z_\lambda(x)} e^{\sum_{n=1}^N \alpha_{nk} K(x_n, x)}$	Online Gradient Descent: Update the parameters using GD after seeing each training instance.
Perceptron	Directly estimate the linear function $y(x)$ by iteratively updating the weight vector when incorrectly classifying a training instance.	Binary, linear classifier: $y(x) = \text{sign}(\mathbf{w}^T x)$... where: $\text{sign}(x) = \begin{cases} +1 & \text{if } x \geq 0 \\ -1 & \text{if } x < 0 \end{cases}$ Multiclass perceptron: $y(x) = \arg \max_{C_k} \mathbf{w}^T \phi(x, C_k)$	Tries to minimise the Error function; the number of incorrectly classified input vectors: $\arg \min_{\mathbf{w}} E_P(\mathbf{w}) = \arg \min_{\mathbf{w}} - \sum_{n \in \mathcal{M}} \mathbf{w}^T x_n t_n$... where \mathcal{M} is the set of misclassified training vectors.	Iterate over each training example x_n , and update the weight vector if misclassification: $\mathbf{w}^{i+1} = \mathbf{w}^i + \eta \Delta E_P(\mathbf{w}) = \mathbf{w}^i + \eta x_n t_n$... where typically $\eta = 1$. For the multiclass perceptron: $\mathbf{w}^{i+1} = \mathbf{w}^i + \phi(x, t) - \phi(x, y(x))$	The Voted Perceptron: run the perceptron i times and store each iteration's weight vector. Then: $y(x) = \text{sign} \left(\sum_i c_i \times \text{sign}(\mathbf{w}_i^T x) \right)$... where c_i is the number of correctly classified training instances for \mathbf{w}_i .	$\mathcal{O}(INML)$, since each combination of instance, class and features must be calculated (see log-linear).	Use a kernel $K(x, x')$, and 1 weight per training instance: $y(x) = \text{sign} \left(\sum_{n=1}^N w_n t_n K(x, x_n) \right)$... and the update: $w_n^{i+1} = w_n^i + 1$	The perceptron is an online algorithm per default.
Support vector machines	A maximum margin classifier: finds the separating hyperplane with the maximum margin to its closest data points.	$y(x) = \sum_{n=1}^N \lambda_n t_n x^T x_n + w_0$	Primal $\arg \min_{\mathbf{w}, w_0} \frac{1}{2} \ \mathbf{w}\ ^2$ $\text{s.t. } t_n (\mathbf{w}^T x_n + w_0) \geq 1 \quad \forall n$ Dual $\tilde{\mathcal{L}}(\lambda) = \sum_{n=1}^N \lambda_n - \sum_{n=1}^N \sum_{m=1}^N \lambda_n \lambda_m t_n t_m x_n^T x_m$ $\text{s.t. } \lambda_n \geq 0, \quad \sum_{n=1}^N \lambda_n t_n = 0, \quad \forall n$	• Quadratic Programming (QP) • SMO, Sequential Minimal Optimisation (chunking).	The soft margin SVM: penalise a hyperplane by the number and distance of misclassified points. Primal $\arg \min_{\mathbf{w}, w_0} \frac{1}{2} \ \mathbf{w}\ ^2 + C \sum_{n=1}^N \xi_n$ $\text{s.t. } t_n (\mathbf{w}^T x_n + w_0) \geq 1 - \xi_n, \quad \xi_n > 0 \quad \forall n$ Dual $\tilde{\mathcal{L}}(\lambda) = \sum_{n=1}^N \lambda_n - \sum_{n=1}^N \sum_{m=1}^N \lambda_n \lambda_m t_n t_m x_n^T x_m$ $\text{s.t. } 0 \leq \lambda_n \leq C, \quad \sum_{n=1}^N \lambda_n t_n = 0, \quad \forall n$	• QP: $\mathcal{O}(n^3)$; • SMO: much more efficient than QP, since computation based only on support vectors.	Use a non-linear kernel $K(x, x')$: $y(x) = \sum_{n=1}^N \lambda_n t_n x^T x_n + w_0$ $= \sum_{n=1}^N \lambda_n t_n K(x, x_n) + w_0$ $\tilde{\mathcal{L}}(\lambda) = \sum_{n=1}^N \lambda_n - \sum_{n=1}^N \sum_{m=1}^N \lambda_n \lambda_m t_n t_m x_n^T x_m$ $= \sum_{n=1}^N \lambda_n - \sum_{n=1}^N \sum_{m=1}^N \lambda_n \lambda_m t_n t_m K(x_n, x_m)$	Online SVM. See, for example: <ul style="list-style-type: none"> <i>The Huller: A Simple and Efficient Online SVM</i>, Bordes & Bottou (2005) <i>Pegasos: Primal Estimated sub-Gradient Solver for SVM</i>, Shalev-Shwartz et al. (2007)
<i>k</i>-means	A hard-margin, geometric clustering algorithm, where each data point is assigned to its closest centroid.	Hard assignments $r_{nk} \in \{0, 1\}$ s.t. $\forall n \sum_k r_{nk} = 1$, i.e. each data point is assigned to one cluster k . Geometric distance: The Euclidean distance, l^2 norm: $\ x_n - \mu_k\ _2 = \sqrt{\sum_{i=1}^D (x_{ni} - \mu_{ki})^2}$	$\arg \min_{r, \mu} \sum_{n=1}^N \sum_{k=1}^K r_{nk} \ x_n - \mu_k\ _2^2$... i.e. minimise the distance from each cluster centre to each of its points.	Expectation: $r_{nk} = \begin{cases} 1 & \text{if } \ x_n - \mu_k\ _2^2 \text{ minimal for } k \\ 0 & \text{o/w} \end{cases}$ Maximisation: $\mu_{\text{MLE}}^{(k)} = \frac{\sum_n r_{nk} x_n}{\sum_n r_{nk}}$... where $\mu^{(k)}$ is the centroid of cluster k .	Only hard-margin assignment to clusters.	To be added.	For non-linearly separable data, use kernel k -means as suggested in: <i>Kernel k-means, Spectral Clustering and Normalized Cuts</i> , Dhillon et al. (2004).	Sequential k -means: update the centroids after processing one point at a time.
Mixture of Gaussians	A probabilistic clustering algorithm, where clusters are modelled as latent Gaussians and each data point is assigned the probability of being drawn from a particular Gaussian.	Assignments to clusters by specifying probabilities $p(x^{(i)}, z^{(i)}) = p(x^{(i)} z^{(i)})p(z^{(i)})$... with $z^{(i)} \sim \text{Multinomial}(\gamma)$, and $\gamma_{nk} \equiv p(k x_n)$ s.t. $\sum_{j=1}^k \gamma_{nj} = 1$. I.e. want to maximise the probability of the observed data \mathbf{x} .	$\mathcal{L}(\mathbf{x}, \pi, \mu, \Sigma) = \log p(\mathbf{x} \pi, \mu, \Sigma)$ $= \sum_{n=1}^N \log \left(\sum_{k=1}^K \pi_k \mathcal{N}(x_n \mu_k, \Sigma_k) \right)$	Expectation: For each n, k set: $\gamma_{nk} = p(z^{(i)} = k x^{(i)}; \gamma, \mu, \Sigma) \quad (= p(k x_n))$ $= \frac{p(x^{(i)} z^{(i)} = k; \mu, \Sigma) p(z^{(i)} = k; \pi)}{\sum_{j=1}^K p(x^{(i)} z^{(i)} = j; \mu, \Sigma) p(z^{(i)} = j; \pi)}$ $= \frac{\pi_k \mathcal{N}(x_n \mu_k, \Sigma_k)}{\sum_{j=1}^K \pi_j \mathcal{N}(x_n \mu_j, \Sigma_j)}$ Maximisation: $\pi_k = \frac{1}{N} \sum_{n=1}^N \gamma_{nk}$ $\Sigma_k = \frac{\sum_{n=1}^N \gamma_{nk} (x_n - \mu_k)(x_n - \mu_k)^T}{\sum_{n=1}^N \gamma_{nk}}$ $\mu_k = \frac{\sum_{n=1}^N \gamma_{nk} x_n}{\sum_{n=1}^N \gamma_{nk}}$	The mixture of Gaussians assigns probabilities for each cluster to each data point, and as such is capable of capturing ambiguities in the data set.	To be added.	Not applicable.	Online Gaussian Mixture Models. A good start is: <i>A View of the EM Algorithm that Justifies Incremental, Sparse, and Other Variants</i> , Neal & Hinton (1998).

¹Created by Emanuel Ferm, HT2011, for semi-procrastinational reasons while studying for a Machine Learning exam. Last updated May 5, 2011.