Bag of words: unordered collection of words ["John": 1,"likes": 7, "viagra": 3} Jerm weightening: tf-sdf > How important a word is to a document . tf = # in document idf(t,D)=log[IDI/KdeD:ted] - D=all docs => tf &f(t,d,D)=tf(t,d) o idf(t,D) [Stop-Words] Similarity of Dectors: Sim (9x,d)= I(tix+tij). Hultinomia): events are independ. DOC IS multinom outcome of words. P(W=n1,..., WK=NK) N, O1,...,OK)=1211-17KI · O11... OKK Classification: KNU: needs a distance metric. Determine K. -> Majority vote ... between the K nearest. Possible weighting: 1/dist. All computation deffered until classification. A popular lable may dominate due to its quantity - try to weight closest items to overcome. Association Rule Mining: Conions, Potatos 3 => [burgers. I = [1,000, Inf I tems (boolean) D= {t1, ..., tm} transactions. Rule: X=> Y where X, YEI and XNY= Ø. MinSup, HinConf. Support Supp(X) = Proportion of transactions containing Hemset X. Confidence Conf (X >> Y) = = Supp(XUY)/supp(X) = P(YIX). Luft = lift (X=>Y) = Supp(XUY)/(Supp(X)·Supp(Y)) Lift>1'= 1011 new, ochit<1=1850, Lift=1=1000 rom. Aprion's Principle: If an stemset is frequent then all of its subsets must also be frequent. {A,B,C,D} Conf(ABC > D) > Conf(AB > CD) > ... (A > B) L_1 ft $(X \rightarrow Y) = P(Y|X)/P(X)$ \rightarrow D(A,B) \leq D(AC) + D(C,B). Manhatan(Q,b)=2'(Q;-b;) Clustering: Distance Properties: D(A,B)=D(B,A); D(A,A)=0; D(A,B)=0=>A=B; results in a dendogram. K-means clustering: aims to partition n observations into K clusters. Given n observations (X1,000, Xn) each x is a dimentional vector, partition to K sets S={S1,000,Sk} minimize the "Within cluster sum of squares" argmin 2 27 1x-M; 12. M; = mean of points in S; The algorithm alternates betweet two steps until convergence : (1) Assignment step: $S_{\tau}^{\tau} = \{x_{\rho}: |x_{\rho}-m_{\tau}^{\tau}|^2 \leq |x_{\rho}-m_{\tau}^{\tau}|^2 \forall 1 \leq J \leq K\}$ (2) Update Step: $m_{\tau}^{\tau+1} = \frac{1}{|S_{\tau}^{\tau}|} \cdot \sum_{j \in S_{\tau}^{\tau}} \cdot \sum_{j \in S_{\tau}$ assumes clusters are of similar sizes. for different sizes the EM is better (genzation of kinearly K-medolds: PAM+> Paratitioning Around Medolds: Silhouette-tool for determining K. Algorithm: (1) Initialize: randomly select K of the n data-points as the medoids 2) Associate each data-point to the closest medoid. (3) for each medoid m: (3.1) For each non-medoid data point 0 (3.1.7) Swap m and o and comput the 391 Silhouette (Clustering): How well each object lies within its cluster. (x(i)) is average distance of object; with all other objects within its own cluster. b(i) is lavest average distance of i to all other clusters. the cluster with distance b(i) is the neighboring cluster (best fit after current cluster) -> S(1) = [b(1) - ox(1)]/max{ox(1), b(1)} -145(9)=1. 1)3 great, (-1) is worst, 0 point is on the Kiss-border between the neighboring clusters Agglomerative hierarchical Clustering: in begining each element is a cluster of its own. Coex, yell Single-linkage: D(X,Y)=mind(x,y). Complete linkage: D(X,Y)=maxd(x,y). Average: AI.IY) · II d(x,y) Decision Tree Learning: Data comes in records of form (X,Y) = (X1,X2,X3,...,XK, Y) Work top-down by choosing a variable at each step that best splits the set of items. Gin, Impurity: 15 a measure of how often a randomly chosen element from the set would be incorrectly labled if it were בא האופציות שמת המשליין שלפין הופקים השום הפיצול כפטיל. לשימ לכ לאופל המכנה בצל המלה ארים הראשונים הנותצים האופן חנקני, הראשין חים הקרוב לכולם ואחרין אותו הקרוב לכולם ואחרין

Bayesian inference: Bayes rule: P(H)E) = P(E) (Evidence There are competing hypotheses from which one chooses the most probable.

NOUSE Bayes Classifiers assumes that the presence or absence of a particular feature is unrelated to any other feature, given the class isacrable. The probability model for a classifier is conditional & P(C Fi,..., Fn) - Class; Feature, Posterior = Prior likelihood P(C|F1,000, Fn) = P(C).P(F1,000, Fn | C)/P(F1,000, Fn) - Denom 15 Const. P(CIFI..., Fn) = const · P(C) · JI, r P(F; IC) = Naive assumption. All model parameters & class Priors and feature probability distributions) can be opporoximated with relative frequencies from the training set. Class prior > 1 #c equally for all or 2 #C9/#C +> Relative Frequency of C9 For discrete features > Multinomial and Bernoulli distributions are popular to assume. Bernouls $\leftrightarrow X \sim \text{Ber}(p) \leftrightarrow P(X=1)=P$. Multinomial $\leftrightarrow P(x_1,...,x_K; n, p_1,...,p_M) = \frac{n!}{x_1!...x_K!} \cdot p_1^{x_1} \cdot p_2^{x_1} \cdot \dots \cdot p_K^{x_K}$ For continuos features Gausian dist assumed. P(X=V)C) = Vall or exp(-[v-Mc]2/202] MAP Sample-Correction is needed when one of the fregs is o in training set. Maximum Aposteriori Decision Rule Combination of features which characterizes or separates two or more classes of objects. Fundemental assumptions the independent variables are normally distributed. For 2 classes: Assumes P(X)y=0) and P(X)y=1) are both normally distributed With mean and colgariance ($M_0 \Sigma y=0$) and ($M_1, \Sigma y=1$). Under this assumption, the Bayes optimal solution is to predict y=1 if the kg of the likelihood ratios is below some treshold $J: (X-M_0)^T \Sigma_{y=0}^{-1} (X-M_0) + \ln(\Sigma_{y=0}) - (X-M_1)^T \Sigma_{y=1}^{-1} (X-M_1) - \ln|\Sigma_{y=0}| < T$ Without further assumptions, the above is aDA. LDA assumes Iy=0= [y=1=5] In this case Several terms cancel and the above becomes W.X.> C for some threshold constant C, where Was; 1 (M,-Mo) ⇒ Criterion of an input X being of class y is purely a function of this Imear combination of the known observations. In geometrical terms we project a space point on onto vector withus we only consider its direction) in other words, the observation belongs to y if core goonding I is located on a certain side of a hyperplane perpendicular to w. the To cortion of the plane is defined by the threshold C. Fisher's linear discriminant & Suppose 2 Classes of observations have mean's My=0, My=1 and covariances Zy=0, Zy=1 then the linear Combination of features w. X will have means w. My= , and variances wisy ; for i = 0,1 Fisher defines the separation as the ratio of the variance between the classes to the parance within the classes: $S = \frac{6}{8}$ Edween $\frac{6}{8}$ Within $= \frac{1}{8} \cdot \frac{1}{8} \cdot \frac{1}{9} \cdot$ approx same distr, a good choice of C will be a hyperplane between the two means: C=W. 1/2 (My=0+ Ply=1)=1/2 Ply= 2-1 Ply= - & My=0 I-1 Ply=0.
Logistic Regression: measures the relationship between a categorical dependant variable and one or more independant variables (usually continuos), by using Probability score as the predicted values of the dependent variable. Logistic regression is used to predict the odds of result success based on the values of the independent variables (predictors). The odds are defined as the probabil that a particular outcome is success devided by the probability that its failure. Logistic Function: $F(t) = e^t/(e^t+1) = 1/(1+e^{-t}) \in [0..1]$ Where t is a linear function of an explanatory variable X. It can be written as : JT(x) = exp(Bo+B1X)/[exp(Bo+B1X)+1] = Probability for success. The inverse of the logistic function: g(x) = ln[f(x)/(y-J(x))] = Bo + B1X

(NC3) 27220 (1860) 2100" (NC 2880) C (NESCO Soft Margin SM 2128: N'26022 DE (NIXNEN תכניא להקלוב בהליה ובגובה כשועות. צבער עונג הגבלה א משביא אני עם המשעובים בהבסצה עת שעש הביא שנט מהטיה נעוכים שנט הפועוט. ברגיסים פועינעות הצפע הפרבה ל ושבים פועוט ועורים אט ההטיה Rus ANN-Artificial Neural Networks: Perceptrons: Outpute (-1,1) O(X1,000, Xn) = 1 If Wo+WiX1+ W2/2+000+ WnXn O and (-1) otherwise. (-Wo-1) is the threshold. We imagine Ko=1=> O(4...xn)=1 If ∑W;X;>0 => O(X)=sgn(W·X). Perceptron training rule: We - Wi + DW, Where DW, = n (t-0) X, t=target; 0=output; n=learning rate-must be small! (Step)

When we need to ensure a differentiable threshold unit: $0 = \sigma(\vec{w}.\vec{x}) = \int_{-\infty}^{\infty} (ta-0a) \cdot (-x_{9a})$ Back-propagation; Multilayer-Network. We redefine $E \neq 0$ sum the errors over all of the network output units: $E(\vec{w}) = \frac{1}{2} \sum_{a=0}^{\infty} \sum_{b=0}^{\infty} x_{b} = \frac{1}{2} \sum_{a=0}^{\infty} x_{b} = \frac{1}{2} x_{b} = \frac{1$ to J is denoted as XJ; , Wy; . (3) for each < x, E> in training examples Do: (1) input the instance X to the network and compute the output Ou for every unit u in the network. (2) Propagate the errors back through the network by: For each network output unit k, calculate its error term 5k $\delta_{\kappa} \leftarrow O_{\kappa} (1-O_{\kappa})(t_{\kappa}-O_{\kappa})(3)$ For each hidden unit h, calculate its error term δ_{h} Sh-Oh (1-Oh): Dikcoutouts Wich Jik 9 Update each network weight Wil- Wji + A Wj; where AWji = 2 Stxj; . Figure se union The FM Algorithms used when only a subset of the relevant instance features are observable. Example: finding the means M, M2 Of K=2 Gaussians where 6,2 are known. The task is to output hypothesis h= KM...MX We want h that maximizes P(DIh). We can think of the field description of each instance as (Xi, Zi, Ziz) Where Xi is the observed value and Z; are indicators of wher xi came from (Which gaussian) Z; are hidden variables. EH algorithm searches for a maximum likelihood hypothesis by repeatedly re-estimating the expected values of the hidden variables Zir given the current hypothesis < M1,000, Mx> then recalculating the ML-hypothesis using these expected values for the hidden variables. First set $h = \langle \mu_1, \mu_2 \rangle$ arbitrarily. Then, iterativity re-estimate h by repeating 2-steps until convergence. Step 18 Calculate expected value $E[Z_{ij}]$ for each hidden variable Z_{ij} assuming $h = \langle \mu_1, \mu_2 \rangle$. Step 2 & Calculate a new ML-hypo h = SMI, ME > assuming value taken by each hidden variable Zij Is E[Zij] calculated in step 1, then replace h with h and iterate. In our example, [Zij] Is just h'revised. EM searches for ML-hypo h' by seeking h' that maximizes E[ln P(Ylh)]. We define a function Q(h'lh) = E[ln P(Ylh')ln, X]. Step 1: Estimation (E) step: Calculate Q(h'lh) using the current hypo h and observed data X to estimat prob-dist over Y. Q(hih) ElenP(YIh) Ih, X) Step 2: Maximization (M) Step: replace h with h that maximizes this Q function htargmax Q(KIH) באשנות אופר באינות אינה אונית שניבא אינית שתישונים משוצים מעובר באינית בים . Unsupervised EM שונוש משוצים מעובר $N(X)/(k, \Sigma_k) = \frac{1}{(2\pi)^{4}} \cdot (2\pi)^{4} \cdot (2\pi)^{4}$. ביכו ביכו בינו ליוניו ליוניו לא של אל אל אל אוצה בינו שומה של אל אל אוצה ליונים הואל הואל אוצה בינו בינו בינו $P(X_i) = \sum P(X_i, r_{jk}) = \sum P(r_{ik}) \cdot P(X_i | r_{jk}) = \sum \mathcal{I}_{ik} \mathcal{N}(X_i | \mu_{ik}, \Sigma_{ik}). \quad \mathcal{T}(r_{jk}) = P(r_{ik} + r_{jk}) = \sum \mathcal{T}(r_{ik}) \cdot P(X_i | r_{jk}) = \sum \mathcal{T}(r_{ik}) \cdot$ $P(X_{i}) = P(Y_{i}, Y_{i}) =$ N-Number of observations, K-Number of classes, Nt-Number : Decision Trees univer Of observations at node t, NE-Number of observations from class k at node t, P(KIt)proportion of observations from class κ at node t. $\hat{P}(\kappa,t) = N_{k}^{\kappa}/N_{k}$. Y(t)-class assigned to the terminal node t. Entropy(t) = $-\sum_{k=1}^{\kappa}\hat{P}(\kappa|t)\log_{2}(\hat{P}(\kappa|t))$, G_{INI} Index(t) = $1-\sum_{k=1}^{\kappa}\hat{P}^{2}(\kappa|t)$ Misclassification error t) = $1-\max_{k=1,\ldots,\kappa}\{\hat{P}(\kappa|t)\}$, IG(s,A) = Entropy(s) - $\sum_{v\in\mathcal{V}_{k}}\sum_{k=1}^{k}\hat{P}(\kappa|t)$.

P(Y=1|X=x)=h(WTx)=1+e-WTX (WTX): 128 P(Y|X) UIC UIDIEI D'87NN & (N'60'C18 D'000) P"ONOPN Sensitivity & Specificity / 8 ere 250 DANI DUDEDRO T-D SO WIC /PON B(X) =1-14 P(Y=1/X)=h(WTX) > ~ Specificity = True negative/(True neg + false pos) Sensitivity = True positive/(true positive + false negative) for the 2 בפשעת כי ת שמעו אינעו אנצטיםי עונגל -> כאשר יש תעון עוני מיעון ושנת כנציות אחסון לפצ"ת חישוב כי ים במתחבת מכל מעופיות ולא מן חלקן. Gausian Uxture ובאור Clustisso ובאור כושל באל ווmit Clustisso ובאורים מסוגברונים Hierarch - Clustering. The BID'E SAPS CHIEFER CHEIRER CHIEFER CHORICE INSTERN ISE AUSTUR DE DE 1251 פיעי בול עתת יכוד להחנים בתוציות שועת שועות באלה ין ולא רך כאכת איוזופי הצובים, 25. תחללית בחללים בפיע שעתיקים בפש לה בשל סכום האגרוביות של ערכי הבשל Y (של + בעל בענת החדשלה העכעה שם האגרובית תוא כעת עכני צין יש צער תצבק תסבצים של האל התפתל לפצעת כעת יצי ק"מיה (בין אה צים חו מו) ענונו ב מרכים של שרכים רציפים אונינים את נתוני השיעון שפי השרק הרצים והותרים את על ההפרבה (צ'י עונוצא) בין ב מרכים של שרט על הפרבים על אינינים אונינים אוני P(1/X)> 2+B PR 1-8 c/105 RM & ONO ST VIOLO ST VIOLO SE DOPUD 18 CUCAIR OUD SES - Bayes Classifier I DICINI SE DIBIDAD GIED JTO \leftarrow NOIVE BOYES. B. P(1|X)>X. P(2|X) DIC 1 DISTONDE ANOS SIDE DE SODE Y-1 X SE D'DEACH D'DYSD SOS P(X,Y)-1 P(X,1Y) UNDEACH SO VIC DENS D'D3. 8 9 IN 11:38 IN ILUIC VILL 8 1011C US et AC DE . J. f(X,X): CISN'OPUN LIC DINOSI 'XOION CIC DENS ανημον ριστη 85 de 18:00 gars / ερνε Ισημον ομορουν μαρ πουρουν του σε του ομον μουρουν μουρουν μουρουν μαρ του ομον μουρουν μαρ του ομονου שלובת תבועת ערופות צרכים בפושעה האנאר של בו היא שם אבל כב פיצוב שלים. (אופרטינים) כבל שניספר מצלים קא יותוכ Y=1 H y > τ : y υχώρδ θρουν τ = υντοισιδιδιούς εξενυνίετοι εδενούς μερο δου μερο δου σου μερο δου μερο δου μερο σου μερο συ μερο εξεν μερο μερο και μερο κα מש ב- אאל ביים אוע שאירה שם אול שונים של ארשונים א"י הורבת מיטונים א"י הורבת מיטונים אול ביים אואלים שקם אולים שקירה עקס בש הפצור העקורי ששתן לצורם שחצור. ADA - רוצים קו ושר שכשת בביות תואנה צאיו הן תוברפות האובן הרול. MICOLON - & ACT MARGINION (XIX) I JUNIA) Y COSO, C SISO) IN MICONION (XIX) בונקציית נכלות ב על של אין מפלים ההסתברעות בפי לעציא עלון כשפטון: משמשים הה לקירום שונשי בונקציה הוקצה. היו ליני על אין בשפטון: משמשים הה לקירום שונשי בונקצה. הוקצה היו ליני מנים בינים ולינים לינים אין הוקנים לינים אין הוקנים לינים אין הוקנים לינים אין הוקנים לינים לינים לינים אין הוקנים לינים ליני : 31/200 105 8188. y=0-8 p'aou D'D) pune 30 Xn+1=Xn-f(Xn)/f(Xn) 8800 128 X 713 UK DUIN K8 PLENDES