A DATA TRANSFORMATION APPROACHFOR DIRECT AND INDIRECT DISCRIMINATION IN DATA MINING

Mafitha John.M
MEComputer Science and Engineering, Infant Jesus College of Engineering and Technology
mafithajohn@gmail.com

Mr.M.Varghese,M.Tech,(Ph.D)

Head of Department, Department of Computer Science and Engineering,

Infant Jesus College of Engineering and Technology

Abstract

Discrimination is the prejudicial treatment which involves denying opportunities to members of one group in favor of other groups. It is unfair to discriminate people because of their gender, religion, nationality, age and so on, especially when those attributes are used for making decisions about them like giving them a job, loan, insurance, etc. If the training data are inherently biased for or against a particular community,discriminatory decisions may ensue. Discovering the potential biases and eliminating them from the training data without harming their decision-making utility is therefore highly desirable which forms the primary goal of anti-discrimination techniques in data mining. Discrimination can be either direct or indirect. Direct discrimination occurs when decisions are made based on sensitive attributes. Indirect discrimination occurs when decisions are made based on non-sensitive attributes which are strongly correlated with biased sensitive ones. This paper aims at identifying potential discrimination using an acceptable level of discrimination.

Key words: discrimination, direct and indirect discrimination

I Introduction

Discrimination is a very important issue when considering the legal and ethical aspects of data mining. It is more than obvious that most people do not want to be discriminated because of their gender, religion, nationality, age and so on, especially when those attributes are used for making decisions about them like giving them a job, loan, insurance, etc. It involves denying to members of one group opportunities that are available to other groups. There is a list of **antidiscrimination** acts, which are laws designed to *prevent* discrimination on the basis of a number of attributes (**e.g., race, religion, gender, nationality, disability, marital status, and age**) in various settings (**e.g., employment and training, access to public services, credit and insurance, etc.**). At first sight, automating decisions may give a sense of fairness: classification rules do not guide themselves by personal preferences. However, at a closer look, one realizes that classification rules are actually learned by the system (**e.g., loan granting**) from the training data. If the training data are inherently biased for or against a particular community (**e.g., foreigners**), the learned model may show a discriminatory prejudiced behavior. In other words, the system may infer that just being foreign is a legitimate reason for loan denial. Discovering such potential biases and eliminating them from the training data without harming their decision making utility is therefore highly desirable. One must prevent data mining from becoming itself a source of discrimination, due to data mining tasks generating discriminatory models from biased data sets as part of the automated decision making. In [12], it is demonstrated that data mining can be both a source of discrimination and a means for discovering discrimination. Discrimination can be either direct or indirect (also called systematic).

1.1 Direct discrimination

Direct discrimination consists of rules or procedures that explicitly mention minority or disadvantaged groups based on sensitive discriminatory attributes related to group membership. Discriminatory (sensitive) attributes like gender, race, religion, etc.,

1.2 Indirect discrimination

Indirect discrimination consists of rules or procedures that, while not explicitly mentioning discriminatory attributes, intentionally or unintentionally could generate discriminatory decisions. Redlining by financial institutions (refusing to grant mortgages or insurances in urban areas they consider as deteriorating) is an archetypal example of indirect discrimination, although certainly not the only one. With a slight abuse of language for the sake of compactness, in this paper indirect discrimination will also be referred to as redlining and rules causing **indirect discrimination will be called redlining rules** [12].

Indirect discrimination could happen because of the availability of some background knowledge (rules), for example, that a certain zip code corresponds to a deteriorating area or an area with mostly black population. The background knowledge might be accessible from publicly available data (e.g., census data) or might be obtained from the original data set itself because of the existence of nondiscriminatory attributes that are highly correlated with the sensitive ones in the original data set. .

1.3Basic definition

Some basic definitions related to data mining [17]. After that, we elaborate on measuring and discovering discrimination.

- **A data set**is a collection of data (records) and their attributes. Let DB be the original data set.
- *An* **item**is an attribute along with its value, e.g., Race = black.
- *An* **item set**is a collection of one or more items, e.g.,{ Foreign worker = Yes, City = NYC}.
- *A* **classification rule**is an expression X -> C, where C is a class item (a yes/no decision), and X is an item set containing no class item, e.g., {**Foreign worker = Yes, City = NYC-> Hire = no**}. X is called the premise of the rule.
- The support of an item set, supp(X), is the fraction of records that contain the item set X. We say that a rule X -> C is completely supported by a record if both X and C appear in the record.
- The confidence of a classification rule, conf(X)->C, measures how often the class item C appears in records that contain X. Hence, if supp(X) > 0 then
- **Conf(X)->C =**$\frac{\boldsymbol{supp(X,C)}}{\boldsymbol{supp(X)}}$

Support and confidence range over (0,1)

- **A frequent classification rule**is a classification rule with support and confidence greater than respective specified lower bounds. Support is a measure of statistical significance, whereas confidence is a measure of the strength of the rule. Let FR be the database of frequent classification rules extracted from DB.
- **Discriminatory attributes and item sets***(protected by law):* Attributes are classified as discriminatory according to the applicable anti-discrimination acts (laws). For instance, U.S. federal laws prohibit discrimination on the basis of the following attributes: race, color, religion, nationality, sex, marital status, age and pregnancy (Pedreschi et al. 2008). Hence these attributes are regarded as discriminatory and the item sets corresponding to them are called discriminatory item sets. {Gender=Female, Race=Black} is just an example of a discriminatory item

set. Let *DAs* be the set of predetermined discriminatory attributes in *DB* and *DIs* be the set of predetermined discriminatory item sets in *DB.*

- **Non-discriminatory attributes and item sets:** If A*s* is the set of all the attributes in *DB* and *Is* the set of all the item sets in *DB,* then *nDA*s (*i.e.* set of *nondiscriminatory attributes*) is *As* - *DAs* and *nDIs* (*i.e.* set of *non-discriminatory item sets*) is *Is* - *DIs*. An example of non-discriminatory item set could be {**Zip= 10451, City=NYC**}.

II RELATED WORK

Some proposals are oriented to the discovery and measure of discrimination. The discovery of discriminatory decisions was first proposed by Pedreschi et al. [12], [15]. The approach is based on mining classification rules (the inductive part) and reasoning on them (the deductive part) on the basis of quantitative measures of discrimination that formalize legal definitions of discrimination. For instance, the US Equal Pay Act [18] states that: "a selection rate for any race, sex, or ethnic group which is less than four-fifths of the rate for the group with the highest rate will generally be regarded as evidence of adverse impact."

This approach has been extended to encompass statistical significance of the extracted patterns of discrimination in [13] and to reason about affirmative action and favoritism [14]. Moreover it has been implemented as an Oracle-based tool in [16]. Current discrimination discovery methods consider each rule individually for measuring discrimination without considering other rules or the relation between them.Three approaches are conceivable:**pre-processing, in processing and post-processing approaches**. We next describe these groups

2.1 Pre processing.

Transform the source data in such a way that the discriminatory biases contained in the original data are removed so that no unfair decision rule can be mined from the transformed data and apply any of the standard data mining algorithms. The preprocessing approaches of data transformation and hierarchy-based generalization can be adapted from the privacy preservation literature. Along this line, [7], [8] perform a controlled distortion of the training data from which a classifier is learned by making minimally intrusive modifications leading to an unbiased data set. The preprocessing approach is useful for applications in which a data set should be published and/or in which data mining needs to be performed also by external parties (and not just by the data holder).

2.2In-processing

Change the data mining algorithms in such a way that the resulting models do not contain unfair decision rules. For example, an alternative approach to cleaning the discrimination from the original data set is proposed in [2] whereby theNon discriminatory constraint is embedded into a decision tree learner by changing its splitting criterion and pruning strategy through a novel leafrelabeling approach. However, it is obvious that in processing discrimination prevention methods rely on new special-purpose data mining algorithms; standard data mining algorithms cannot be used.

2.3Post processing.

Modify the resulting data mining models, instead of cleaning the original data set or changing the data mining algorithms. For example, in [13], a confidence-altering approach is proposed for classification rules inferred by the CPAR algorithm. The post processing approach does not allow the data set to be published, only the modified data models can be published (knowledge publishing), hence data mining can be performed by the data holder only. One might think of a straightforward pre processing approach consisting of just removing the discriminatory attributes from the data set. Although this would solve the direct discrimination problem, it would cause much information loss and in general it would not solve indirect discrimination. As stated in [12] there may be other attributes (e.g., Zip) that are highly correlated with the sensitive ones (e.g., Race) and allow inferring discriminatory rules.

Preprocessing approach seems to be the most flexible one, it does not require changing the standard data mining algorithms, unlike the in processing approach, and it allows data publishing (rather than just knowledge publishing), unlike the post processing approach.

IIISYSTEM MODEL

There are two types of rules:
1. PD Rule
2. PND Rule

3.1Potentially discriminatory rule

A classification rule X -> C is potentially discriminatory (PD) when X = A, B with A is a discriminatory item set and B a nondiscriminatory item set

For example,(Foreign worker = Yes, City = NYC-> Hire = No). The word "potentially" means that a PD rule could probably lead to discriminatory decisions. Therefore, some measures are needed to quantify the direct discrimination potential.

3.2Direct discrimination measure

One of these measures is the extended lift (elift)Let A, B -> C be a classification rule such that Conf (B -> C) > 0. The extended lift of the rule is elift **(A, B -> C)** = $\frac{conf(A,B\rightarrow C)}{conf(B\rightarrow C)}$

The idea here is to evaluate the discrimination of a rule as the gain of confidence due to the presence of the discriminatory items (i.e., A) in the premise of the rule. Whether the rule is to be considered discriminatory can be assessed by thresholding elift as follows.

Let α € R be a fixed threshold and let A be a discriminatory item set. A PD classification rule c = A, B ->C c is taken as α - discriminatory w.r.t. elift if elift(c) > α. Otherwise is taken as α-protective,. The purpose of direct discrimination discovery is to identify α -discriminatory rules. In fact, α – discriminatory rules indicate biased rules that are directly inferred from discriminatory items (e.g., Foreign worker = Yes).

We call these rules direct α -discriminatory rules. In addition to elift, two other measures slift and olift were proposed by Pedreschi et al. in [13].

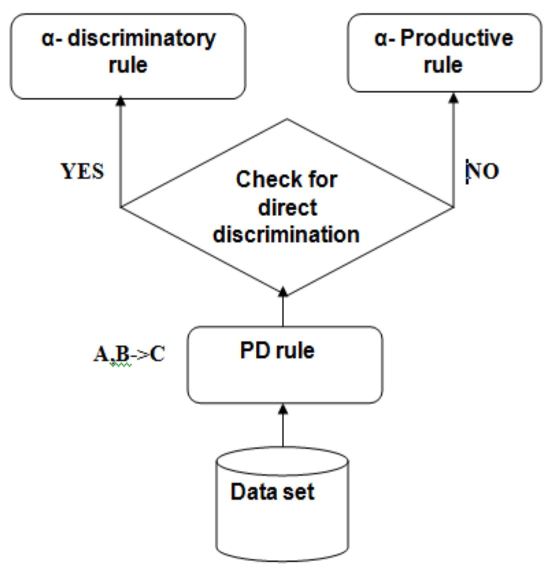

Fig:3.1 Direct discrimination measure

Fig 3.1 says that PD rule of the form A,B->C A is a discriminatory item set and check for direct discrimination such a measure is called Elift .if Elift>α α-discriminatory rule otherwise α-protective rule

3.3Potentially non discriminatory rule

A classification rule X -> C is potentially nondiscriminatory (PND) when X = D, B is a nondiscriminatory item set. For example,

{Zip =10451, City = NYC -> Hire = No} or

{Experience = Low, City = NYC -> Hire = No}
PND rule could lead to discriminatory decisions in combination with some background knowledge.
e.g., if the premise of the PND rule contains the zip code as an attribute and one knows that zip code 10451 is mostly inhabited by foreign people. Hence, measures are needed to quantify the indirect discrimination potential as well

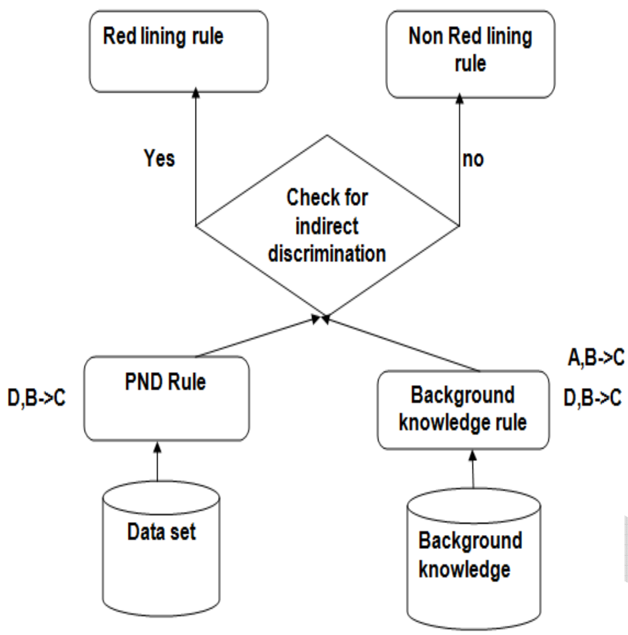

Fig:3.2Indirect discrimination measure

Fig: 3.2 says that PND rule of the form D,B->C where D is not directly discriminated but highly correlated with discriminatory attribute A . and check for indirect discrimination such a measure is called elb. If elb>α is taken as readlining rule otherwise is taken as non lining rule.

3.5Indirect discrimination measure

The purpose of indirect discrimination discovery is to identify redlining rules. In fact, redlining rules indicate biased rules that are indirectly inferred from nondiscriminatory items (e.g., Zip = 10451) because of their correlation with discriminatory ones. To determine the redlining rules, Pedreschi et al. in [12] stated the theorem below which gives a lower bound for α discrimination of PD classification rules, given information available in PND rules (γ , δ), and information available from background rules (β1, β2). They assume that background knowledge takes the form of classification rules relating a nondiscriminatory item set D to a discriminatory item set A within the context B.

Theorem Let r: D,B -> C be a PND classification rule, and let
γ= conf(r: D, B -> C) δ= conf(B -> C) > 0:
Let A be a discriminatory item set, and letβ1, β2 such that
Conf (rb1: A, B -> D) ≥β1
Conf (rb2: D, B -> A) ≥β2

$$\mathbf{F\ (x)}=\frac{\beta 1}{\beta 2}(\beta 2 + x - 1)$$

$$\mathbf{elb\ (x, y)}=\begin{cases} \frac{f(x)}{y} \text{ if } f(x) > 0 \\ 0 \text{ otherwise} \end{cases}$$

It holds that, for α ≥ 0, if elb (γ, δ) ≥ α, the PD classification rule r’: A,B -> C is α-discriminatory. Based on the above theorem, the following formal definitions of redlining and non redlining rules are presented:
A PND classification rule r : D,B -> C is a redlining rule if it could yield an α-discriminatory rule r' : A,B -> C in combination with currently available background knowledge rules of the form rb1 : A,B -> D and rb2 : D,B -> A, where A is a discriminatory item set. For example,
{Zip= 10451, City = NYC}-> Hire = No}.

A PND classification rule r : D,B -> C is a non redlining or legitimate rule if it cannot yield any α- discriminatory rule r': A,B -> C in combination with currently available background knowledge rules of the form
rb1 : A,B -> D and rb2 : D,B -> A, where A is a discriminatory item set. For example,
{Experience = Low, City = NYC} -> Hire= No}.

IV EXPERIMENTAL RESULTS AND ANALYSIS

4.1 Data sets

Two data sets are considered: adult and German credit data set.

4.1.1 Adult data set: We used the Adult data set [10], also known as Census Income, in our experiments. This data set consists of 48,842 records, split into a "train" part with 32,561 records and a "test" part with 16,281 records. The data set has 14 attributes (without class attribute). We used the "train" part in our experiments. The prediction task associated with the Adult data set is to determine whether a person makes more than 50K$ a year based on census and demographic information about people. The data set contains both categorical and numerical attributes. For our experiments with the Adult data set, we set DIs= {Sex = Female, Age = Young}. Although the Age attribute in the Adult data set is numerical, we converted it to categorical by partitioning its domain into two fixed intervals: Age <= 30 is renamed as Young and Age > 30 is renamed as old.

4.1.2 German credit data set: we also used the German Credit data set [11]. This data set consists of 1,000 records and 20 attributes (without class attribute) of bank account holders. This is a well-known real-life data set, containing both numerical and categorical attributes. It has been frequently used in the antidiscrimination literature [12], [7]. The class attribute in the German Credit data set takes values representing good or bad classification of the bank account holders. For our experiments with this data set, we set DIs = {Foreign worker = Yes, Personal Status =Female and not Single, Age = Old}; (cut-off for Age = Old: 50 years old).

4.2 Experimental result for adult data set

```
Rule : martial_status='Never-married' and gender='Male' => salary='<=50K'

A : martial_status='Never-married'
B : gender='Male'
C : salary='<=50K'

Number of tuples which satisfy A , B and C : 5591
Number of tuples which satisfy A and B : 5916
Number of tuples which satisfy B and C : 15128
Number of tuples which satisfy B : 21790

Confidence of A , B -> C : 0.9450642325895876
Confidence of B -> C : 0.6942634235888022
Elift : 1.3612473313145896
Rule : relationship='Own-child' and race='White' => salary='<=50K'

A : relationship='Own-child'
B : race='White'
C : salary='<=50K'

Number of tuples which satisfy A , B and C : 4196
Number of tuples which satisfy A and B : 4255
Number of tuples which satisfy B and C : 20699
Number of tuples which satisfy B : 27816

Confidence of A , B -> C : 0.9861339600470035
Confidence of B -> C : 0.7441400632729365
Elift : 1.3251993928531547
Rule : age='old' and captial_loss='Zero' => salary='<=50K'

A : age='old'
B : captial_loss='Zero'
C : salary='<=50K'

Number of tuples which satisfy A , B and C : 9652
Number of tuples which satisfy A and B : 10275
Number of tuples which satisfy B and C : 23974
Number of tuples which satisfy B : 31041

Confidence of A , B -> C : 0.939367396593674
Confidence of B -> C : 0.7723333655487903
Elift : 1.216271934498383
```

Fig4.1 Direct discrimination measure for adult data set

The α- discriminatory rule identification has been done using the following series of steps:

1. For the given data set, association rules have been generated.

2. From the set of rules, PD rules have been extracted where each PD rule contains at least one discriminatory attribute.

3. For the given data set, elift measure for all the PD rules has been calculated.

if elift value is > α,

α discriminatory rule

else

α-protective rule Table 4.1 gives the number of discriminatory rule for various value of threshold

Table 4.1 Result of Adult data set

Threshold	No of discriminatory rule
0.0	392
0.4	392
0.8	352
0.9	346
1.0	222
1.1	157
1.2	63
1.3	35
1.4	16
1.5	6
1.7	6
1.8	0

From the above table it is seen that as the threshold value is increased, number of discriminatory rule decreases. The maximum threshold value is 1.8 because it gives 0 discriminatory rule. The maximum number of discriminatory rule is 392 because it taken at threshold value 0. Taking an approximately intermediate value for both threshold and number of discriminatory rule the α value has been chosen to be 1and also analyzing set of discriminatory rule for both the previous threshold value 0.9 and next threshold value 1.1seems to be a suitable choice for α

4.3 Experimental result for German credit data set

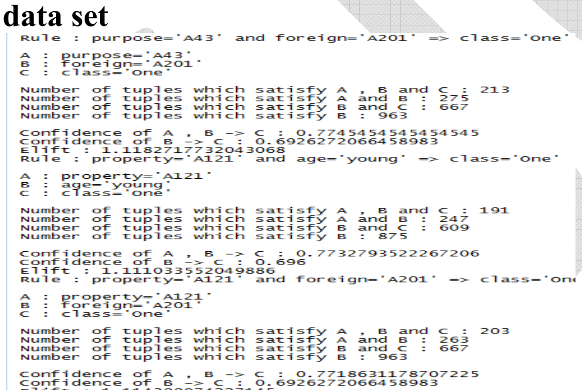

```
Rule : purpose='A43' and foreign='A201' => class='One'

A : purpose='A43'
B : foreign='A201'
C : class='One'

Number of tuples which satisfy A , B and C : 213
Number of tuples which satisfy A and B : 275
Number of tuples which satisfy B and C : 667
Number of tuples which satisfy B : 963

Confidence of A , B -> C : 0.7745454545454545
Confidence of B -> C : 0.6926272066458983
Elift : 1.1182717732043068
Rule : property='A121' and age='young' => class='One'

A : property='A121'
B : age='young'
C : class='One'

Number of tuples which satisfy A , B and C : 191
Number of tuples which satisfy A and B : 247
Number of tuples which satisfy B and C : 609
Number of tuples which satisfy B : 875

Confidence of A , B -> C : 0.7732793522267206
Confidence of B -> C : 0.696
Elift : 1.1110335520498866
Rule : property='A121' and foreign='A201' => class='One

A : property='A121'
B : foreign='A201'
C : class='One'

Number of tuples which satisfy A , B and C : 203
Number of tuples which satisfy A and B : 263
Number of tuples which satisfy B and C : 667
Number of tuples which satisfy B : 963

Confidence of A , B -> C : 0.7718631178707225
Confidence of B -> C : 0.6926272066458983
Elift : 1.114399074227145
```

Fig4.2 direct discrimination measure for German credit data set

Table 4.2: Result of German credit data set

Threshold	No of discriminatory rule
0.0	68
0.4	68
0.8	68
0.9	66
1.0	37
1.1	10
1.2	3
1.3	0

From the above table it is seen that as the threshold value is increased, number of discriminatory rule decreases. The maximum threshold value is 1.3 because it gives 0 discriminatory rule. The maximum number of discriminatory rule is 68 because it taken at threshold value 0. Taking an approximately intermediate value for both threshold and number of discriminatory rule the α value has been chosen to be 1and also analyzing set of discriminatory rule for both the previous threshold value 0.9 and next threshold value 1.1 seems to be a suitable choice for α

4.4 Indirect discrimination measure for adult data set

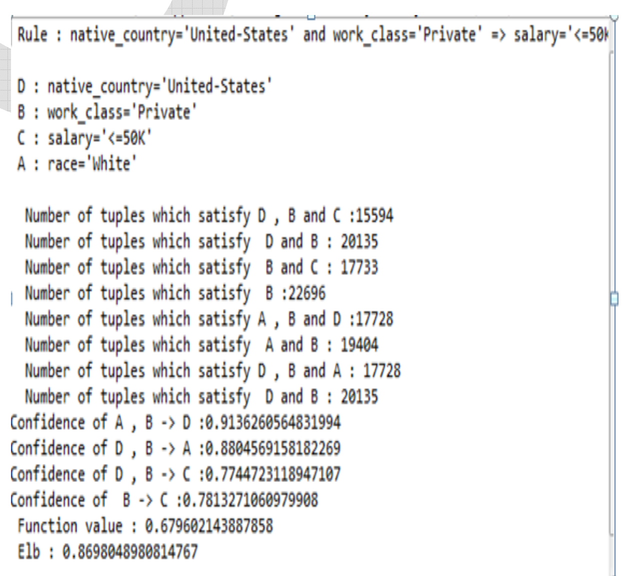

```
Rule : native_country='United-States' and work_class='Private' => salary='<=50K

D : native_country='United-States'
B : work_class='Private'
C : salary='<=50K'
A : race='White'

Number of tuples which satisfy D , B and C :15594
Number of tuples which satisfy  D and B : 20135
Number of tuples which satisfy  B and C : 17733
Number of tuples which satisfy  B :22696
Number of tuples which satisfy A , B and D :17728
Number of tuples which satisfy  A and B : 19404
Number of tuples which satisfy D , B and A : 17728
Number of tuples which satisfy  D and B : 20135
Confidence of A , B -> D :0.9136260564831994
Confidence of D , B -> A :0.8804569158182269
Confidence of D , B -> C :0.7744723118947107
Confidence of  B -> C :0.7813271060979908
Function value : 0.679602143887858
Elb : 0.8698048980814767
```

Fig4.3 Indirect Discrimination measure for adult data set

Fig 4.3 shows that indirect discrimination measure for Adult data set

4.5 Indirect discrimination measure for German credit data set

```
Rule : exp='A73' and no_of_credit='One' => class='One'

D : exp='A73'
B : no_of_credit='One'
C : class='One'
A : age='young'

Number of tuples which satisfy D , B and C : 154
Number of tuples which satisfy D and B : 226
Number of tuples which satisfy B and C : 433
Number of tuples which satisfy B : 633
Number of tuples which satisfy A , B and D : 209
Number of tuples which satisfy A and B : 558
Number of tuples which satisfy D , B and A : 209
Number of tuples which satisfy D , B: 226

Confidence of A , B -> D : 0.37455197132616486
Confidence of D , B -> A : 0.9247787610619469
Confidence of D , B -> C : 0.6814159292035398
Confidence of B -> C : 0.684044233807267
Function value : 0.24551971326164873
Elb : 0.3589237378628722
```

Fig 4.4 shows that indirect discrimination measure for German credit data set

VCONCLUSION

The purpose of this paper is to measure direct and indirect discrimination and identify categories and groups of individuals that have been directly discriminatory in the decision-making processes. The choice of the acceptable level of discrimination has been made. By analyzing the measures of each classification rule, direct discriminatory decision rules have been identified in order to convert them into legitimate (nondiscriminatory) classification rules. The experimental results reported demonstrate that the proposed techniques are quite successful.

REFERENCES

[1] S. Hajian and J. Domingo-Ferrer. A methodology for direct and indirect discrimination prevention in data mining. Manuscript, 2012.

R. Agrawal and R. Srikant, “Fast Algorithms for MiningAssociation Rules in Large Databases,” Proc. 20th Int’l Conf. Very Large Data Bases, pp. 487-499, 1994.

[2] T. Calders and S. Verwer, “Three Naive Bayes Approaches for Discrimination-Free Classification,” Data Mining and Knowledge Discovery, vol. 21, no. 2, pp. 277-292, 2010.

[3] European Commission, “EU Directive 2004/113/EC on Anti- Discrimination,” http://eur-lex.europa.eu/LexUriServ/ LexUriServ.do?uri=OJ:L:2004:373:0037:0043:EN:PDF, 2004.

[4] European Commission, “EU Directive 2006/54/EC on Anti- Discrimination,” http://eur-lex.europa.eu/LexUriServ/ LexUriServ.do?uri=OJ:L:2006:204:0023:0036:en:PDF, 2006.

[5] S. Hajian, J. Domingo-Ferrer, and A. Martı´nez-Balleste´, “Discrimination Prevention in Data Mining for Intrusion and Crime Detection,” Proc. IEEE Symp. Computational Intelligence in Cyber Security (CICS ’11), pp. 47-54, 2011.

[6] S. Hajian, J. Domingo-Ferrer, and A. Martı´nez-Balleste´, “Rule Protection for Indirect Discrimination Prevention in Data Mining,” Proc. Eighth Int’l Conf. Modeling Decisions for Artificial Intelligence (MDAI ’11), pp. 211-222, 2011.

[7] F. Kamiran and T. Calders, “Classification without Discrimination,” Proc. IEEE Second Int’l Conf. Computer, Control and Comm. (IC4 ’09), 2009.

[8] F. Kamiran and T. Calders, “Classification with no Discrimination by Preferential Sampling,” Proc. 19th Machine Learning Conf. Belgium and The Netherlands, 2010.

[9] F. Kamiran, T. Calders, and M. Pechenizkiy, “Discrimination Aware Decision Tree Learning,” Proc. IEEE Int’l Conf. Data Mining (ICDM ’10), pp. 869-874, 2010.

[10] R. Kohavi and B. Becker, “UCI Repository of Machine Learning Databases,” http://archive.ics.uci.edu/ml/datasets/Adult, 1996.

[11] D.J. Newman, S. Hettich, C.L. Blake, and C.J. Merz, “UCI Repository of Machine Learning Databases,” http://archive. ics.uci.edu/ml, 1998.

[12] D. Pedreschi, S. Ruggieri, and F. Turini, “Discrimination-Aware Data Mining,” Proc. 14th ACM Int’l Conf. Knowledge Discovery and Data Mining (KDD ’08), pp. 560-568, 2008.

[13] D. Pedreschi, S. Ruggieri, and F. Turini, “Measuring Discrimination in Socially-Sensitive Decision Records,” Proc. Ninth SIAM Data Mining Conf. (SDM ’09), pp. 581-592, 2009.

[14] D. Pedreschi, S. Ruggieri, and F. Turini, “Integrating Induction and Deduction for Finding Evidence of Discrimination,” Proc. 12th ACM Int’l Conf. Artificial Intelligence and Law (ICAIL ’09), pp. 157-166, 2009.

[15] S. Ruggieri, D. Pedreschi, and F. Turini, “Data Mining for Discrimination Discovery,” ACM Trans. Knowledge Discovery from Data, vol. 4, no. 2, article 9, 2010.

[16] S. Ruggieri, D. Pedreschi, and F. Turini, "DCUBE: Discrimination Discovery in Databases," Proc. ACM Int'l Conf. Management of Data (SIGMOD '10), pp. 1127-1130, 2010.
[17] P.N. Tan, M. Steinbach, and V. Kumar, Introduction to Data Mining. Addison-Wesley, 2006.
[18] United States Congress, US Equal Pay Act, http://archive. eeoc.gov/epa/anniversary/epa-40.html, 1963.
[19] V. Verykios and A. Gkoulalas-Divanis, "A Survey of Association Rule Hiding Methods for Privacy," Privacy-Preserving Data Mining: Models and Algorithms, C.C. Aggarwal and P.S. Yu, eds., Springer, 2008. Sara Hajian is

Priya.P received BE degree in Computer Science and Engineering fromJayaraj Anna packiam CSI College of engineering, Nazareth.She is currently doing ME in the Department of Computer Science and Engineering of Government College of Technology, Coimbatore.