

Data Mining Methods used in Recommender Systems with R

Mudassir Makhdoomi

Abstract— Recommender Systems are processing tools which provide recommendations to people on various products. In this paper we aim to study some common data mining methods that have been successfully used in the Recommender Systems and simultaneously illustrate the methods by plotting them using various packages of R statistical programming language. Our focus will be on some commonly used classification methods: *Entropy and Information Gain* for selecting the most informative attribute(s) of the given data set, *Naïve Bayesian Classifiers* for predicting the class label when the attributes of the data set are independent of each other and finally *Support Vector Machines*, a geometric classification method.

Index Terms— Classification, Entropy, Naïve Bayesian, Recommender Systems, Similarity Measures and Data reduction, SVM.

I. INTRODUCTION

Recommender Systems (RS) are processing tools which provide recommendations to people on various products like books, movies, music and several other shopping products [1]. They are simply software tools which provide suggestions to customers which suit their needs most. RS work on two strategies: *content filtering*: RS creates a separate profile for each customer, reflecting his nature and *collaborative filtering*: RS uses the past transactions done by the customer to provide recommendations [1] [2] [3]. The paper is organized as follows. We start with a brief overview of Recommender Systems and their workings. Next, we describe some important data preprocessing methods focusing on use of similarity measures in Recommender Systems (Section A) and data reduction techniques (Section B). While describing data reduction strategies, we'll focus on one of the most important data reduction strategies: Dimensionality Reduction (Discrete Wavelet Transforms (Section B.1) and Principal Components Analysis or *Karhunen-Loeve Method* (Section B.2)). After that, various classification data mining techniques used in Recommender Systems are described (Section C). The classification techniques explained are *Entropy and Information Gain* (Section C.1), *Naïve Bayesian Classifiers* (Section C.2) and finally *Support Vector Machines* (Section C.3). Finally in the last section, Section D, we conclude the paper with emphasis on future work.

II. DATA PREPROCESSING METHODS

Data mining deals with large volumes of data [1]. There are various kinds of data that can be mined: database data,

transactional data, time-related or sequence data, data streams, spatial data, etc. [1]. In all kinds of data, a datum or a data object represents an entity. A datum is described by a set of attributes or characteristics. An attribute of a data object represents its characteristic or a particular feature. Ideally, all the attributes of all the data objects are expected to have all the corresponding values. However, the real-world data is incomplete and messy. It contains a lot of noise and needs to be preprocessed before it can be used in data mining and machine learning algorithms [1], [2]. In this section, we discuss 3 issues that are important for designing a recommender system. First, we discuss similarity or proximity measures; next we take up sampling of the data in case the data set is very large and finally we discuss some of the data reduction techniques.

A. Similarity or Proximity Measures

In Recommender Systems, we need to know how similar or alike or how dissimilar or different the data object are with respect to one another. Similarity measures are used to determine how similar or dissimilar the data objects are in comparison to one another [2]. Which similarity measure to use, depends on the type of data under consideration. For data objects with numeric attributes, the most commonly used similarity measure is the *Euclidean Distance*:

$$d(x, y) = \sqrt{\sum_{i=1}^n (x_i - y_i)^2} = \| \mathbf{x} - \mathbf{y} \|_2$$

In the above equation, \mathbf{x} and \mathbf{y} are the data objects with n attributes, x_i and y_i are the i^{th} attributes of the two data objects [1], [2]. Another well-known similarity measure is the *Manhattan or city block distance*:

$$d(x, y) = \sum_{i=1}^n |x_i - y_i| = \| \mathbf{x} - \mathbf{y} \|_1$$

Yet another similarity measure is the *Minkowski Distance*. It is actually the generalization of the Euclidean and Manhattan distance. It is given by:

$$d(x, y) = \sqrt[r]{\sum_{i=1}^n |x_i - y_i|^r} \text{ Where } r \in \mathbb{R}^+$$

Above, r is called the degree of the distance [2]. If we substitute $r = 1$, we get the *Manhattan Distance* and if $r = 2$ we get the *Euclidean Distance*. This distance measure is also called *L_r Norm*. Hence, *Manhattan Distance* is also called *L_1 Norm* and *Euclidean Distance* is also called *L_2 Norm* [3]. The distance between data objects with binary attributes is measured using a different metric. Data objects can be viewed as sets of features or attributes or

Manuscript received August 19, 2014.

Mudassir Makhdoomi, Department of Computer Applications, Islamia College of Science & Commerce, Srinagar, India, +919419754269.

characteristics. This is exactly the approach taken by another similarity measure called *Jaccard distance* or *co-efficient*. Two common operations on sets are the *union* and *intersection* of the sets. Suppose there are two data objects **X** and **Y**. Viewing the two data objects as two sets, the cardinal numbers of the union and intersection of **X** and **Y** are given by $|\mathbf{X} \cup \mathbf{Y}|$ and $|\mathbf{X} \cap \mathbf{Y}|$, respectively. The Jaccard distance gives the proportion of all the attributes or characteristics that are shared by the two data objects [3]. It is calculated as:

$$d_{jaccard}(x, y) = 1 - \frac{|\mathbf{X} \cap \mathbf{Y}|}{|\mathbf{X} \cup \mathbf{Y}|}$$

The above equation is one of the many forms of Jaccard Distance. Consider the following; let $A01 =$ the number of attributes where x is 0 and y is 1, $A10 =$ the number of attributes where x is 1 and y is 0, $A11 =$ the number of attributes where both x and y are 1 and $A00 =$ the number of attributes where both x and y are 0. Then the following metrics are available for calculating the similarity between data objects having binary attributes: *Simple Matching coefficient (SMC)*, *The Jaccard coefficient (JC)* and *The Extended Jaccard coefficient (or Tanimoto coefficient) (TC)* [1], [2].

$$SMC = \frac{A11 + A00}{A01 + A10 + A11 + A00},$$

$$JS = \frac{A11}{A01 + A10 + A11}, \text{ and,}$$

$$TC = \frac{x \cdot y}{\|x\|^2 + \|y\|^2 - x \cdot y}$$

Where \bullet is the vector dot products of two sets of attributed possessed by the data objects [2]. The matrix whose elements are the distance values of the set of data object pairs from the given data set is called a *Distance Matrix*. In order to visualize a Distance Matrix, a special diagram is used called *Voronoi Diagram* which divides a plane, containing n points, into cell, edges and vertices [4]. Fig. 1 shows a Voronoi Diagram of 10 numeric data points with 10 attributes each. The data points were generated using a uniform distribution. The diagram was generated using “*deldir*” package of *R statistical programming language* [5].

B.Data reduction techniques

High-traffic e-commerce websites like Amazon.com, eBay.com or Walmart.com or social networking sites like Facebook.com or Twitter.com (which use Recommender systems for suggesting friends, pages, followers and ads) generate a huge amount of data. It will take a long time to perform data analysis and mining on such amounts of data [1], [6]. Hence, methods need to be developed that can be used to represent data in a much more compact way, yet convey the same meaning as the original data. In other words, the data needs to be reduced in volume without any loss of information. Data reduction techniques are used for the above stated purposes. A Commonly used data reduction is called *Dimensionality*. The method of reducing the number of attributes or features of the data objects under consideration is called *Dimensionality Reduction*. Dimensionality of the data refers to the number of attributes of the give dataset. If we consider our data set represented as a 2-dimensional table

with columns representing the values of a particular attribute and the rows representing a specific data point/object, then replacing some columns of the data set with a few or even just one column is called *Dimensionality Reduction* [7]. Popular dimensionality reduction techniques include *Discrete Wavelet Transforms (DWT)* and *Principle Component Analysis (PCA)*.




Fig. 1: Voronoi Diagram of 10 numeric data points with 10 attributes each. The 2-D plane has been divided into 10 regions.




Fig. 2: A Haar Discrete Wavelet Transform applied to 1024 random numbers generated from a lognormal distribution

C.1 Discrete Wavelet Transform (DWT)

The Discrete Wavelet Transform (DWT) comes from the field of signal processing. However, it has been used widely in many statistical applications and also in data mining fields. In Recommender Systems, DWTs can be used as a data reduction technique. DWT is linear signal processing technique. When a DWT is applied to a given data object/vector (represented as a vector of features or attributes) it produces a “cardinally” equivalent, but numerically different data vector of *wavelet coefficients* [1]. Albeit, the two data vector are cardinally (lengthwise) equal, the usefulness of DWT arises in the fact that the wavelet transformed data can be trimmed by storing only a fraction of the strongest wavelet coefficients using some threshold value.

Once the wavelet coefficients are trimmed, the original data can be approximated by applying the inverse of the used DWT [1], [8]. Popular DWTs include *Haar-2*, *Daubechies-4* and *Daubechies-6* [1]. Fig. 2 shows a *Haar Discrete Wavelet Transform* applied to 1024 random numbers generated from a lognormal distribution using “wavelets” package of *R statistical programming language* [9].




Fig. 3: Screenplot of PCA of the *Arrests per 100,000 residents* which has 4 attributes; each bar corresponding to each attribute.

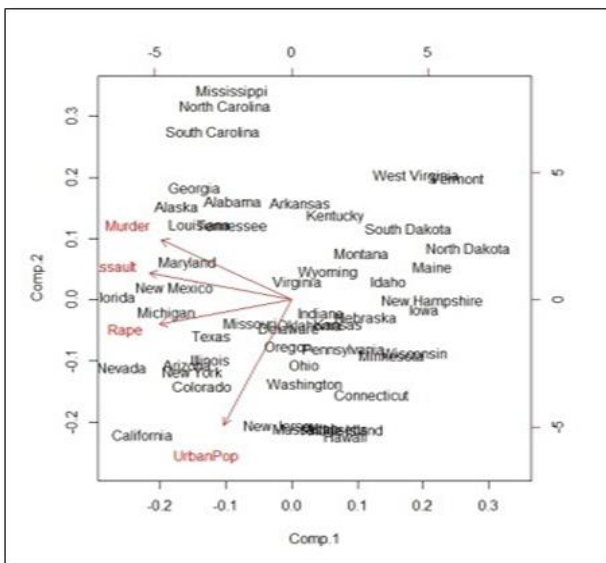


Fig. 4: Biplot of PCA of the *Arrests per 100,000 residents* which has 4 attributes. Notice the orthonormal vectors: principle components of this data set.

C.2 Principle Component Analysis or Karhunen-Loeve or K-L Method

In order to understand Principle Component Analysis (PCA) at an intuitive level, we will proceed first by using an example from our everyday life. People devise concepts like “he is a ‘good’ student”, but we can’t directly measure the concept of “goodness” or how “good” is someone. This, however, means that we internally reduce many attributes of someone to just one attribute defining them all - “good” [10]. This is exactly

how PCA works. PCA or K-L Method is the principle technique for dimensionality reduction in multivariate problems like recommending movies to a customer [11]. For recommending books, we’ve to consider many different attributes of the customer including his past history of movies. PCA works by finding k attributes among the n attributes of the given data objects, where $k \leq n$ such that the k attributes best represent the given data objects [1]. This way, the given data are interpolated to a much smaller dimensionality space. First of all, the given data are “normalized” to a within a common range, such that each all data fall within the chosen range. Next k orthonormal vectors are computed. These vectors are called *principle components*. These components essentially provide us with new axes for the given data. If these principle components are sorted in non-increasing order, they provide important information about the data variance. We used R programming language to do a PCA on a given data set of “*Arrests per 100,000 residents in 50 US states in 1973*” which has 4 attributes (Murder, Assault, Urban Population, and Rape). Fig. 3 shows a screenplot and Fig. 4 shows a biplot of the given data set.

C. Classification

Classification is a data analysis technique which classifies the given data set into various similar sub-classes. It creates a model called *classifier* which is used to predict (predictive modeling) the *class label*. Classification is a stepwise process which starts by creating a model (classifier) from the given data. The data used in this step is called *training set*. Training set has a number of tuples; with each tuple having the form $X = (x_1, x_2, \dots, x_n)$. This step is called *learning* or *training step*. The classifier “learns” from the training set which is made of up tuples and their associated class label(s). The associated class labels are called *target attributes*. The next step uses this model to predict the target attribute for the data instance which is to be classified. If the target attribute is available in the training set, the learning is called *Supervised learning* otherwise it’s called *Unsupervised learning* [1], [3], [7], [10], [2].

C.1 Entropy and Information Gain

In Supervised Classification problems, it is necessary to select the most valuable attribute(s). In other words, for predictive modeling, a set of attribute(s), among the given attributes, needs to be selected such that the selected set contains attributes providing important information about the target variable [3]. The selected set may contain a *single attribute* or it may contain *multiple attributes*. After selecting the informative attribute(s), we can divide or “classify” or “segment” the given data set into groups, in such a way that the resulting groups are distinguished on the basis of the target variable. It is preferable that the resulting groups be as pure or “homogenous with respect to the target variable” as possible [3]. For real-world data it is, usually, difficult and, sometimes, impossible to find such informative variables which result in pure groups. A method or technique which enables us to find such a set of attribute(s) is called a *purity measure*. An important purity measure, called *Entropy*, borrowed from Information Theory, pioneered by Claude Shannon, is widely used as a purity measure of the resulting groups [12]. Entropy is the measure of disorder of a system. In

our case, the system is a particular group formed from the given data set. Each group member (data instance or data point) will have a collection of attributes. In supervised classification, these properties link with the values of the target variable. Entropy of this group means how mixed or “impure” this group is with respect to target variables. Entropy can be defined as:

$$\text{entropy}, E = -p_1 \log(p_1) - p_2 \log(p_2) - \dots$$

In the above equation P_i is the probability (the relative percentage) of attribute i within the group. The probability, P_i ranges from 1 (when all group members have same values for attribute i) and 0 (when no group members have same values for attribute i) [3]. Once we’ve calculated the entropies of all the “child” groups, we’d also like to know how “informative” an attribute is w.r.t. the target variable. Toward this end, we calculate a metric called *Information Gain* which measures the amount of “improvement” in entropy due to the attribute used in grouping the data set. In other words, it measures the amount of decrease of entropy of a group or how much information is added to a group due to the data set partitioning attribute. Information Gain (IG) can be calculated as:

$$IG = E(\text{original}) - [p(g_1) \times E(g_1) + p(g_2) \times E(g_2)]$$

In the above equation $E(\text{original})$ is the entropy of the original data set $E(g_i)$ is the entropy of the i^{th} group and, $p(g_i)$ is the fraction/proportion of data instances belonging to that group [3] [10].

C.2 Bayesian classification

Bayesian Classifiers take a probabilistic approach to classification by predicting the class label of a given data object or tuple [11]. The basic working idea of the Bayesian classification is Bayes’ Theorem. The probability calculated to a Bayesian classifier is the conditional probability which is based on some background information [13]. We’ll provide a brief overview of Bayes’ Theorem followed by a simplified form of Bayesian classifier called Naïve Bayesian Classifier. Bayes’ theorem is based on conditional probability. A conditional probability is a probability that is based on some background information [13]. Using the provided information, we can calculate the probability of some other event. If A is the required background information for calculating some other event B , then the “conditional” probability of event B given A is written as $P(B|A)$. The probability $P(B|A)$ is calculated using Bayes’ Theorem as:

$$P(B|A) = \frac{P(A|B) \cdot P(B)}{P(A)}$$

$P(A|B)$, is *posterior probability* of A conditioned on B . Let X be data tuple described by the observations made on the set of n attributes. We can write X as $X = \{x_1, x_2, \dots, x_i, \dots, x_n\}$, where x_i is the value of i^{th} attribute of the give data tuple. A Bayesian classifier calculates the probability of the class label, C , to which the tuple, X , belongs to. Since X is the given background information, we can use, Bayes’ Theorem to calculate $P(C|X)$ as:

$$P(C|X) = \frac{P(X|C) \cdot P(C)}{P(X)}$$

The conditional probability in the numerator of the above equation $P(X|C)$ can be calculated empirically from the given training set. It is simply the frequency with which X occurs among the tuples/instances belonging to class C . This probability is really hard to compute in practice. Even if each attribute is only symmetrical binary (both the values 0 and 1 are equally important), the number of combinations for X is 2^n and grows with number of values an attribute can take [11]. In order to solve this problem, a simplification is done where it is assumed that all the attributes are independent of each other that $P(X|C)$ can be calculated as:
 $P(X|C) = P(\{x_1, x_2, \dots, x_i, \dots, x_n\} | C) = P(x_1|C) \cdot P(x_2|C) \dots P(x_n|C)$
 This is the basis of a Bayesian classifier called Naïve Bayesian Classifier [14]. The following algorithm summarizes how a Naïve Bayesian Classifier works [1].

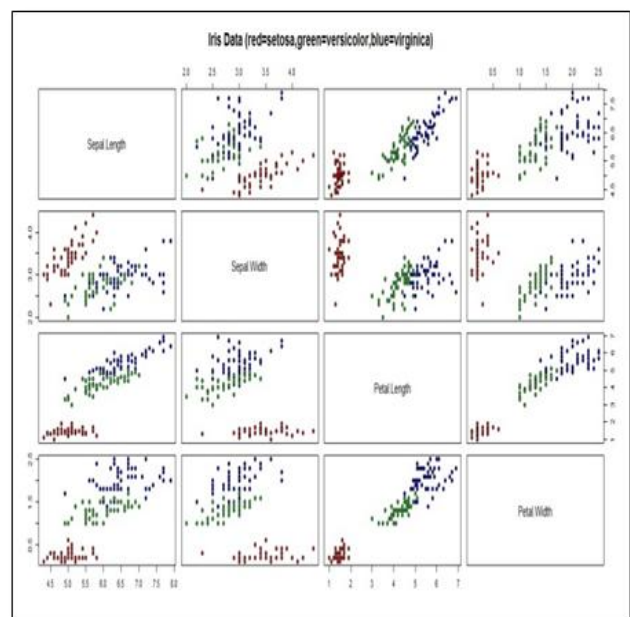


Fig. 5: Plot of the Iris dataset. Red is for Setosa, green for Versicolor and blue for Virginica.

a) Let T be the set of training tuples and their associated class labels. Let $X = \{x_1, x_2, \dots, x_n\}$ be a training tuple with n attributes $\{A_1, A_2, \dots, A_m\}$ and C_1, C_2, \dots, C_m be m class labels.

b) Calculate $P(X)$ which is constant for all classes.

c) Calculate

$$P(C_i) = \frac{|C_{i,T}|}{|T|}$$

Where $|C_{i,T}|$ is number of training tuples of class C_i in T .

d) If attribute A_k is a categorical attribute, calculate,

$$P(x_k|C_i) = \frac{|x_{k,T}|}{|C_{i,T}|}$$

Otherwise, if attribute A_k is continuous-valued, it is assumed to have a Gaussian distribution with a mean μ and standard deviation σ . Calculate:

$$g(x, \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

So that

$$P(x_k|C_i) = g(x_k, \mu_{C_i}, \sigma_{C_i}).$$

e) Calculate $P(\mathbf{X}|C_i) = \prod_{t=1}^n P(x_t|C_i)$.

This assumes all the attributes are independent of each other. This is called *class-conditional independence*.

f) Calculate

$$P(C_i|\mathbf{X}) = \frac{P(\mathbf{X}|C_i) \cdot P(C_i)}{P(\mathbf{X})}$$

Such that

$$P(C_i|\mathbf{X}) > P(C_j)$$

for $1 \leq j \leq m, j \neq i$

This maximizes $P(C_i|\mathbf{X})$.

Fig. 5 shows a plot, using *R programming language*, of the famous dataset called *Fisher's Iris data set* [15] also called *Anderson's Iris data set* [16]. The Iris dataset contains 50 samples of three species of Iris (*Iris setosa*, *Iris virginica* and *Iris versicolor*) along with observations of 4 attributes for each species: petal length, petal width, sepal width and sepal length. Although the data set is not related to recommender systems, it clearly illustrates how a Naïve Bayesian Classifier is used to predict the class label. The aim here is to predict the species (target class label) given the attribute values (a data tuple). For achieving this, we used "e1071" package of *R statistical programming language* [17]. After training our Naïve Bayesian Classifier with the Iris data set, we predicted the species using the 4 attributes. Fig. 6 shows the plot for the same.




Fig. 6: Iris species classified into 3 respective classes, each represented by a variable shade of gray color.




Fig. 7: Linearly separable two class classification with SVM. The hyperplane is associated with a small margin.




Fig. 8: Linearly separable two-class classification with SVM. The hyperplane is associated with a large margin.

C.3 Support Vector Machines (SVMs)

An SVM classifier is a geometrical method which can be used to classify both linear as well as non-linear data. An SVM classifier uses a non-linear function or mapping to transform a given data set into some other higher dimension where it, optimally, searches for a linear hyperplane or a decision boundary. This hyperplane is then used to classify (separate) the data set into classes. If we consider a simple linearly-separable 2-D two class classification problem, as shown in Fig. 7 and 8, it can be observed that several hyperplanes are possible. Notice that each hyperplane has an associated margin. The goal of an SVM is choose the hyperplane that maximizes the margin. This optimal hyperplane ensures misclassification is kept to minimum, if not completely eliminated [2]. For finding such maximum margin hyperplanes, a dot product between the two vectors is calculated. However, it is not always possible to linearly separate the given data set. Sometimes the data instances overlap in such a way that it is not possible to draw a straight hyperplane. For such data sets, we can either let the stray data instances to be misclassified upto a certain tolerance error

rate. This is done by introducing *Slack variables*. These variables associate a cost with each misclassified instance. Other solution is to draw a curved hyperplane, instead of a straight line. This is achieved, usually, by using *Kernelization* or “*kernel trick*”. The main idea is to replace the dot product by a more general function. Some commonly used kernel functions are Polynomial, Sigmoid and a family of Radial Basis Function (RBF) [1] [2] [11]. We will use a data set called “cats” for illustrating SVMs in R. The data set contains 144 rows and 3 columns (Sex, Bwt (Body Weight in kg) and Hwt (Heart Weight in kg)). 144 adult cats were used for experiments with the drug *digitalis*. Their heart and body weights were recorded. 97 of the cats were male and 47 were female [18]. We will use the same “*e1071*” package of *R statistical programming language* [17]. The results are shown in Fig. 9.

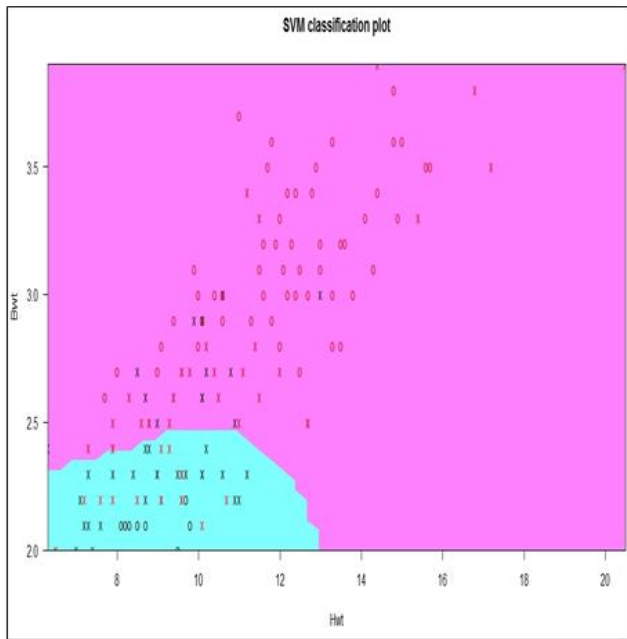


Fig. 9: SVM classification of “cats” dataset with *Hwt* as x-axis and *Bwt* as y-axis. The pink area represents the male cats and the sky-bluish color represents the female cats. The boundary between the two colors is the hyperplane with RBF used as kernel function.

D. Conclusion and Future work

The paper only enumerated and explained some of the commonly used data mining methods used in RS. There are many other methods and techniques which we have not covered. Bayesian Belief Networks, Decision Tree Induction, k-nearest neighbor, Artificial Neural Networks, Association Rule mining and many clustering techniques have an excellent scope for future study. There is also a future scope for evaluating the efficiency, performance and accuracy of various mining methods with relative advantages and disadvantages.

REFERENCES

[1] J. Han, M. Kamber and J. Pei, *Data Mining: Concepts and Techniques*, MA, USA: Morgan Kaufman, 2012.
 [2] F. Ricci, L. Rokach, B. Shapira and P. B. Kantor,

Recommender Systems Handbook, New York, USA: Springer, 2011.
 [3] F. Provost and T. Fawcett, *Data Science for Business*, Sebastopol, CA: O'Reilly, 2014.
 [4] F. Aurenhammer and R. Klein, *Handbook of Computational Geometry*, Amsterdam, Netherlands: North-Holland, 2000.
 [5] R. Turner, "deldir: Delaunay Triangulation and Dirichlet (Voronoi) Tessellation," 2 Febuary 2014. [Online]. Available: <http://cran.r-project.org/web/packages/deldir/index.html>. [Accessed 2 July 2014].
 [6] P. Berkhin, *Grouping Multidimensional Data*, Springer Berlin Heidelberg, 2006.
 [7] D. Conway and J. M. White, *Machine Learning for Hackers*, Sebastopol, CA: O'Reilly, 2012.
 [8] B. Vidakovic and P. Mueller, "WAVELETS FOR KIDS: A Tutorial Introduction," Institute of Statistics and Decision Sciences, Duke University, Durham, N, 1991.
 [9] E. Aldrich, "wavelets: A package of functions for computing wavelet filters, wavelet transforms and multiresolution analyses," 18 December 2013. [Online]. Available: <http://cran.r-project.org/web/packages/wavelets/wavelets.pdf>. [Accessed 5 July 2014].
 [10] R. Schutt and C. O'Neil, *Doing Data Science - Straight talk from the frontline*, Sebastopol, CA: O'Reilly Media, 2013.
 [11] P. K. Janert, *Data Analysis with Open Source Tools*, Sebastopol, CA: O'Reilly Media, 2012.
 [12] C. E. Shannon, "A Mathematical Theory of Communication," *The Bell System Technical Journal*, vol. 27, pp. 379-423, 1948.
 [13] A. B. Downey, *Think Bayes*, Sebastopol, CA: O'Reilly Media, 2013.
 [14] K. P. Murphy, *Naive Bayes classifiers*, University of British Columbia, 2006.
 [15] R. A. Fisher, "The use of multiple measurements in taxonomic problems," *Annals of Eugenics*, vol. 7, no. 2, p. 179-188, 1936.
 [16] E. Anderson, "The species problem in Iris," *Annals of the Missouri Botanical Garden*, vol. 23, no. 3, p. 457-509, 1936.
 [17] D. Meyer, E. Dimitriadou, K. Hornik, A. Weingessel, F. Leisch, C.-C. Chang and C.-C. Lin, "e1071: Misc Functions of the Department of Statistics (e1071), TU Wien," 17 March 2014. [Online]. Available: <http://cran.r-project.org/web/packages/e1071/index.html>. [Accessed 10 July 2014].
 [18] R Documentation, "Anatomical Data from Domestic Cats," [Online]. Available: <http://astrostatistics.psu.edu/su07/R/html/MASS/html/cats.html>. [Accessed 2014 July 2014].



Mudassir Makhdoomi is an Assistant Professor in Department of Computer Applications, Islamia College of Science & Commerce, Srinagar, India, +919419754269