Data Mining
Similarities and Distance Metrics

Instructor: Sharma Chakravarthy
sharmac@cse.uta.edu
The University of Texas at Arlington

Similarity and Dissimilarity

- Similarity is an extremely important concept in mining
  - Different from "exact" that is used in databases, search, etc.
- Humans do it without thinking much about it!
- Humans understand it qualitatively, computers cannot!
- So, how do we incorporate the notion of "similarity" into our computation?
- We need to map it quantitatively and map back the results!
- This means
  - Need to have a way to express this quantitatively!
  - The results depend upon how good this mapping is!
  - This is further complicated by multiple representations of input
    - Attribute types, binary and non-binary vectors, strings, matrix (sparse and otherwise), etc.

2

What is Similarity?

Similarity is hard to define, but… *"We know it when we see it"*

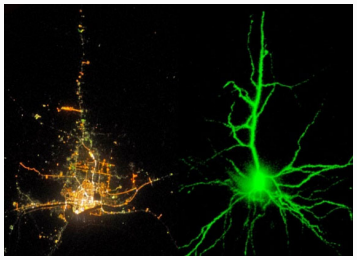

Slide based on one by Eamonn Keogh

What is Similarity?

Similarity is hard to define, but… *"We know it when we see it"*

Slide based on one by Eamonn Keogh

What is Similarity?

Slide based on one by Eamonn Keogh

Similarity is hard to define, but… "*We know it when we see it*"

Similarity and Dissimilarity

- So, how do we do it using computers?
  - Using the notion of "distance" between two objects
    - If they are "similar", the distance should be around 0 or 1
    - If they are "dissimilar", the distance should be opposite (1 or 0)
- Easily computable for a given data/attribute type
- Appropriate for a given data/attribute type
- So, we have several "distance" metrics
  - "city block" or "taxi", Euclidian, L-forms, edit distance, similarity matching, Jaccard, Pearson correlation, Spearman, to name a few
- This also means you have to choose the correct one (burden on the analyst)

6

Similarity and Dissimilarity

Similarity
- Numerical measure of how alike two data objects are
- Is higher (close to 1) when objects are more alike
- Often falls in the range [0,1] or [-1 to 1]

Dissimilarity
- Numerical measure of how different are two data objects
- Lower when objects are more alike
- Minimum dissimilarity is often 0
- Upper limit varies

Proximity refers to similarity or dissimilarity

7

Similarity and Dissimilarity

Similarity and Dissimilarity of Simple Attributes

Dissimilarity between Objects
- Distance
- Set Difference
- …

Similarity between Objects
- Binary Vectors
- Vectors
- …

8

Similarity/Dissimilarity for Simple Attributes

p and q are the attribute values for two data objects. d(p, q) and s(p, q) are dissimilarity and similarity between p and q.

Attribute Type	Dissimilarity	Similarity
Nominal	$d = \begin{cases} 0 & \text{if } p = q \\ 1 & \text{if } p \neq q \end{cases}$	$s = \begin{cases} 1 & \text{if } p = q \\ 0 & \text{if } p \neq q \end{cases}$
Ordinal	$d = \frac{\lvert p-q \rvert}{n-1}$ (values mapped to integers 0 to $n-1$, where n is the number of values)	$s = 1 - \frac{\lvert p-q \rvert}{n-1}$
Interval or Ratio	$d = \lvert p - q \rvert$	$s = -d$, $s = \frac{1}{1+d}$ or $s = 1 - \frac{d - min\_d}{max\_d - min\_d}$

Table 5.1. Similarity and dissimilarity for simple attributes

9

Euclidean Distance

Euclidean Distance

$$dist = \sqrt{\sum_{k=1}^{n} (p_k - q_k)^2}$$

Where n is the number of dimensions (attributes) and p_k and q_k are, respectively, the k[th] attributes (components) or data objects p and q.

Standardization is necessary, if scales differ.

10

Euclidean Distance

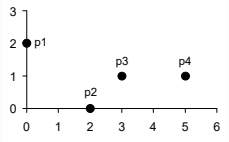

point	x	y
p1	0	2
p2	2	0
p3	3	1
p4	5	1

L2	p1	p2	p3	p4
p1	0	2.828	3.162	5.099
p2	2.828	0	1.414	3.162
p3	3.162	1.414	0	2
p4	5.099	3.162	2	0

L2 is Euclidian distance
Also know as $\|L_2\|$ norm

L1	p1	p2	p3	p4
p1	0	4	4	6
p2	4	0	2	4
p3	4	2	0	2
p4	6	4	2	0

L1 is Manhattan distance

Distance Matrix

11

Minkowski Distance

Minkowski Distance is a generalization of Euclidean Distance

$$dist = \left(\sum_{k=1}^{n} |p_k - q_k|^r\right)^{\frac{1}{r}}$$

$\|Lr\|$ norm in n-dimensions

Where r is a parameter, n is the number of dimensions (attributes) and p_k and q_k are, respectively, the kth attributes (components) or data objects p and q.

As r gets bigger, the contributions of smaller terms get less and less, until only the maximum term is left. For instance, $\|(3,4)\|_1=7$, $\|(3,4)\|_2=5$, $\|(3,4)\|_3 \sim 4.5$, and $\|(3,4)\|_\infty = 4$. So raising r "damps out" all the terms but the largest (also known as "chessboard distance")

12

Minkowski Distance: Examples

- $r = 1$. City block (Manhattan, taxicab, L_1 norm) distance.

  A common example of this is the Hamming distance, which is just the number of bits that are different between two binary vectors

- $r = 2$. Euclidean distance (L_2 norm)
- $r \to \infty$. "supremum" or "chessboard" (L_{max} norm, L_∞ norm) distance. (movement of the king)

  This is the maximum difference between any component of the vectors

- Do not confuse r with n, i.e., all these distances are defined for all numbers of dimensions.

13

Minkowski Distance

point	x	y
p1	0	2
p2	2	0
p3	3	1
p4	5	1

L1	p1	p2	p3	p4
p1	0	4	4	6
p2	4	0	2	4
p3	4	2	0	2
p4	6	4	2	0

L2	p1	p2	p3	p4
p1	0	2.828	3.162	5.099
p2	2.828	0	1.414	3.162
p3	3.162	1.414	0	2
p4	5.099	3.162	2	0

Distance Matrix

L_∞	p1	p2	p3	p4
p1	0	2	3	5
p2	2	0	1	3
p3	3	1	0	2
p4	5	3	2	0

14

Common Properties of a Distance

- Distances, such as the Euclidean distance, have some well known properties.

  1. $d(p, q) \geq 0$ for all p and q and $d(p, q) = 0$ only if $p = q$. (Positive definiteness)
  2. $d(p, q) = d(q, p)$ for all p and q. (Symmetry)
  3. $d(p, r) \leq d(p, q) + d(q, r)$ for all points p, q, and r. (Triangle Inequality) We know it is true for scalars. but it is also true for vectors!

  where $d(p, q)$ is the distance (similarity) between points (data objects), p and q.

- A distance that satisfies these properties is a metric

15

Slide based on one by Eamonn Keogh

Intuitions behind desirable distance measure properties

$D(A,B) = D(B,A)$
Otherwise you could claim "Alex looks like Bob, but Bob looks nothing like Alex."

$D(A,A) = 0$
Otherwise you could claim "Alex does not look like Alex."

Slide based on one by Eamonn Keogh

Intuitions behind desirable distance measure properties (continued)

$D(A,B) = 0$ Iff A=B

Otherwise there are objects in your world that are different, but you cannot tell apart.

$D(A,B) \leq D(A,C) + D(B,C)$

Otherwise you could claim "Alex is very like Carl and Bob is very like Carl, but Alex is very unlike Bob."

Common Properties of a Similarity

- Similarities, also have some well known properties.

1. $s(p, q) = 1$ (or maximum similarity) only if $p = q$.
2. $s(p, q) = s(q, p)$ for all p and q. (Symmetry)

where $s(p, q)$ is the similarity between points (data objects), p and q.

18

Similarity Between Binary Vectors

- Common situation is that objects, p and q, have only binary attributes
- Compute similarities using the following quantities

M_{01} = the number of attributes where p was 0 and q was 1
M_{10} = the number of attributes where p was 1 and q was 0
M_{00} = the number of attributes where p was 0 and q was 0
M_{11} = the number of attributes where p was 1 and q was 1

- Simple Matching and Jaccard Coefficients (SMC and JC)

SMC = number of matches / number of attributes
$= (M_{11} + M_{00}) / (M_{01} + M_{10} + M_{11} + M_{00})$

J = number of 11 matches / number of not-both-zero attributes values
$= (M_{11}) / (M_{01} + M_{10} + M_{11})$

can also be calculated for sets using intersection and union

19

SMC versus Jaccard: Example

$p = 1\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0$
$q = 0\ 0\ 0\ 0\ 0\ 0\ 1\ 0\ 0\ 1$

$M_{01} = 2$ (the number of attributes where p was 0 and q was 1)
$M_{10} = 1$ (the number of attributes where p was 1 and q was 0)
$M_{00} = 7$ (the number of attributes where p was 0 and q was 0)
$M_{11} = 0$ (the number of attributes where p was 1 and q was 1)

$SMC = (M_{11} + M_{00})/(M_{01} + M_{10} + M_{11} + M_{00}) = (0+7) / (2+1+0+7) = 0.7$

$J = (M_{11}) / (M_{01} + M_{10} + M_{11}) = 0 / (2 + 1 + 0) = 0$

20

SMC versus Jaccard: which one to use?

p = 1 0 0 0 0 0 0 0 0 0 0 *SMC is 0.7 and J is 0*

q = 0 0 0 0 0 0 0 1 0 0 1

- Suppose we are looking at Amazon customers who bought CDs
- If each user's purchase of CD is a vector and we want to know who is similar to who in purchases,
  - Should we use SMC or Jaccard? And why?
- Jaccard will look for the matches!
- Whereas, SMC will include all CDs sold by Amazon! If 0-0 matches are counted, most of the customers will be similar to others!
- Which one will give a better value for similarity?
- Jaccard!
- This understanding and use the correct metric is is key to mining!

Cosine Similarity

- Why is it called Cosine similarity?
  If d_1 and d_2 are two (e.g., document) vectors, then
  $\cos(d_1, d_2) = (d_1 \bullet d_2) / ||d_1|| ||d_2||$,
  where $\bullet$ indicates vector dot product (or scalar product) and $||d||$ is the length (or magnitude) of vector d *(Euclidean distance)*

Dot product is different From cross product! Denoted by a x b, where a and b are vectors

- Note that dot product is a scalar! But is computed on a vector!
- Also, note that dot product is different from the cross product (or vector product) of vectors

- The dot product can be used to compute the length of a vector or the angle between two vectors
- The cross product is used to find a vector which is perpendicular to the plane spanned by two vectors (binary operation)

22

Cosine Similarity

Example:

d_1 = **3 2 0 5 0 0 0 2 0 0**

d_2 = **1 0 0 0 0 0 0 1 0 2**

Dot product is different From cross product! Denoted by a x b, where a and b are vectors

$d_1 \bullet d_2$ = 3*1 + 2*0 + 0*0 + 5*0 + 0*0 + 0*0 + 0*0 + 0*1 + 0*0 + 0*2 = 5

$||d_1||$ = $(3*3+2*2+0*0+5*5+0*0+0*0+0*0+2*2+0*0+0*0)^{0.5}$ = $(42)^{0.5}$ = 6.481

$||d_2||$ = $(1*1+0*0+0*0+0*0+0*0+0*0+0*0+1*1+0*0+2*2)^{0.5}$ = $(6)^{0.5}$ = 2.245

$\cos(d_1, d_2)$ = .3150 (closer to 0, here means they are not similar)

- Is there a difference between dot product as a similarity measure and cosine as a similarity measure?

  Yes. cosine only cares about angle difference where as dot product cares about angle and magnitude

  If you normalize your data to have the same magnitude, the two are indistinguishable

  Technically, both are not "distance metric"

23

Cosine Similarity

Let us understand it in a bit more detail! Consider vectors *a, b, c*

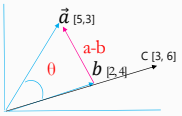

Is *a* similar to *b*? Is *b* similar to *c*? or is *a* similar to *c*?

What if you ignore the magnitude and look at how close they are!

What is the length of vector a? note that length is a scalar

We can compute $\|a\|^2 = 5^2 + 3^2$ using Euclidian distance formula (Pythagoras theorem!)

We also know that cos(θ) is adjacent/hypotenuse

24

Cosine Similarity

We also know that

cos(0) is 1 and cos(90) is 0, and cos (180) is -1

The general formula for non-right-angled triangles is

$\| a\text{-}b \|^2 = \| a \|^2 + \| b \|^2 - 2 \| a \| \| b \| \cos(\theta)$

When this is simplified for Cos(θ), you get $(a \bullet b) / ||a|| ||b||$ (try simplifying!)

If the vectors are orthogonal, the value is 0 showing they are perpendicular (more separation)

If the vectors overlap (have 0 degree between them) the value is 1, showing complete similarity!

If the value is -1, they are opposite vectors

- Easy to generalize to n-dimensions
- Very widely used similarity metric!

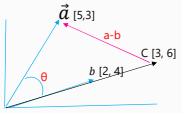

Cosine similarity value examples

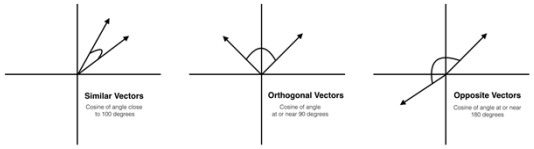

Cosine Similarity: why we ignore magnitude?

For example, vectors b and c below have 0 angle between them, but different magnitudes

So, they are similar!

However, vectors a and b look similar in magnitude, but have a large angle between them

They are considered dissimilar! Why?

Think of a, b, and c as document vectors

	X	Y
Doc a	5	3
Doc b	2	4
Doc c	3	6

HW: what if the vectors are of different length? How do you compute cos(theta)?

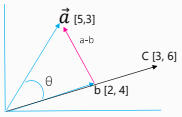

Pearson Correlation Coefficient

- Correlation measures the linear relationship between objects
- Pearson correlation is used between two sets of numerical values!

population correlation

$$\rho_{X,Y} = \operatorname{corr}(X,Y) = \frac{\operatorname{cov}(X,Y)}{\sigma_X \sigma_Y} = \frac{E[(X-\mu_X)(Y-\mu_Y)]}{\sigma_X \sigma_Y},$$

$$E[X] = \mu. \qquad \sigma = \sqrt{E[(X-\mu)^2]}$$
$$= \sqrt{E[X^2] + E[-2\mu X] + E[\mu^2]} = \sqrt{E[X^2] - 2\mu E[X] + \mu^2}$$
$$= \sqrt{E[X^2] - 2\mu^2 + \mu^2} = \sqrt{E[X^2] - \mu^2}$$
$$= \sqrt{E[X^2] - (E[X])^2}$$

covariance

$$\operatorname{cov}(X,Y) = E[(X - E[X])(Y - E[Y])]$$
$$= E[XY - XE[Y] - E[X]Y + E[X]E[Y]]$$
$$= E[XY] - E[X]E[Y] - E[X]E[Y] + E[X]E[Y]$$
$$= E[XY] - E[X]E[Y].$$

sample correlation

$$r = \frac{\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{n}(x_i - \bar{x})^2}\sqrt{\sum_{i=1}^{n}(y_i - \bar{y})^2}}$$

$$r = r_{xy} = \frac{n\sum x_i y_i - \sum x_i \sum y_i}{\sqrt{n\sum x_i^2 - (\sum x_i)^2}\sqrt{n\sum y_i^2 - (\sum y_i)^2}}.$$

$$r = r_{xy} = \frac{\sum x_i y_i - n\bar{x}\bar{y}}{\sqrt{(\sum x_i^2 - n\bar{x}^2)}\sqrt{(\sum y_i^2 - n\bar{y}^2)}}.$$

$$r = r_{xy} = \frac{1}{n-1}\sum_{i=1}^{n}\left(\frac{x_i - \bar{x}}{s_x}\right)\left(\frac{y_i - \bar{y}}{s_y}\right)$$ ← dot product

$$r = r_{xy} = \frac{\sum x_i y_i - n\bar{x}\bar{y}}{(n-1)s_x s_y} \qquad s_x = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(x_i - \bar{x})^2}$$

Mean, median, and Mode, Standard Deviation

- The *sample mean* is the average and is computed as the sum of all the observed outcomes from the sample divided by the total number of events. We use x as the symbol for the sample mean. In math terms, $\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x$
- The *mode* of a set of data is the number with the highest frequency (voting, consensus)
- Mean is strongly affected by outliers.

Mean of 10, 20, 30, 30, 40 is 65
Mean of 10, 20, 30, 30, 40, 170 is 150
Median of above is 35
Mode of the above is 30

- The *median* is the middle score. If we have an even number of events we take the average of the two middles. The median is better for describing the typical value. It is often used for income and home prices.
- The above indicate where the center of the data set is.

29

When to use which one?

Name & Meaning	Formula / Example	Used for
Arithmetic Mean [average]	$\frac{sum}{size} = \frac{a+b+c}{3}$	Most situations ("average item")
Median [middle value]	Middle of sorted list (2 middles? Average 'em)	Wildly varying samples (houses, incomes)
Mode [most popular]	Most popular value	No compromises (winner takes all)
Geometric Mean [average factor]	$\sqrt[3]{abc}$	Investments, growth, area, volume
Harmonic Mean [average rate]	$\frac{3}{\frac{1}{a}+\frac{1}{b}+\frac{1}{c}}$	Speed, production, cost

30

Variance and standard deviation

- If you want the data spread, *variance* and *standard deviation* can be used. Variance measures the average degree to which each point differs from the mean (variability) We define the variance to be $s^2 = \frac{1}{n-1}\sum_{i=1}^{n}(x-\bar{x})^2$ (why n-1 and not n?)
- The units of variance is not the same as that of data
- A variance of 0 means all data points are identical!
- A high variance indicates that data points are very spread out from the mean and from one another.
- A distinction is made between sample mean and population mean represented as mu (μ) (n-1) is used fro sample mean and n is used for population mean;
- There is also the notion of average deviation or Mean absolute deviation
  - Uses absolute values instead of squaring to circumvent negative differences
  - Used less frequently because the use of absolute values makes further calculations more complicated

31

Variance and standard deviation

- Standard deviation is the square root of variance. $s^2 = \frac{1}{n-1}\sum_{i=1}^{n}(x-\bar{x})^2$ (σ)
- Units of standard deviation is the same as units of data.
- Standard deviation is a "measure of spread". Indicates how spread is your data about the center.
- Bounds: on how much of the data must lie close to the mean. In particular for any k > 1, the proportion of the data that lies within k standard deviations of the mean is at least $1 - 1/k^2$ (depends on the distribution)
- Chebyshev's theorem

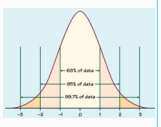

32

Use of standard deviation

- Use: One of the most common methods of determining the risk an investment poses is standard deviation.
- Standard deviation helps determine market volatility or the spread of asset prices from their average price.
- When prices move wildly, standard deviation is high, meaning an investment will be risky. Low standard deviation means prices are calm, so investments come with low risk.
- Average spending in a restaurant can be better estimated by using standard deviation rather than the mean

33

Use of standard deviation, percentile

- The most common definition of a percentile is a number where **a certain percentage of scores fall below that number.**
- Percentile: Let p be any integer between 0 and 100. The *p*th *percentile* of data set is the data value at which *p* percent of the value in the data set are less than or equal to this value
  - Remember GRE, SAT, LSAT
- You might know that you scored 67 out of 90 on a test. But that figure may not be helpful unless you know what percentile you fall into. If you know that your score is in the **90th percentile**, that means you scored better than 90% of people who took the test.

34

Covariance

- *Covariance* is a statistical tool that is used to determine the relationship between the movement of two data sets or random variables (e.g., stock prices)
- Covariance measures the directional relationship between the two data sets. A positive covariance means that asset returns move together while a negative covariance means they move inversely.
  - Used for **diversifying** portfolios in an investment!
  - For choosing a portfolio where all stocks that do not behave the same way in prices/returns (not putting all eggs in one basket!)
- **$Cov(x,y) = SUM [(x_i - x_m) * (y_i - y_m)] / (n - 1)$**
- While the covariance does measure the directional relationship between two assets, it does not show the strength of the relationship between the two assets; the coefficient of correlation is a more appropriate indicator of this strength.
- Has units. Does not measure the strength or the dependency between variables!
- Strength is measured using a coefficient; dependency by regression!

35

Covariance and correlation

- On the other hand, correlation measures the strength of the relationship between variables. Correlation is the scaled measure of covariance. It is dimensionless. In other words, the correlation coefficient is always a pure value and not measured in any units.
- The relationship between two concepts can be expressed using the formula
- $\rho(X,Y) = \frac{Cov(X,Y)}{\sigma_X \sigma_Y}$ (Pearson Correlation coefficient)
- Where:
- r(X,Y) – the correlation between the variables X and Y
- Cov(X,Y) – the covariance between the variables X and Y
- σX – the standard deviation of the X-variable
- σY – the standard deviation of the Y-variable

36

Example

$$\rho(X,Y) = \frac{\mathrm{Cov}(X,Y)}{\sigma_X \sigma_Y} = r_{xy} = \frac{\sum(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum(x_i - \bar{x})^2 \sum(y_i - \bar{y})^2}}$$

(a = $x_i - \bar{x}$, b = $y_i - \bar{y}$)

	S&P 500	Apple
2013	1691.75	68.96
2014	1977.80	100.11
2015	1884.09	109.06
2016	2151.13	112.18
2017	2519.36	154.12
Mean	2044.83	108.89

	S&P 500	Apple	Step 3: a	Step 4: b	Step 5: a x b	a^2	b^2
2013	1691.75	68.96	-353.08	-39.93	14,096.91	124,662.66	1,594.09
2014	1977.80	100.11	-67.03	-8.78	588.22	4,492.48	77.02
2015	1884.09	109.06	-160.74	0.17	-27.97	25,836.07	0.03
2016	2151.13	112.18	106.30	3.29	350.16	11,300.52	10.85
2017	2519.36	154.12	474.53	45.23	21,465.08	225,182.62	2,046.11
Mean	2044.83	108.89	Sums		36,472.40	391,474.35	3,728.10

$$r_{xy} = \frac{\text{numerator}}{\sqrt{391,474.35 \times 3,728.10}} = \frac{36,272.40}{\sqrt{391,474.35 \times 3,728.10}} = 0.95$$

Positive, close to 1

37

Visually Evaluating Correlation

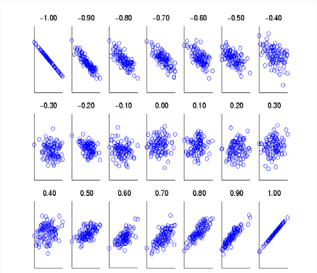

Scatter plots showing the similarity from –1 to 1.

38

Correlation coefficient

- It is a value between -1 and 1
- Both -1 or 1 indicate very strong (or perfect) positive or negative correlation
- Values between 0.3 and 0.7 (-0.3 and -0.7) indicate a moderate positive (negative) linear relationship
- The relationship between two variable is considered strong when their r value is large than 0.7 (-0.7)
- A r value of 0 (or near 0) indicates no linear relationship!
- When the r value is 0, non-linear relationships could still exist.
- The book has an example. Take a look at it.
- This is used when both variables being studied are normally distributed. This correlation is affected by extreme values
- Otherwise, use Spearman's rank correlation coefficient which is better if one or both are not normally distributed!
- Look up Spearman's coefficient on your own

39

Correlation coefficient

- What is the difference between uniform and normal distribution?

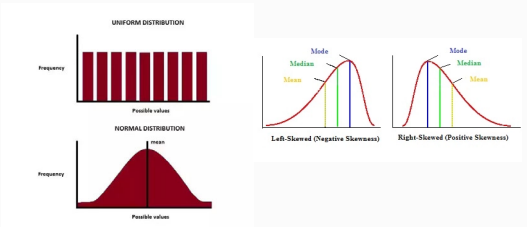

40

General Approach for Combining Similarities

- Sometimes attributes are of many different types, but an overall similarity is needed.

1. For the k^{th} attribute, compute a similarity, s_k, in the range $[0, 1]$.
2. Define an indicator variable, δ_k, for the k_{th} attribute as follows:

$$\delta_k = \begin{cases} 0 & \text{if the } k^{th} \text{ attribute is a binary asymmetric attribute and both objects have a value of 0, or if one of the objects has a missing values for the } k^{th} \text{ attribute} \\ 1 & \text{otherwise} \end{cases}$$

3. Compute the overall similarity between the two objects using the following formula:

$$similarity(p, q) = \frac{\sum_{k=1}^{n} \delta_k s_k}{\sum_{k=1}^{n} \delta_k}$$

41

Using Weights to Combine Similarities

- May not want to treat all attributes the same.
  - Use weights w_k which are between 0 and 1 and sum to 1.

$$similarity(p, q) = \frac{\sum_{k=1}^{n} w_k \delta_k s_k}{\sum_{k=1}^{n} \delta_k}$$

$$distance(p, q) = \left(\sum_{k=1}^{n} w_k |p_k - q_k|^r \right)^{1/r}$$

42

Selecting the right proximity/similarity measure

- For many types of dense and continuous data, metric distance measures (e.g., Euclidian distance) are used
- Scale issues need to be dealt with using normalization and weighting attributes
- For sparse data, which often consist of asymmetric attributes similarity measures that ignore 0-0 matches are used. Cosine, Jaccard are appropriate for such data
- Invariance to scaling (multiplication), translation (addition) makes cosine, Euclidian, and correlation useful for sparse data
- Correlation works better for time series where both scaling and translation are important
- In some cases, normalization and transformation of data may be needed (for periodicity data)
- In summary, proper choice of a proximity measure can be a time consuming task requiring both domain knowledge and the purpose for which the measure is being used!

43

Understand a few functions and their behavior

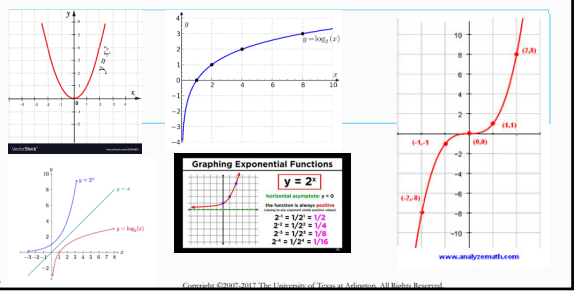

44