Carnegie Mellon

15-826: Multimedia Databases and Data Mining

Lecture #25: Time series mining and forecasting

Christos Faloutsos

1

Carnegie Mellon

Problem:

Q: mine/forecast (one, or more) time sequences

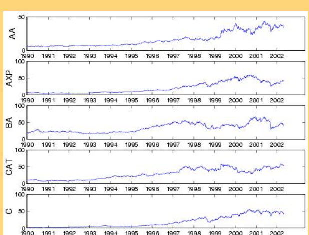

15-826 Copyright (c) 2019 C. Faloutsos 2

2

Carnegie Mellon

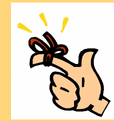

Answers

- Similarity search: **Euclidean**/time-warping; **feature extraction** and **SAMs**
- Linear Forecasting: **AR** (Box-Jenkins)
- Non-linear forecasting: **lag-plots**
- Gray-box modeling: **Lotka-Volterra**

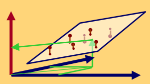

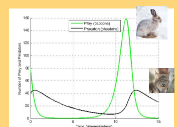

3

Carnegie Mellon

Must-Read Material

- Byong-Kee Yi, Nikolaos D. Sidiropoulos, Theodore Johnson, H.V. Jagadish, Christos Faloutsos and Alex Biliris, *Online Data Mining for Co-Evolving Time Sequences*, ICDE, Feb 2000.
- Chungmin Melvin Chen and Nick Roussopoulos, *Adaptive Selectivity Estimation Using Query Feedbacks*, SIGMOD 1994

4

Carnegie Mellon

Thanks


Deepay Chakrabarti (UT-Austin)


Spiros Papadimitriou (Rutgers)


Prof. Byoung-Kee Yi (Samsung)

5

Carnegie Mellon

Outline

- Motivation
- Similarity search – distance functions
- Linear Forecasting
- Bursty traffic - fractals and multifractals
- Non-linear forecasting
- Gray box modeling – Lotka Volterra eq's
- Conclusions

6

Carnegie Mellon

Problem definition

- Given: one or more sequences
  x_1 , x_2 , … , x_t , …
  (y_1, y_2, … , y_t, …
  …)
- Find
  - similar sequences; forecasts
  - patterns; clusters; outliers

15-826 Copyright (c) 2019 C. Faloutsos 7

7

Carnegie Mellon

Motivation - Applications

- Financial, sales, economic series
- Medical
  - ECGs +; blood pressure etc monitoring
  - reactions to new drugs
  - elderly care

15-826 Copyright (c) 2019 C. Faloutsos 8

8

Carnegie Mellon

Motivation - Applications (cont' d)

- 'Smart house'
  - sensors monitor temperature, humidity, air quality
- video surveillance

9

Carnegie Mellon

Motivation - Applications (cont' d)

- civil/automobile infrastructure
  - bridge vibrations [Oppenheim+02]
  - road conditions / traffic monitoring

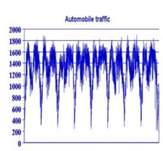

10

Carnegie Mellon

Motivation - Applications (cont' d)

- Weather, environment/anti-pollution
  - volcano monitoring
  - air/water pollutant monitoring

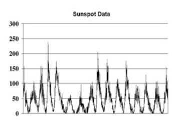

15-826 11

11

Carnegie Mellon

Motivation - Applications (cont' d)

- Computer systems
  - 'Active Disks' (buffering, prefetching)
  - web servers (ditto)
  - network traffic monitoring
  - ...

15-826 12

12

Carnegie Mellon

Stream Data: Disk accesses

#bytes

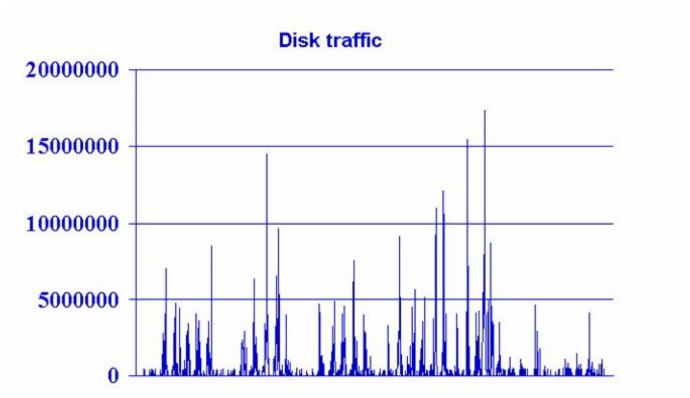

time

Carnegie Mellon

Problem #1:

Goal: given a signal (e.g.., #packets over time)

Find: patterns, periodicities, and/or compress

count

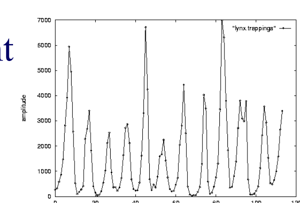

lynx caught per year
(packets per day;
temperature per day)

year

Carnegie Mellon

Problem#2: Forecast

Given x_t, x_{t-1}, …, forecast x_{t+1}

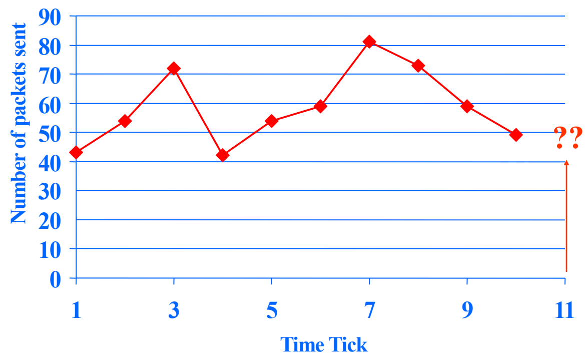

15

Carnegie Mellon

Problem#2’ : Similarity search

E.g.., Find a 3-tick pattern, similar to the last one

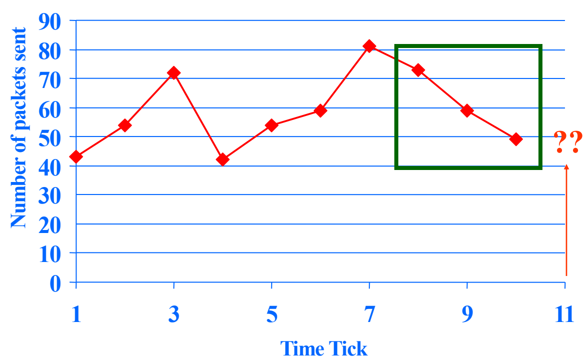

16

Carnegie Mellon

Problem #3:

- Given: A set of **correlated** time sequences
- Forecast 'Sent(t)'

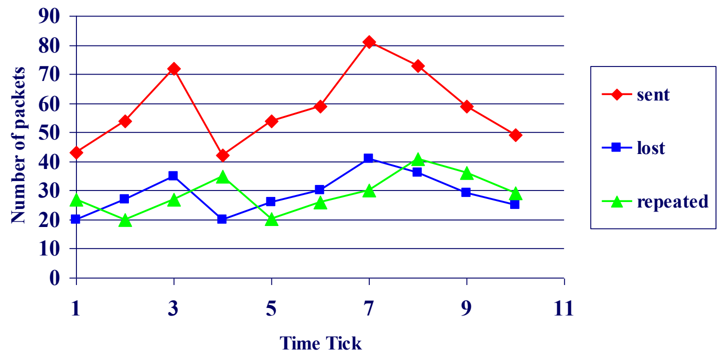

15-826 Copyright (c) 2019 C. Faloutsos 17

17

Carnegie Mellon

Important observations

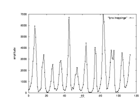

Patterns, rules, forecasting and similarity indexing are closely related:

- To do forecasting, we need
  - to find patterns/rules
  - **compress**
  - to find similar settings in the past
- to find outliers, we need to have forecasts
  - (outlier = too far away from our forecast)

15-826 Copyright (c) 2019 C. Faloutsos 18

18

Carnegie Mellon

Outline

- Motivation
- ➡ Similarity Search and Indexing
- Linear Forecasting
- Bursty traffic - fractals and multifractals
- Non-linear forecasting
- Gray box modeling – Lotka Volterra eq's
- Conclusions

Carnegie Mellon

Detailed Outline

- Motivation
- ➡ Similarity search and distance functions
  – Euclidean
  – Time-warping
- ...

Carnegie Mellon

Problem:

Q: How similar are two sequences?

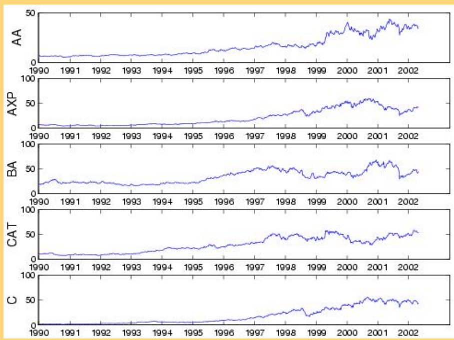

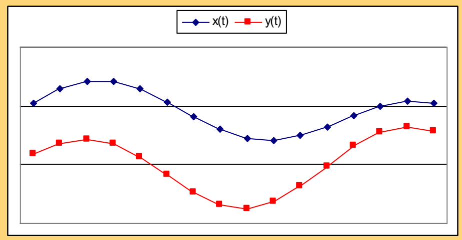

15-826 Copyright (c) 2019 C. Faloutsos 21

21

Carnegie Mellon

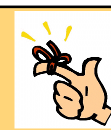

Answer:

Q: How similar are two sequences?

A: Euclidean distance (<-> cosine similarity)

$$D(\vec{x},\vec{y}) = \sum_{i=1}^{n} (x_i - y_i)^2$$

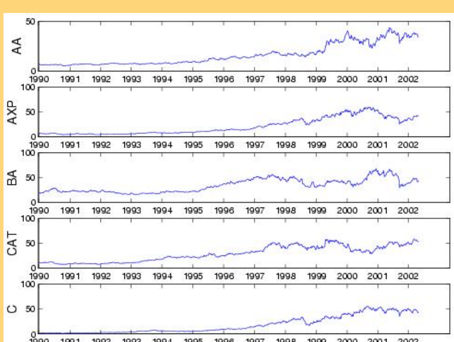

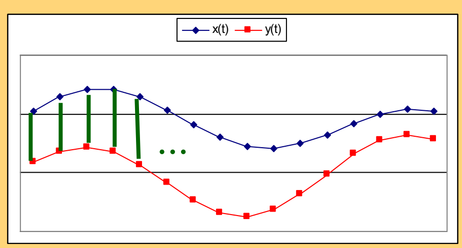

15-826 Copyright (c) 2019 C. Faloutsos 22

22

Carnegie Mellon

Importance of distance functions

Subtle, but **absolutely necessary**:

- A 'must' for similarity indexing (-> forecasting)
- A 'must' for clustering

Two major families

- – Euclidean and Lp norms
- – Time warping and variations

Carnegie Mellon

Euclidean and Lp

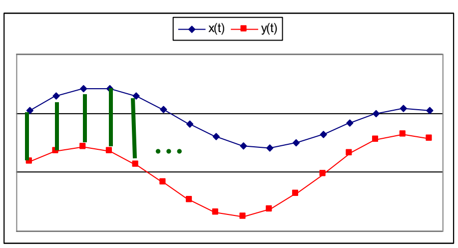

$$D(\vec{x},\vec{y}) = \sum_{i=1}^{n} (x_i - y_i)^2$$

$$L_p(\vec{x},\vec{y}) = \sum_{i=1}^{n} | x_i - y_i |^p$$

- L_1: city-block = Manhattan
- L_2 = Euclidean
- L_∞

Carnegie Mellon

Observation #1

- **Time sequence -> n-d vector**

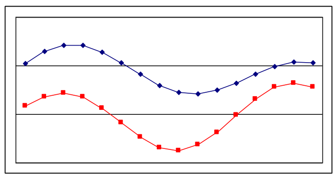

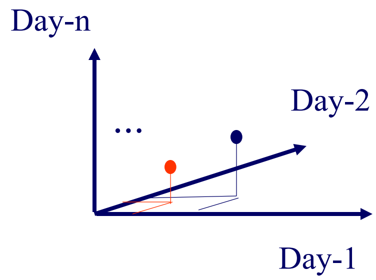

25

Carnegie Mellon

Observation #2

Euclidean distance is closely related to
- cosine similarity
- dot product
- ‘cross-correlation’ function

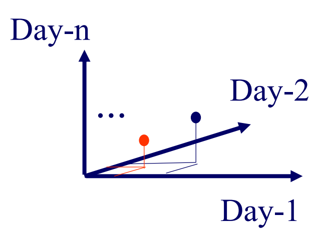

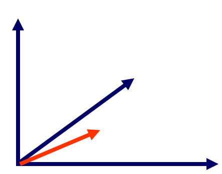

26

Carnegie Mellon

Time Warping

- allow accelerations - decelerations
  - (with or w/o penalty)
- THEN compute the (Euclidean) distance (+ penalty)
- related to the string-editing distance

27

Carnegie Mellon

Time Warping

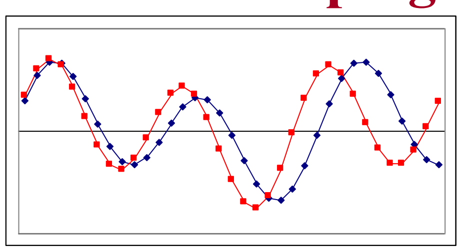

'stutters' :

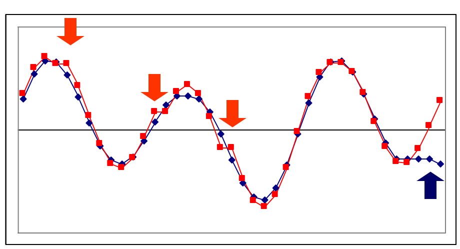

28

Carnegie Mellon

Time warping

Q: how to compute it?

A: dynamic programming

D(i, j) = cost to match

prefix of length *i* of first sequence *x* with prefix of length *j* of second sequence *y*

15-826 29

29

Carnegie Mellon

reminder

Full-text scanning

- Approximate matching - **string editing** distance:

d(‘survey’, ‘surgery’) = 2

= min # of insertions, deletions, substitutions to transform the first string into the second

SURVEY

SURGERY

15-826 30

30

Carnegie Mellon

Time warping

Thus, with no penalty for stutter, for sequences

$x_1, x_2, \ldots, x_{i,;} \qquad y_1, y_2, \ldots, y_j$

$$D(i,j) = \|x[i] - y[j]\| + \min \begin{cases} D(i-1, j-1) & \text{no stutter} \\ D(i, j-1) & \text{x-stutter} \\ D(i-1, j) & \text{y-stutter} \end{cases}$$

15-826 31

31

Carnegie Mellon

Time warping

VERY SIMILAR to the string-editing distance

$$D(i,j) = \|x[i] - y[j]\| + \min \begin{cases} D(i-1, j-1) & \text{no stutter} \\ D(i, j-1) & \text{x-stutter} \\ D(i-1, j) & \text{y-stutter} \end{cases}$$

15-826 32

32

Carnegie Mellon

reminder

Full-text scanning

```
if s[i] = t[j] then
    cost( i, j ) = cost(i-1, j-1)
else
    cost(i, j ) = min (
            1 + cost(i, j-1) // deletion
            1 + cost(i-1, j-1) // substitution
            1 + cost(i-1, j) // insertion
            )
```

15-826 Copyright (c) 2019 C. Faloutsos 33

33

Carnegie Mellon

Time warping

VERY SIMILAR to the string-editing distance

Time-warping

$$D(i,j) = \|x[i] - y[j]\| + \min \begin{cases} D(i-1, j-1) \\ D(i, j-1) \\ D(i-1, j) \end{cases}$$

String editing

cost(i,j) = min {
1 + cost(i-1, j-1) // sub.
1 + cost(i, j-1) // del.
1 + cost(i-1, j) // ins.
}

15-826 Copyright (c) 2019 C. Faloutsos 34

34

Carnegie Mellon

Time warping

- Complexity: O(M*N) - quadratic on the length of the strings
- **Many** variations (penalty for stutters; limit on the number/percentage of stutters; …)
- popular in voice processing [Rabiner + Juang]

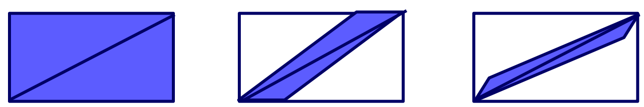

35

Carnegie Mellon

Other Distance functions

- piece-wise linear/flat approx.; compare pieces [Keogh+01] [Faloutsos+97]
- 'cepstrum' (for voice [Rabiner+Juang])
  - do DFT; take log of amplitude; do DFT again!
- Allow for small gaps [Agrawal+95]

See tutorial by [Gunopulos + Das, SIGMOD01]

36

Carnegie Mellon

Other Distance functions

- In [Keogh+, KDD' 04]: parameter-free, MDL based

37

Carnegie Mellon

Conclusions

Prevailing distances:

- Euclidean and
- time-warping

38

Carnegie Mellon

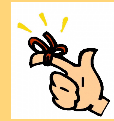

Answer:

Q: How similar are two sequences?

A: Euclidean distance (<-> cosine similarity)

$$D(\vec{x}, \vec{y}) = \sum_{i=1}^{n} (x_i - y_i)^2$$

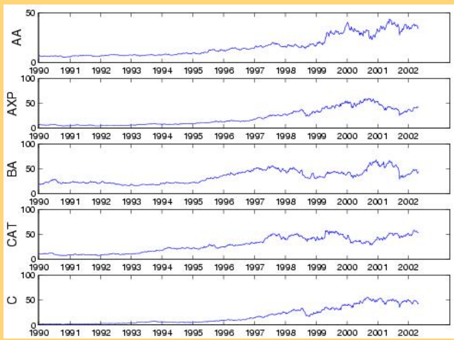

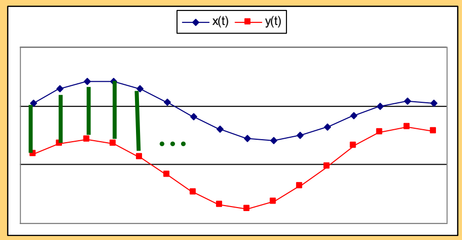

39

Carnegie Mellon

Outline

- Motivation
- Similarity search and distance functions
- Linear Forecasting
- Bursty traffic - fractals and multifractals
- Non-linear forecasting
- Gray box modeling – Lotka Volterra eq’s
- Conclusions

40

Carnegie Mellon

Linear Forecasting

41

Carnegie Mellon

Forecasting

"Prediction is very difficult, especially about the future." - Nils Bohr

http://www.hfac.uh.edu/MediaFutures/thoughts.html

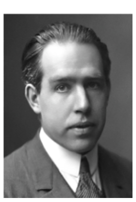

42

Carnegie Mellon

Problem#2: Forecast

- given x_{t-1}, x_{t-2}, …,
- Q: forecast x_t

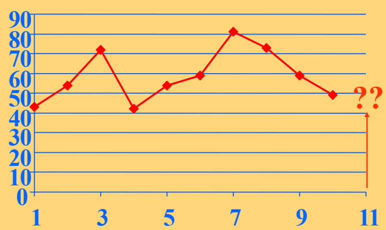

43

Carnegie Mellon

Solution: AR(IMA)

- given x_{t-1}, x_{t-2}, …,
- Q: forecast x_t
- A: AR(IMA) = Box-Jenkins (< Holt-Winters, Kalman)

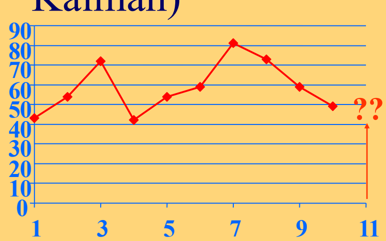

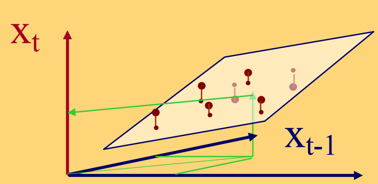

44

Carnegie Mellon

Detailed Outline

- Motivation
- ...
- Linear Forecasting
  - – Auto-regression: Least Squares; RLS
  - – Co-evolving time sequences
  - – Examples
  - – Conclusions

15-826 Copyright (c) 2019 C. Faloutsos 45

45

Carnegie Mellon

Reference

[Yi+00] Byoung-Kee Yi et al.: *Online Data Mining for Co-Evolving Time Sequences*, ICDE 2000. (Describes MUSCLES and Recursive Least Squares)

15-826 Copyright (c) 2019 C. Faloutsos 46

46

Problem#2: Forecast

- Example: give x_{t-1}, x_{t-2}, …, forecast x_t

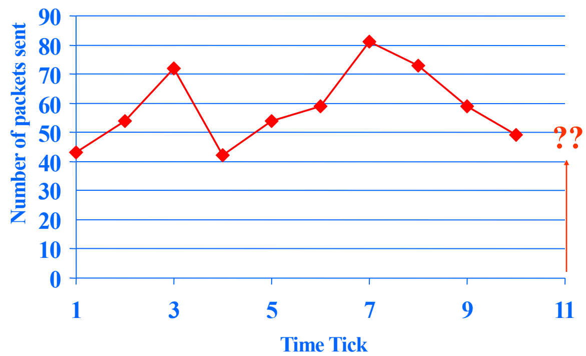

47

Forecasting: Preprocessing

MANUALLY:

remove trends

spot periodicities

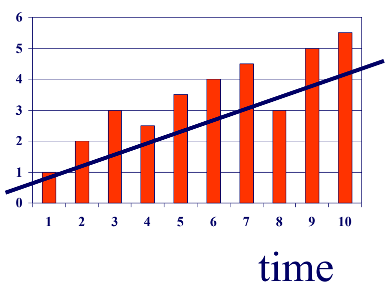

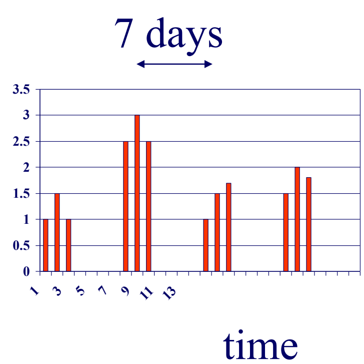

48

Carnegie Mellon

Problem#2: Forecast

- Solution: try to express

  x_t

  as a linear function of the past: x_{t-2}, x_{t-2}, …,

  (up to a window of w)

Formally:

$$x_t \approx a_1 x_{t-1} + \ldots + a_w x_{t-w} + noise$$

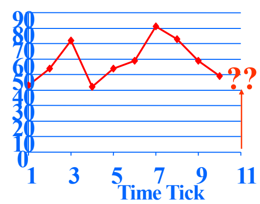

15-826 Copyright (c) 2019 C. Faloutsos 49

49

Carnegie Mellon

(Problem: Back-cast; interpolate)

- Solution - interpolate: try to express

  x_t

  as a linear function of the past AND the future:

  x_{t+1}, x_{t+2}, … $x_{t+wfuture;}$ x_{t-1}, … $x_{t-wpast}$

  (up to windows of w_{past} , w_{future})
- EXACTLY the same algo’s

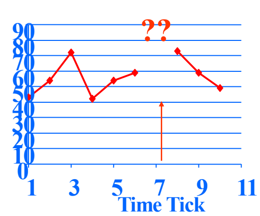

15-826 Copyright (c) 2019 C. Faloutsos 50

50

Carnegie Mellon

Linear Regression: idea

patient	*weight*	*height*
1	27	43
2	43	54
3	54	72
…	…	…
N	25	??

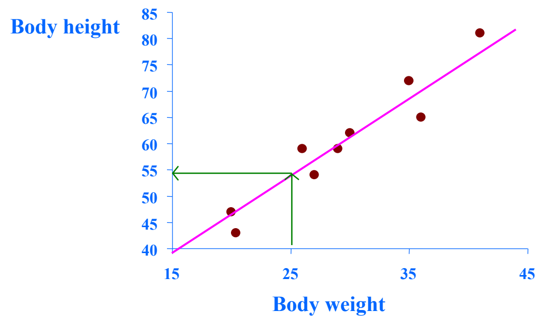

- express what we don't know (= 'dependent variable')
- as a linear function of what we know (= 'indep. variable(s)')

51

Carnegie Mellon

Linear Auto Regression:

Time	*Packets Sent(t)*
1	43
2	54
3	72
…	…
N	??

52

Carnegie Mellon

Linear Auto Regression:

Time	Packets Sent (t-1)	Packets Sent(t)
1	-	43
2	43	54
3	54	72
	…	
…		…
N	25	??

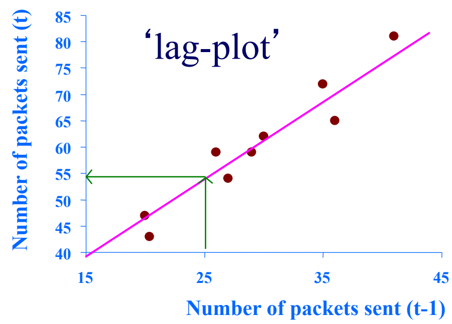

- lag *w*=1
- Dependent variable = # of packets sent (S [t])
- Independent variable = # of packets sent (S[t-1])

Carnegie Mellon

Detailed Outline

- Motivation
- ...
- Linear Forecasting
  - Auto-regression: **Least Squares; RLS**
  - Co-evolving time sequences
  - Examples
  - Conclusions

Carnegie Mellon

More details:

- Q1: Can it work with window *w*>1?
- A1: YES!

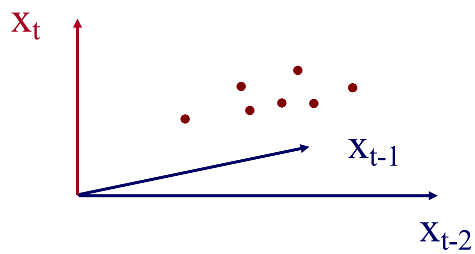

55

Carnegie Mellon

More details:

- Q1: Can it work with window *w*>1?
- A1: YES! (we' ll fit a hyper-plane, then!)

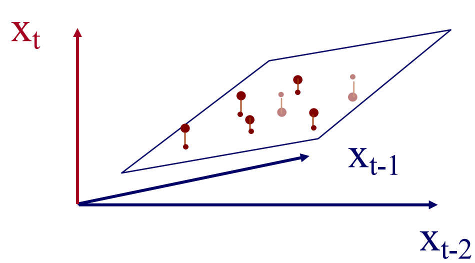

56

Carnegie Mellon

More details:

- Q1: Can it work with window *w*>1?
- A1: YES! (we'll fit a hyper-plane, then!)

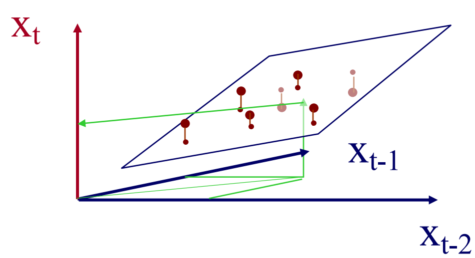

57

Carnegie Mellon

More details:

- Q1: Can it work with window *w*>1?
- A1: YES! The problem becomes:

$$\mathbf{X}_{[N \times w]} \times \mathbf{a}_{[w \times 1]} = \mathbf{y}_{[N \times 1]}$$

- OVER-CONSTRAINED
  - **a** is the vector of the regression coefficients
  - **X** has the *N* values of the *w* indep. variables
  - **y** has the N values of the dependent variable

58

Carnegie Mellon

More details:

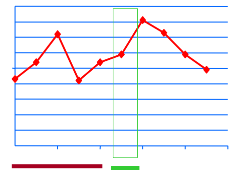

- $\mathbf{X}_{[N \times w]} \times \mathbf{a}_{[w \times 1]} = \mathbf{y}_{[N \times 1]}$

Ind-var1 Ind-var-w

time

$$\begin{bmatrix} X_{11}, X_{12}, \cdots, X_{1w} \\ X_{21}, X_{22}, \ldots, X_{2w} \\ \vdots \\ \vdots \\ \vdots \\ X_{N1}, X_{N2}, \ldots, X_{Nw} \end{bmatrix} \times \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_w \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ \vdots \\ \vdots \\ y_N \end{bmatrix}$$

15-826 59

59

Carnegie Mellon

More details:

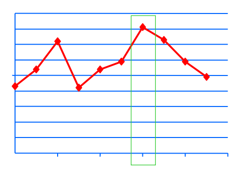

- $\mathbf{X}_{[N \times w]} \times \mathbf{a}_{[w \times 1]} = \mathbf{y}_{[N \times 1]}$

Ind-var1 Ind-var-w

time

$$\begin{bmatrix} X_{11}, X_{12}, \cdots, X_{1w} \\ X_{21}, X_{22}, \ldots, X_{2w} \\ \vdots \\ \vdots \\ \vdots \\ X_{N1}, X_{N2}, \ldots, X_{Nw} \end{bmatrix} \times \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_w \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ \vdots \\ \vdots \\ y_N \end{bmatrix}$$

15-826 60

60

Carnegie Mellon

More details

- Q2: How to estimate $a_1, a_2, \ldots a_w = \mathbf{a}$?
- A2: with Least Squares fit

$$\mathbf{a} = (\mathbf{X}^T \times \mathbf{X})^{-1} \times (\mathbf{X}^T \times \mathbf{y})$$

- (Moore-Penrose pseudo-inverse)
- **a** is the vector that minimizes the RMSE from **y**
- <**identical** math with 'query feedbacks'>

Carnegie Mellon

More details

- Q2: How to estimate $a_1, a_2, \ldots a_w = \mathbf{a}$?
- A2: with Least Squares fit

$$\mathbf{a} = (\mathbf{X}^T \times \mathbf{X})^{-1} \times (\mathbf{X}^T \times \mathbf{y})$$

Identical to earlier formula (proof?)

$$\mathbf{a} = \mathbf{V} \text{ x } \mathbf{\Lambda}^{(-1)} \text{ x } \mathbf{U}^{\mathrm{T}} \times \mathbf{y}$$

Where

$$\mathbf{X} = \mathbf{U} \text{ x } \mathbf{\Lambda} \text{ x } \mathbf{V}^{\mathrm{T}}$$

Carnegie Mellon

More details

- Straightforward solution:

$$\mathbf{a} = (\mathbf{X}^T \times \mathbf{X})^{-1} \times (\mathbf{X}^T \times \mathbf{y})$$

a : Regression Coeff. Vector
X : Sample Matrix

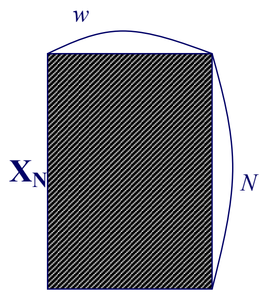

- Observations:
  - Sample matrix X grows over time
  - needs matrix inversion
  - **O**($N \times w^2$) computation
  - **O**($N \times w$) storage

15-826 Copyright (c) 2019 C. Faloutsos 63

63

Carnegie Mellon

Even more details

- Q3: Can we estimate **a** incrementally?
- A3: Yes, with the brilliant, classic method of 'Recursive Least Squares' (RLS) (see, e.g., [Yi+00], for details).
- We can do the matrix inversion, WITHOUT inversion! (How is that possible?!)

15-826 Copyright (c) 2019 C. Faloutsos 64

64

Carnegie Mellon

Sub-Problem: matrix inversion

- How to invert: $(X^T X)^{-1}$
- Incrementally
- WITHOUT inverting (!)

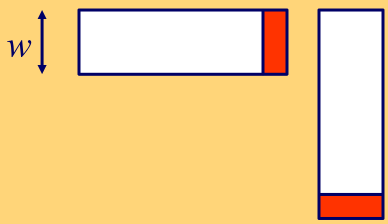

15-826 65

65

Carnegie Mellon

Sub-Problem: matrix inversion

- How to invert: $(\mathbf{X}^T \mathbf{X})^{-1}$
- Incrementally
- WITHOUT inverting (!)

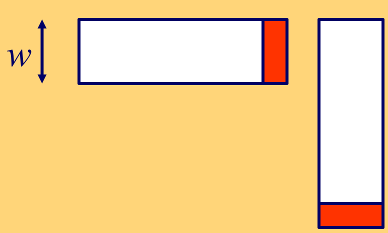

A: keep a $w \times w$ matrix, & update it

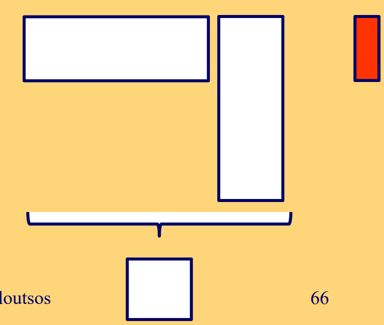

15-826 66

66

Carnegie Mellon

Even more details

- Q3: Can we estimate **a** incrementally?
- A3: Yes, with the brilliant, classic method of 'Recursive Least Squares' (RLS) (see, e.g., [Yi+00], for details).
- We can do the matrix inversion, WITHOUT inversion! (How is that possible?!)
- A: our matrix has special form: $(X^T X)$

15-826 67

67

Carnegie Mellon

Intuition:

- How to compute the average of $x_1, x_2, \dots x_n$
- Incrementally

- Solution: 'sufficient statistics'
  - Count 'k' so far : k += 1
  - Sum 's_k' so far : $s_{k+1} \leftarrow s_k + x_{k+1}$

  Return $s_k \div k$

15-826 68

68

Carnegie Mellon

More details

At the $N+1$ time tick:

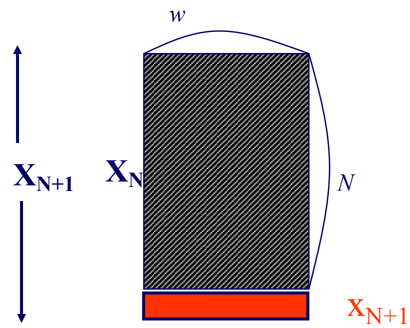

Carnegie Mellon

$\mathbf{a} = (\mathbf{X}^T \times \mathbf{X})^{-1} \times (\mathbf{X}^T \times \mathbf{y})$

More details

- Let $\mathbf{G}_N = (\mathbf{X}_N^T \times \mathbf{X}_N)^{-1}$ (``gain matrix'')
- $\mathbf{G}_{N+1}$ can be computed recursively from $\mathbf{G}_N$

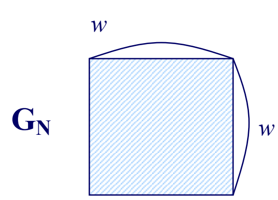

Carnegie Mellon

EVEN more details:

$$G_{N+1} = G_N - [c]^{-1} \times [G_N \times x_{N+1}{}^T] \times x_{N+1} \times G_N$$

1 x *w* row vector

$$c = [1 + x_{N+1} \times G_N \times x_{N+1}{}^T]$$

Let's elaborate
(VERY IMPORTANT, VERY VALUABLE!)

71

Carnegie Mellon

EVEN more details:

$$a = [X_{N+1}{}^T \times X_{N+1}]^{-1} \times [X_{N+1}{}^T \times y_{N+1}]$$

72

Carnegie Mellon

EVEN more details:

$$a = [X_{N+1}^T \times X_{N+1}]^{-1} \times [X_{N+1}^T \times y_{N+1}]$$

[w x 1]

[(N+1) x w]

[(N+1) x 1]

[w x (N+1)]

[w x (N+1)]

Carnegie Mellon

EVEN more details:

$$a = [X_{N+1}^T \times X_{N+1}]^{-1} \times [X_{N+1}^T \times y_{N+1}]$$

[(N+1) x w]

[w x (N+1)]

Carnegie Mellon

EVEN more details:

$$a = [X_{N+1}^T \times X_{N+1}]^{-1} \times [X_{N+1}^T \times y_{N+1}]$$

'gain matrix' $$G_{N+1} \equiv [X_{N+1}^T \times X_{N+1}]^{-1}$$

1 x w row vector

$$G_{N+1} = G_N - [c]^{-1} \times [G_N \times x_{N+1}^T] \times x_{N+1} \times G_N$$

$$c = [1 + x_{N+1} \times G_N \times x_{N+1}^T]$$

15-826 Copyright (c) 2019 C. Faloutsos 75

75

Carnegie Mellon

EVEN more details:

$$G_{N+1} = G_N - [c]^{-1} \times [G_N \times x_{N+1}^T] \times x_{N+1} \times G_N$$

$$c = [1 + x_{N+1} \times G_N \times x_{N+1}^T]$$

15-826 Copyright (c) 2019 C. Faloutsos 76

76

Carnegie Mellon

EVEN more details:

1x1

1xw

wxw wxw wxw wx1 wxw

$$G_{N+1} = G_N - [c]^{-1} \times [G_N \times x_{N+1}{}^T] \times x_{N+1} \times G_N$$

SCALAR! $$c = [1 + x_{N+1} \times G_N \times x_{N+1}{}^T]$$

15-826 Copyright (c) 2019 C. Faloutsos 77

77

Carnegie Mellon

$\mathbf{a} = (\mathbf{X}^T \times \mathbf{X})^{-1} \times (\mathbf{X}^T \times \mathbf{y})$

Altogether:

$$a = [X_{N+1}{}^T \times X_{N+1}]^{-1} \times [X_{N+1}{}^T \times y_{N+1}]$$

$$G_{N+1} \equiv [X_{N+1}{}^T \times X_{N+1}]^{-1}$$

$$G_{N+1} = G_N - [c]^{-1} \times [G_N \times x_{N+1}{}^T] \times x_{N+1} \times G_N$$

$$c = [1 + x_{N+1} \times G_N \times x_{N+1}{}^T]$$

15-826 Copyright (c) 2019 C. Faloutsos 78

78

Carnegie Mellon

Altogether:

$$G_0 \equiv \delta I$$ **IMPORTANT!**

where
I: w x w identity matrix
δ: a large positive number (say, 10^4)

79

Carnegie Mellon

Comparison:

- Straightforward Least Squares
  - Needs huge matrix (**growing** in size) $(O(N \times w)$
  - Costly matrix operation $O(N \times w^2)$
- Recursive LS
  - Need much smaller, fixed size matrix $O(w \times w)$
  - Fast, incremental computation $(O(1 \times w^2)$
  - **no matrix inversion**

$N = 10^6, \quad w = 1\text{-}100$

80

Carnegie Mellon

Pictorially:

- Given:

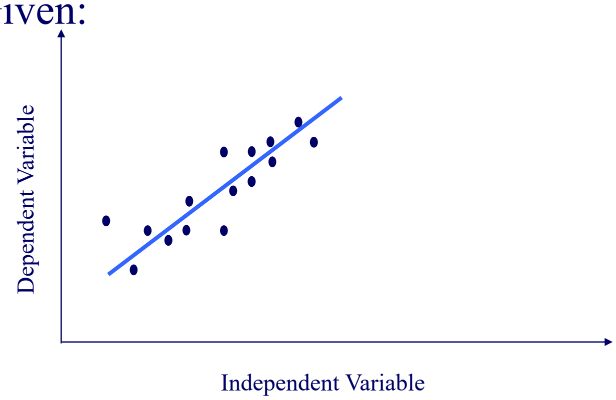

81

Carnegie Mellon

Pictorially:

•

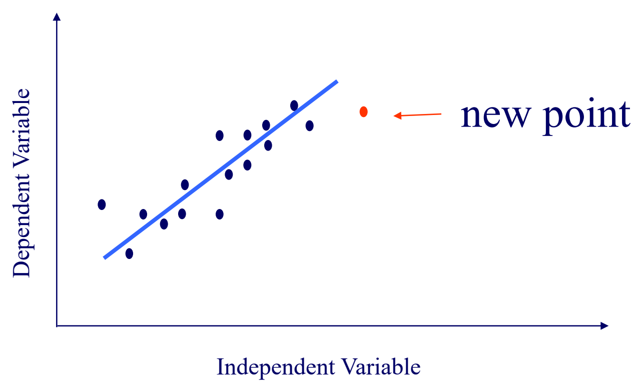

82

Carnegie Mellon

Pictorially:

RLS: quickly compute new best fit

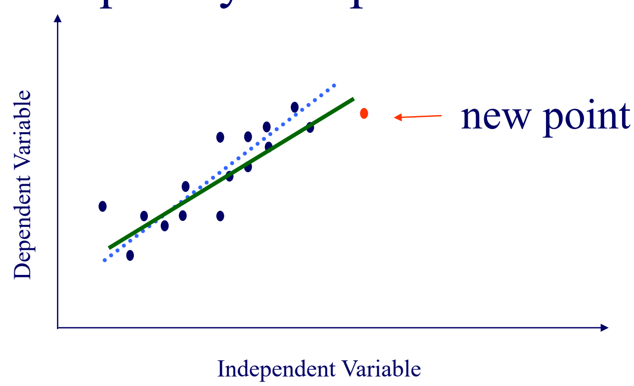

83

Carnegie Mellon

Even more details

- Q4: can we ‘forget’ the older samples?
- A4: Yes - RLS can easily handle that [Yi+00]:

84

Carnegie Mellon

Adaptability - 'forgetting'

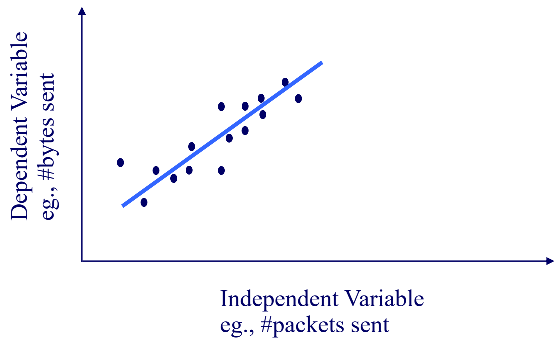

15-826 Copyright (c) 2019 C. Faloutsos 85

85

Carnegie Mellon

Adaptability - 'forgetting'

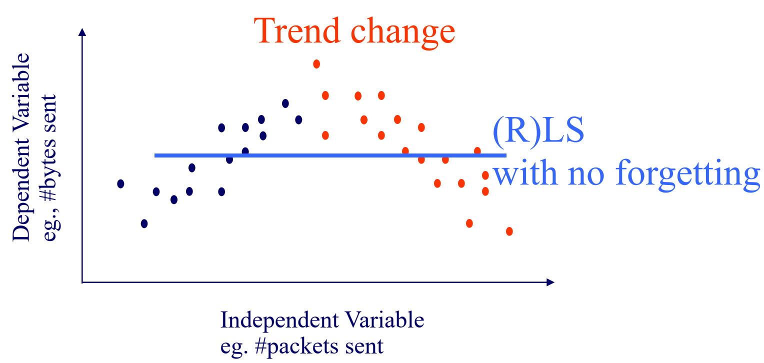

15-826 Copyright (c) 2019 C. Faloutsos 86

86

Carnegie Mellon

Adaptability - ‘forgetting’

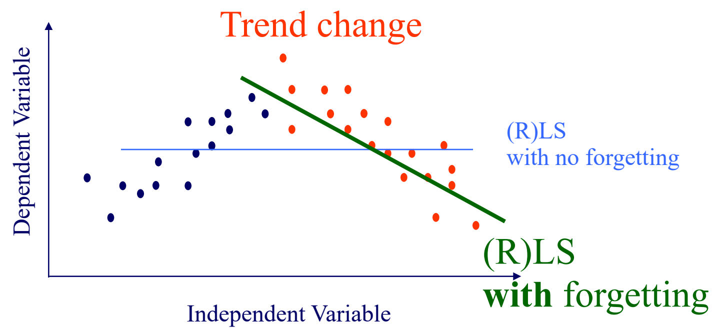

- RLS: can *trivially* handle ‘forgetting’ (see [Yi+,2000])

15-826 Copyright (c) 2019 C. Faloutsos 87

87

Carnegie Mellon

Detailed Outline

- Motivation
- ...
- Linear Forecasting
  - Auto-regression: Least Squares; RLS
  - Co-evolving time sequences
  - Examples
  - Conclusions

15-826 Copyright (c) 2019 C. Faloutsos 88

88

Carnegie Mellon

Co-Evolving Time Sequences

- Given: A set of **correlated** time sequences
- Forecast '**Repeated(t)**'

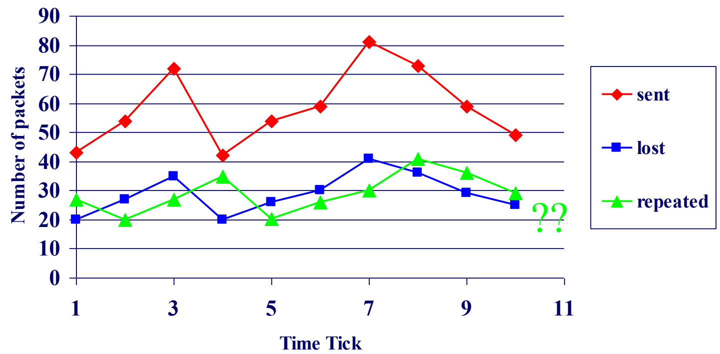

Carnegie Mellon

Solution:

Q: what should we do?

Carnegie Mellon

Solution:

Least Squares, with

- Dep. Variable: Repeated(t)
- Indep. Variables: Sent(t-1) … Sent(t-w); Lost(t-1) …Lost(t-w); Repeated(t-1), ...
- (named: 'MUSCLES' [Yi+00])

15-826 Copyright (c) 2019 C. Faloutsos 91

91

Carnegie Mellon

Forecasting - Outline

- Auto-regression
- Least Squares; recursive least squares
- Co-evolving time sequences
- ➡ Examples
- Conclusions

15-826 Copyright (c) 2019 C. Faloutsos 92

92

Carnegie Mellon

Examples - Experiments

- Datasets
  - Modem pool traffic (14 modems, 1500 time-ticks; #packets per time unit)
  - AT&T WorldNet internet usage (several data streams; 980 time-ticks)
- Measures of success
  - Accuracy : Root Mean Square Error (RMSE)

Carnegie Mellon

Accuracy - "Modem"

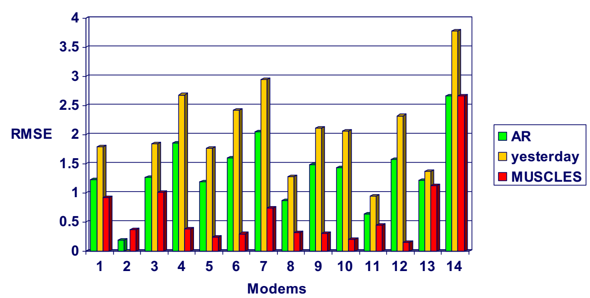

MUSCLES outperforms AR & "yesterday"

Carnegie Mellon

Accuracy - "Internet"

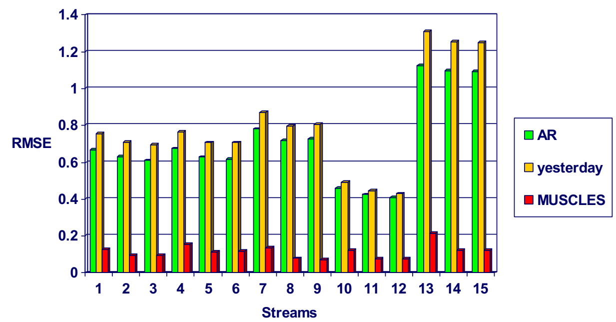

MUSCLES consistently outperforms AR & "yesterday"

Carnegie Mellon

Linear forecasting - Outline

- Auto-regression
- Least Squares; recursive least squares
- Co-evolving time sequences
- Examples
- Conclusions

Carnegie Mellon

Conclusions - Practitioner's guide

- AR(IMA) methodology: prevailing method for linear forecasting
- Brilliant method of Recursive Least Squares for fast, incremental estimation.
- See [Box-Jenkins]

Carnegie Mellon

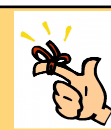

Solution: AR(IMA)

- given x_{t-1}, x_{t-2}, …,
- Q: forecast x_t
- A: AR(IMA) = Box-Jenkins (< Holt-Winters, Kalman)

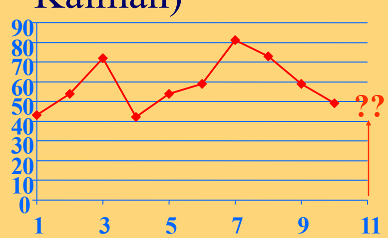

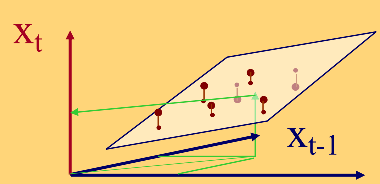

Carnegie Mellon

Resources: software and urls

- free-ware: 'R' for stat. analysis
  (clone of Splus)
  http://cran.r-project.org/
- python script for RLS
  http://www.cs.cmu.edu/~christos/SRC/rls-all.tar

15-826 99

99

Carnegie Mellon

Books

- George E.P. Box and Gwilym M. Jenkins and Gregory C. Reinsel, *Time Series Analysis: Forecasting and Control*, Prentice Hall, 1994 (the classic book on ARIMA, 3rd ed.)
- Brockwell, P. J. and R. A. Davis (1987). Time Series: Theory and Methods. New York, Springer Verlag.

15-826 100

100

Carnegie Mellon

Additional Reading

- [Papadimitriou+ vldb2003] Spiros Papadimitriou, Anthony Brockwell and Christos Faloutsos *Adaptive, Hands-Off Stream Mining* VLDB 2003, Berlin, Germany, Sept. 2003
- [Yi+00] Byoung-Kee Yi et al.: *Online Data Mining for Co-Evolving Time Sequences*, ICDE 2000. (Describes MUSCLES and Recursive Least Squares)

Carnegie Mellon

Outline

- Motivation
- Similarity search and distance functions
- Linear Forecasting
- ➡ Bursty traffic - fractals
- Non-linear forecasting
- Gray box modeling – Lotka Volterra eq's
- Conclusions

Carnegie Mellon

Bursty Traffic & fractals

103

Carnegie Mellon

Detailed Outline

- Motivation
- ...
- Linear Forecasting
- Bursty traffic - fractals
  - Problem
  - Main idea (80/20, Hurst exponent)
  - Results

104

Carnegie Mellon

Reference:

[Wang+02] Mengzhi Wang, Tara Madhyastha, Ngai Hang Chang, Spiros Papadimitriou and Christos Faloutsos, *Data Mining Meets Performance Evaluation: Fast Algorithms for Modeling Bursty Traffic*, ICDE 2002, San Jose, CA, 2/26/2002 - 3/1/2002.

Full thesis: CMU-CS-05-185 *Performance Modeling of Storage Devices using Machine Learning* Mengzhi Wang, Ph.D. Thesis Abstract, .ps.gz, .pdf

105

Carnegie Mellon

Bursty traffic

- Q: Any pattern?

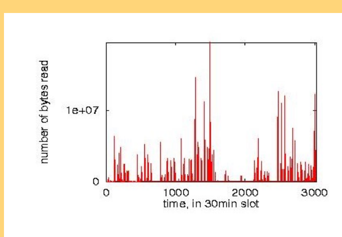

106

Carnegie Mellon

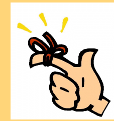

Bursty traffic

- Q: Any pattern?
- A: fractal dimension (‘*b*’ model e.g. 80/20)

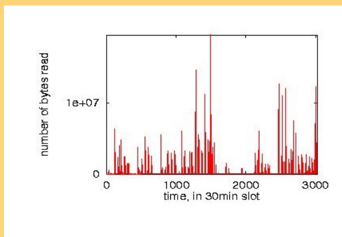

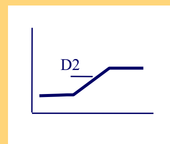

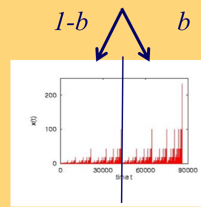

107

Carnegie Mellon

Recall: Problem #1:

Goal: given a signal (eg., #bytes over time)

Find: patterns, periodicities, and/or compress

#bytes

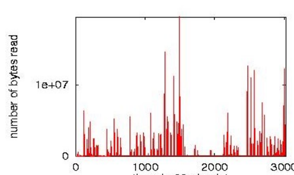

Bytes per 30’
(packets per day;
earthquakes per year)

time

108

Carnegie Mellon

Problem #1

- model bursty traffic
- generate realistic traces
- (Poisson does not work)

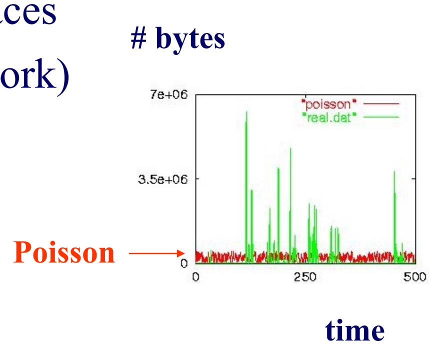

109

Carnegie Mellon

Motivation

- predict queue length distributions (e.g., to give probabilistic guarantees)
- “learn” traffic, for buffering, prefetching, ‘active disks’, web servers

110

Carnegie Mellon

But:

- Q1: How to generate realistic traces; extrapolate; give guarantees?
- Q2: How to estimate the model parameters?

111

Carnegie Mellon

Outline

- Motivation
- ...
- Linear Forecasting
- Bursty traffic - fractals and multifractals
  - Problem
  - Main idea (80/20, Hurst exponent)
  - Results

112

Carnegie Mellon

Approach

- Q1: How to generate a sequence, that is
  - bursty
  - self-similar
  - and has similar queue length distributions

Carnegie Mellon

reminder

Approach

- A: 'binomial multifractal' [Wang+02]
- ~ 80-20 'law':
  - 80% of bytes/queries etc on first half
  - repeat recursively
- *b*: bias factor (eg., 80%)

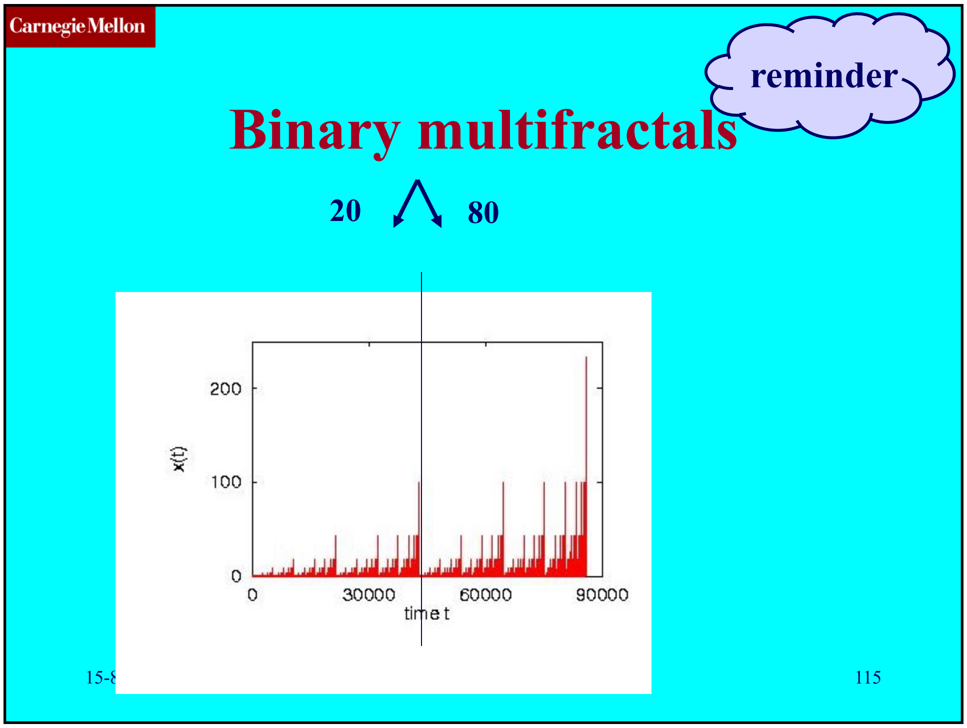
Carnegie Mellon
reminder
Binary multifractals
20
80
x(t)
200
100
0
0
30000
60000
90000
time t
15-8
115

115

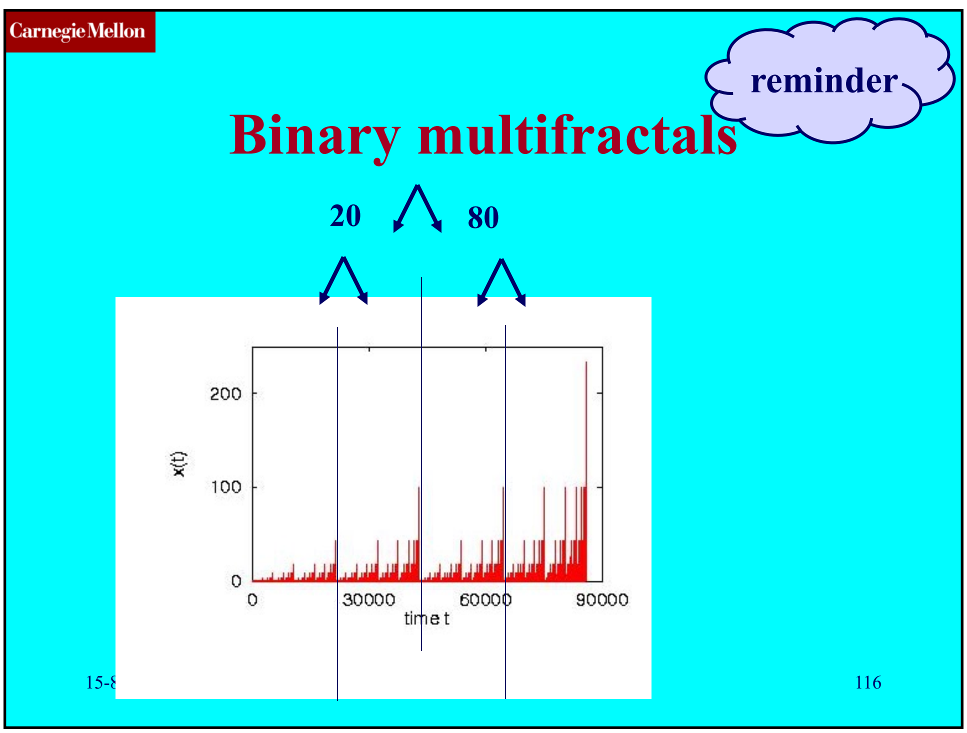
Carnegie Mellon
reminder
Binary multifractals
20
80
x(t)
200
100
0
0
30000
60000
90000
time t
15-8
116

116

Carnegie Mellon

Could you use IFS?

To generate such traffic?

117

Carnegie Mellon

Could you use IFS?

To generate such traffic?

A: Yes – which transformations?

118

Carnegie Mellon

Could you use IFS?

To generate such traffic?

A: Yes – which transformations?

A:

$x' = x / 2$ (p = 0.2)

$x' = x / 2 + 0.5$ (p = 0.8)

Carnegie Mellon

Parameter estimation

- Q2: How to estimate the bias factor *b*?

Carnegie Mellon

Parameter estimation

- Q2: How to estimate the bias factor *b*?
- A: MANY ways [Crovella+96]
  - Hurst exponent
  - variance plot
  - even DFT amplitude spectrum! (‘periodogram’)
  - Fractal dimension (D2)
    - Or D1 (‘entropy plot’ [Wang+02])

15-826 Copyright (c) 2019 C. Faloutsos 121

121

Carnegie Mellon

Fractal dimension

- Real (and 80-20) datasets can be in-between: bursts, gaps, smaller bursts, smaller gaps, at every scale

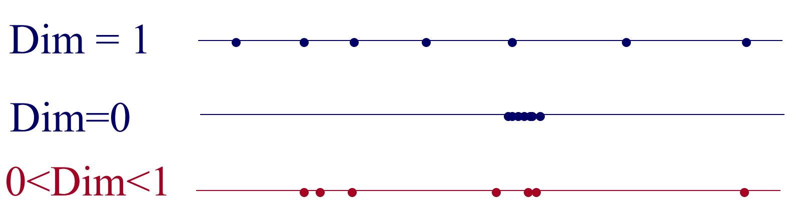

15-826 Copyright (c) 2019 C. Faloutsos 122

122

Carnegie Mellon

Estimating 'b'

- **Exercise**: Show that

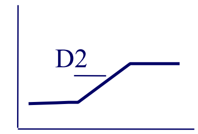

$$D_2 = -\log_2 (b^2 + (1-b)^2)$$

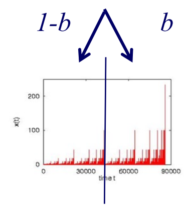

Sanity checks:

- b =1.0 D_2 = ??
- b =0.5 D_2 = ??

123

Carnegie Mellon

(Fractals, again)

- What set of points could have behavior between point and line?

124

Carnegie Mellon

Cantor dust

- Eliminate the middle third
- Recursively!

15-826 Copyright (c) 2019 C. Faloutsos 125

125

Carnegie Mellon

Cantor dust

15-826 Copyright (c) 2019 C. Faloutsos 126

126

Carnegie Mellon

Cantor dust

127

Carnegie Mellon

Cantor dust

128

Carnegie Mellon

Cantor dust

Carnegie Mellon

Cantor dust

Dimensionality?
(no length; infinite # points!)
Answer: log2 / log3 = 0.6

Carnegie Mellon

Some plots:

- Poisson vs real

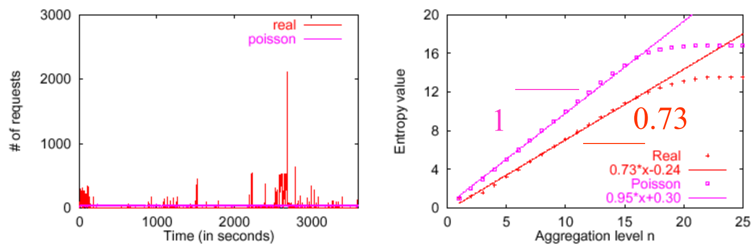

Poisson: slope = ~1 -> uniformly distributed

15-826 Copyright (c) 2019 C. Faloutsos 131

131

Carnegie Mellon

Conclusions

- 80/20, '*b*-model', Multiplicative Wavelet Model (MWM), for analysis and synthesis of bursty traffic

15-826 Copyright (c) 2019 C. Faloutsos 132

132

Bursty traffic

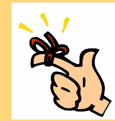

- Q: Any pattern?

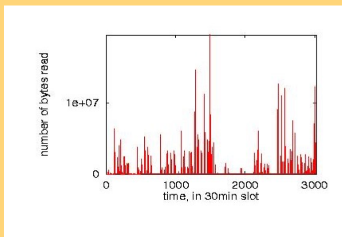

- A: fractal dimension (‘*b*’ model e.g. 80/20)

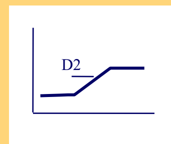

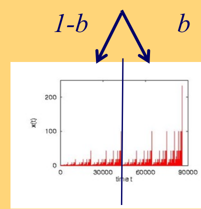

Books

- <u>Fractals</u>: Manfred Schroeder: *Fractals, Chaos, Power Laws: Minutes from an Infinite Paradise* W.H. Freeman and Company, 1991 (Probably the BEST book on fractals!)

Carnegie Mellon

Further reading:

- Crovella, M. and A. Bestavros (1996). Self-Similarity in World Wide Web Traffic, Evidence and Possible Causes. Sigmetrics.
- [ieeeTN94] W. E. Leland, M.S. Taqqu, W. Willinger, D.V. Wilson, *On the Self-Similar Nature of Ethernet Traffic,* IEEE Transactions on Networking, 2, 1, pp 1-15, Feb. 1994.

15-826 135

135

Carnegie Mellon

Further reading

Entropy plots

- [Riedi+99] R. H. Riedi, M. S. Crouse, V. J. Ribeiro, and R. G. Baraniuk, *A Multifractal Wavelet Model with Application to Network Traffic*, IEEE Special Issue on Information Theory, 45. (April 1999), 992-1018.
- [Wang+02] Mengzhi Wang, Tara Madhyastha, Ngai Hang Chang, Spiros Papadimitriou and Christos Faloutsos, *Data Mining Meets Performance Evaluation: Fast Algorithms for Modeling Bursty Traffic*, ICDE 2002, San Jose, CA, 2/26/2002 - 3/1/2002.

15-826 136

136

Carnegie Mellon

Outline

- Motivation
- ...
- Linear Forecasting
- Bursty traffic - fractals and multifractals
- ➡ Non-linear forecasting
- Gray box modeling – Lotka Volterra eq's
- Conclusions

Carnegie Mellon

Detailed Outline

- Non-linear forecasting
  - Problem
  - Idea
  - How-to
  - Experiments
  - Conclusions

Carnegie Mellon

Reference:

[Deepay Chakrabarti and Christos Faloutsos *F4: Large-Scale Automated Forecasting using Fractals* CIKM 2002, Washington DC, Nov. 2002.]

139

Carnegie Mellon

Problem: Forecast

- given x_{t-1}, x_{t-2}, …, (‘chaotic’/non-linear)
- Q: forecast x_t

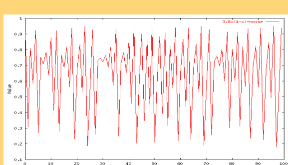

140

Carnegie Mellon

Solution

- given x_{t-1}, x_{t-2}, …, (‘chaotic’/non-linear)
- Q: forecast x_t
- A: lag-plots + sim. search (= ‘Delayed Coordinate Embedding’)

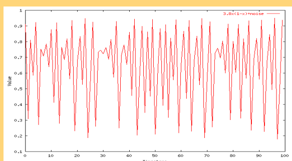

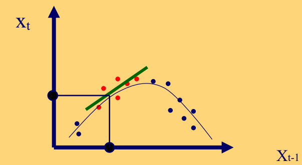

15-826 Copyright (c) 2019 C. Faloutsos 141

141

Carnegie Mellon

Recall: Problem #1

Value

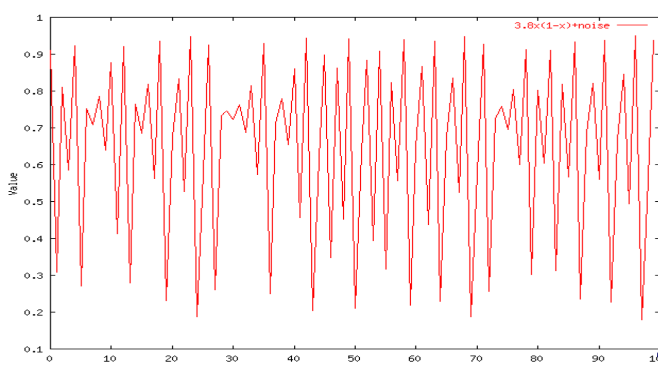

Time

Given a time series $\{x_t\}$, predict its future course, that is, x_{t+1}, x_{t+2}, ...

15-826 Copyright (c) 2019 C. Faloutsos 142

142

How to forecast?

- ARIMA - but: linearity assumption
- ANSWER: 'Delayed Coordinate Embedding' = Lag Plots [Sauer94]

143

ARIMA pitfall

Example: logistic parabola

Models population of flies [R. May/1976]

$$x_{t+1} = a x_t \cdot (1 - x_t)$$

Time-series plot

x_t

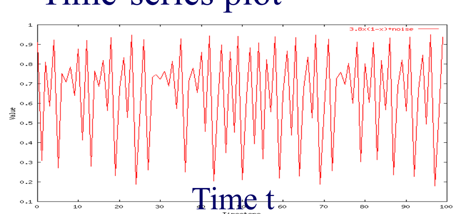

Time t

Logistic map

x_{t+1}

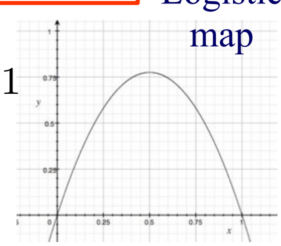

x_t

144

Carnegie Mellon

ARIMA pitfall

Example: logistic parabola

Models population of flies [R. May/1976]

$$x_{t+1} = a x_t \cdot (1 - x_t)$$

Logistic map

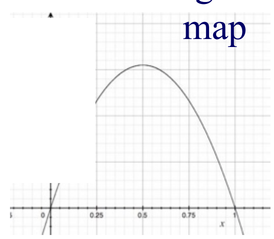

- = SI virus prop. model
- ~ Bass equation (market penetration)
- Special case of Lotka-Volterra

x_t

15-826 Copyright (c) 2019 C. Faloutsos 145

145

Carnegie Mellon

ARIMA pitfall

Linear equations, e.g., AR, ARIMA, …

x_{t+1}

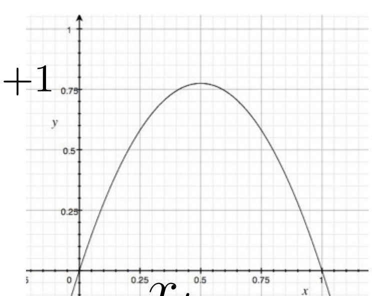

x_t

15-826 Copyright (c) 2019 C. Faloutsos 146

146

Carnegie Mellon

ARIMA pitfall

Linear equations, e.g., AR, ARIMA, …

e.g., AR(1)

$$x_{t+1} = a x_t + \epsilon$$

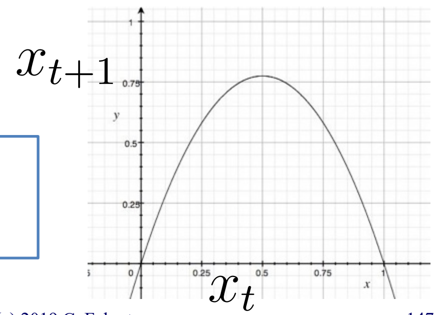

147

Carnegie Mellon

ARIMA pitfall

Linear equations, e.g., AR, ARIMA, …

e.g., AR(1)

$$x_{t+1} = a x_t + \epsilon$$

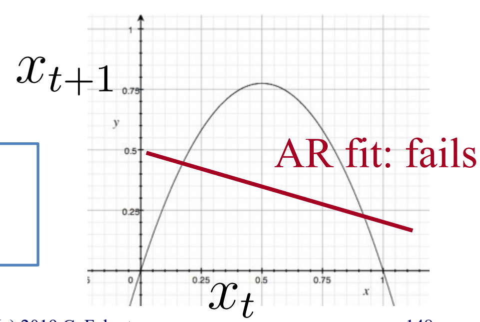

148

Carnegie Mellon

Solution?

"Delayed Coordinate Embedding"

= Lag Plots

[Sauer94]

k-nearest neighbor search

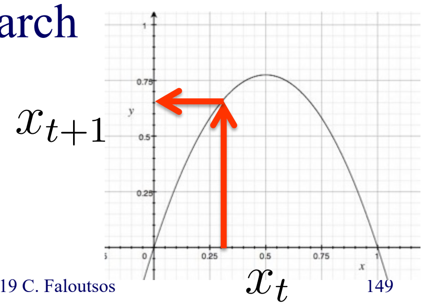

149

Carnegie Mellon

General Intuition (Lag Plot)

Lag = 1,
k = 4 NN

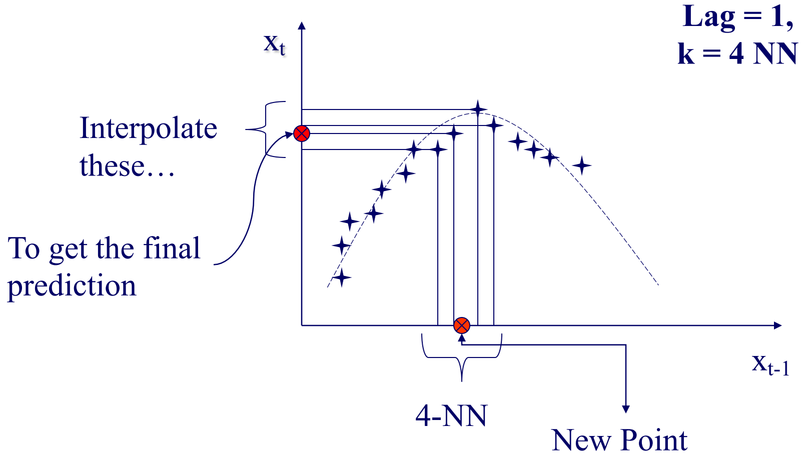

150

Carnegie Mellon

Questions:

- Q1: How to choose lag *L*?
- Q2: How to choose *k* (the # of NN)?
- Q3: How to interpolate?
- Q4: why should this work at all?

Carnegie Mellon

Q1: Choosing lag *L*

- Manually (16, in award winning system by [Sauer94])

Carnegie Mellon

Q2: Choosing number of neighbors *k*

- Manually (typically ~ 1-10)

15-826 153

153

Carnegie Mellon

Q3: How to interpolate?

How do we interpolate between the *k* nearest neighbors?

A1: Average

A2: Weighted average (weights drop with distance - how?)

15-826 154

154

Carnegie Mellon

Q3: How to interpolate?

A3: Using SVD - seems to perform best ([Sauer94] - first place in the Santa Fe forecasting competition)

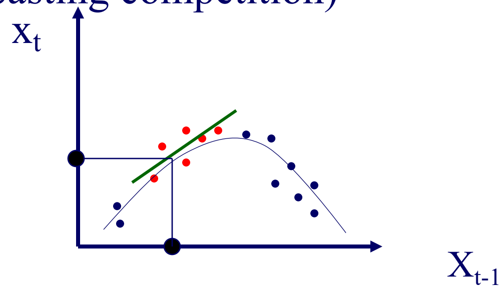

15-826 Copyright (c) 2019 C. Faloutsos 155

155

Carnegie Mellon

Q4: Any theory behind it?

A4: YES!

15-826 Copyright (c) 2019 C. Faloutsos 156

156

Carnegie Mellon

Details

Theoretical foundation

- Based on the "Takens' Theorem" [Takens81]
- which says that long enough delay vectors can do prediction, even if there are unobserved variables in the dynamical system (= diff. equations)

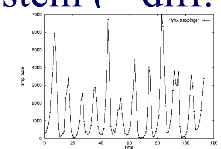

157

Carnegie Mellon

Details

Theoretical foundation

Example: Lotka-Volterra equations

$$dH/dt = r\,H - a\,H{*}P$$
$$dP/dt = b\,H{*}P - m\,P$$

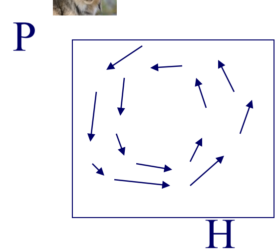

H is count of prey (e.g., hare)
P is count of predators (e.g., lynx)

Suppose only P(t) is observed (t=1, 2, …).

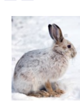

158

Carnegie Mellon

Details

Theoretical foundation

- But the delay vector space is a faithful reconstruction of the internal system state
- So prediction in **delay vector space** is as good as prediction in **state space**

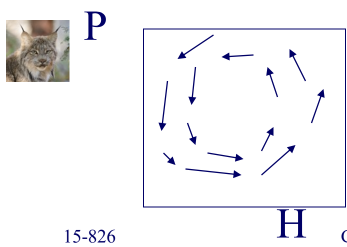

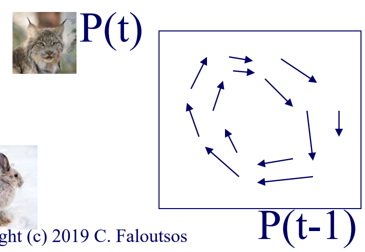

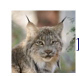

15-826 Copyright (c) 2019 C. Faloutsos 159

159

Carnegie Mellon

Detailed Outline

- Non-linear forecasting
  - Problem
  - Idea
  - How-to
  - Experiments
  - Conclusions

15-826 Copyright (c) 2019 C. Faloutsos 160

160

Datasets

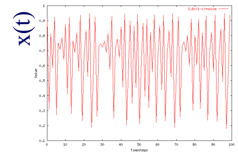

Logistic Parabola:

$x_t = ax_{t-1}(1-x_{t-1})$ + noise

Models population of flies [R. May/1976]

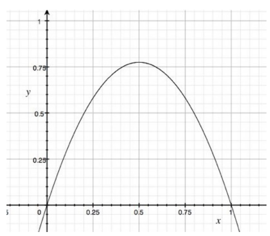

Lag-plot

15-826 Copyright (c) 2019 C. Faloutsos 161

161

Logistic Parabola

Our Prediction from here

Value

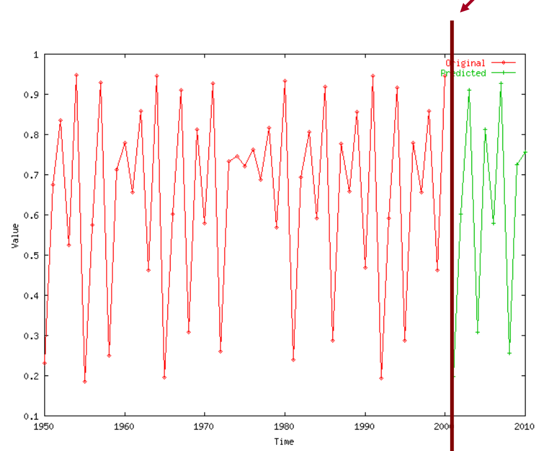

Timesteps

15-826 Copyright (c) 2019 C. Faloutsos 162

162

Carnegie Mellon

Logistic Parabola

Comparison of prediction to correct values

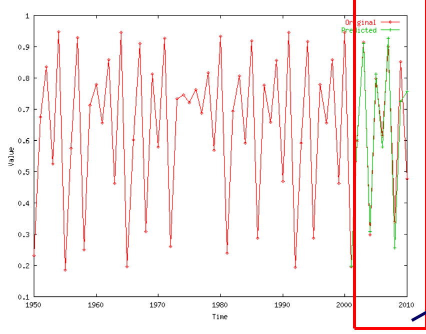

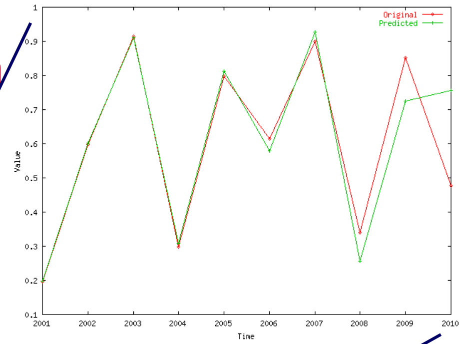

15-826 Copyright (c) 2019 C. Faloutsos 163

163

Carnegie Mellon

Datasets

LORENZ: Models convection currents in the air

dx / dt = a (y - x)

dy / dt = x (b - z) - y

dz / dt = xy - c z

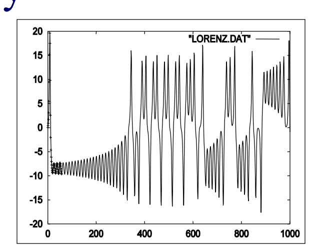

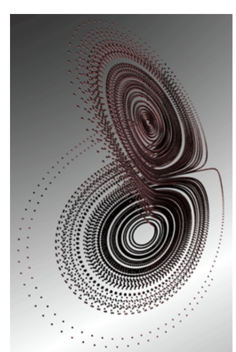

15-826 Copyright (c) 2019 C. Faloutsos 164

164

Carnegie Mellon

Datasets

LORENZ: Models convection currents in the air

dx / dt = a (y - x)

dy / dt = x (b - z) - y

dz / dt = xy - c z

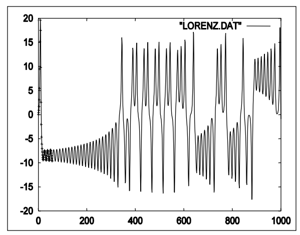

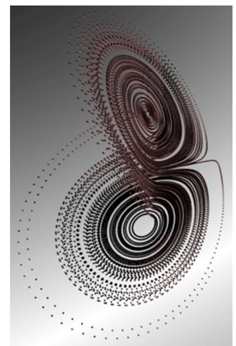

15-826 Copyright (c) 2019 C. Faloutsos 165

165

Carnegie Mellon

LORENZ

Comparison of prediction to correct values

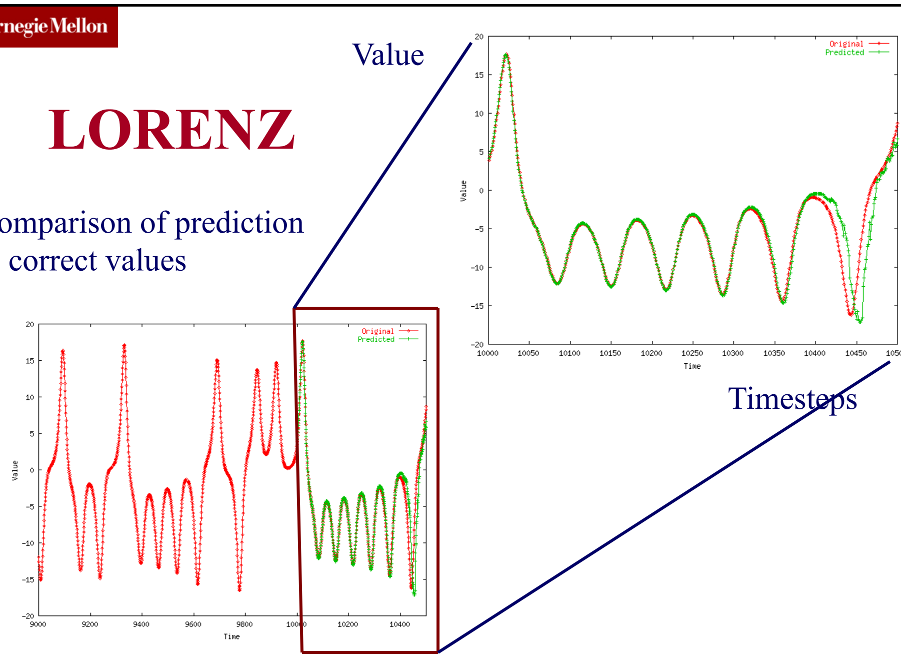

15-826 Copyright (c) 2019 C. Faloutsos 166

166

Value

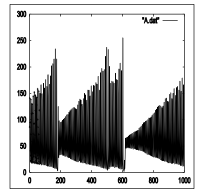

Datasets

- LASER: fluctuations in a Laser over time (used in Santa Fe competition)

Time

Value

Laser

Comparison of prediction to correct values

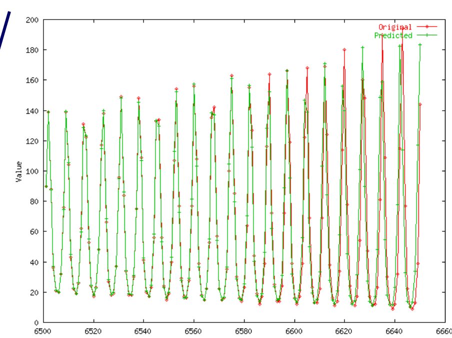

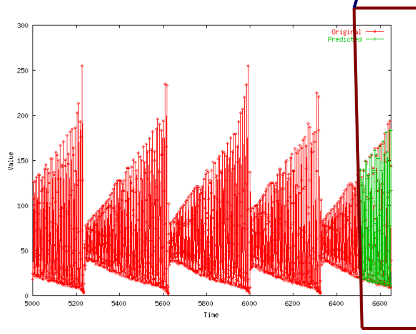

Timesteps

Carnegie Mellon

Conclusions

- Lag plots for non-linear forecasting (Takens' theorem)
- suitable for 'chaotic' signals

Carnegie Mellon

Outline

- Motivation
- ...
- Linear Forecasting
- Bursty traffic - fractals and multifractals
- Non-linear forecasting
- ➡ Gray box modeling – Lotka Volterra eq's
- Conclusions

Carnegie Mellon

Problem: Gray-box forecast

Q: How to model (competing) species/products/ideas?

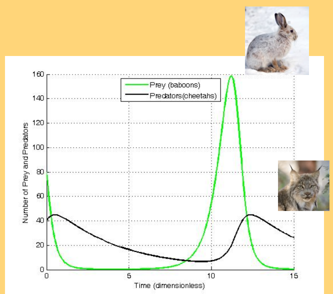

15-826 Copyright (c) 2019 C. Faloutsos 171

171

Carnegie Mellon

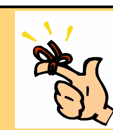

Answer

Q: How to model (competing) species/products/ideas?

A: Lotka-Volterra

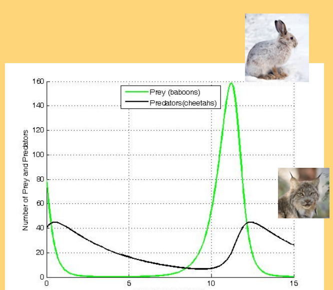

$$P_i(t+1) = P_i(t)\left[1 + r_i\left(1 - \frac{\sum_{j=1}^{d} a_{ij}P_j(t)}{K_i}\right)\right], \quad (i = 1, \cdots, d), \quad (3)$$

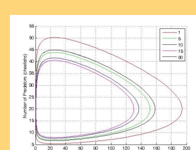

15-826 Copyright (c) 2019 C. Faloutsos 172

172

Carnegie Mellon

Details

Theoretical foundation

Example: Lotka-Volterra equations

$$dH/dt = r\,H - a\,H{*}P$$
$$dP/dt = b\,H{*}P - m\,P$$

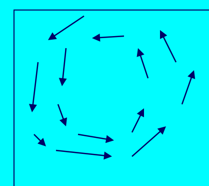

H is count of prey (e.g., hare)
P is count of predators (e.g., lynx)

Suppose only P(t) is observed (t=1, 2, …).

15-826 Copyright (c) 2019 C. Faloutsos 173

173

Carnegie Mellon

Solution to Lotka-Volterra eq.

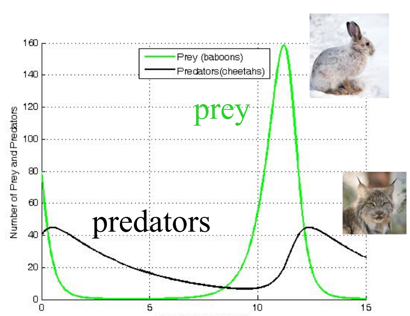

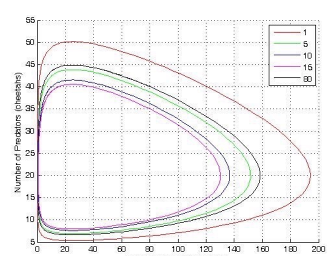

15-826 174

from wikipedia

174

Carnegie Mellon

Compare to reality:

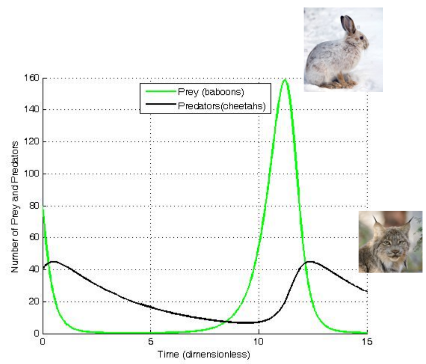

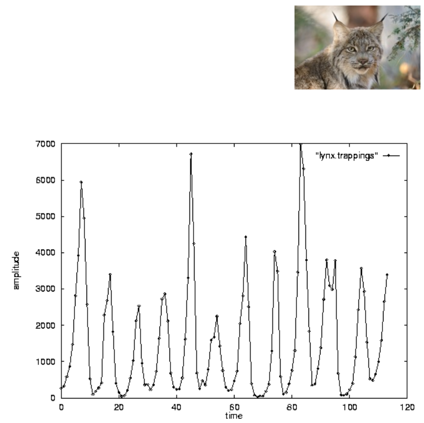

175

Carnegie Mellon

Notice: LV are vital!

Example: Lotka-Volterra equations

$$dH/dt = r\,H - a\,H*P$$
$$dP/dt = b\,H*P - m\,P$$

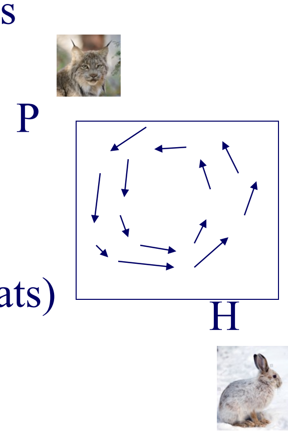

- Prey-predator
- Competing animals (rabbits/goats)
- Self-competition (Bass model)
- Competing products (stocks/bonds)

176

Carnegie Mellon

The Web as a Jungle: Non-Linear Dynamical Systems for Co-evolving Online Activities


Yasuko Matsubara (Kumamoto University)
Yasushi Sakurai (Kumamoto University)
Christos Faloutsos (CMU)

Open source code: here

177

Carnegie Mellon

Given: online user activities

Xbox, PlayStation, Wii, Android

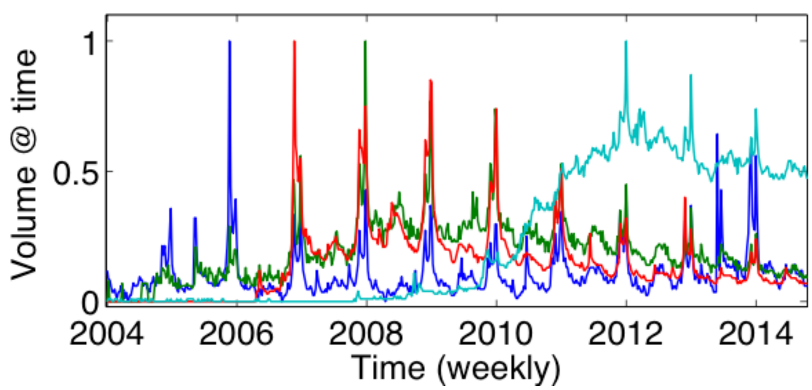

178

Carnegie Mellon

The Web as a jungle

Ecosystem on the **Web**

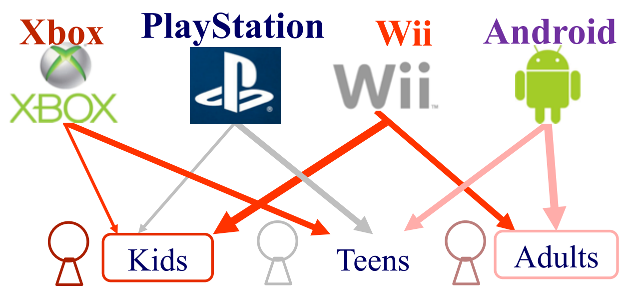

15-826 Copyright (c) 2019 C. Faloutsos 179

179

Carnegie Mellon

The Web as a jungle

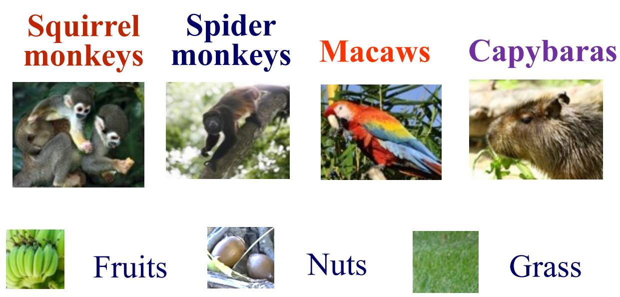

Ecosystem in the **Jungle**

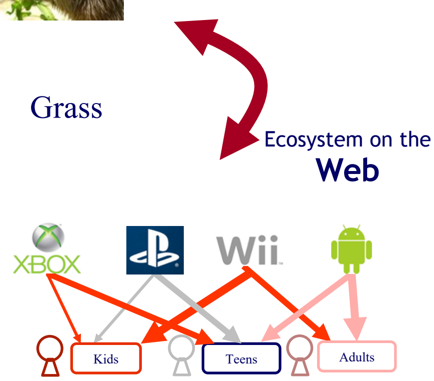

15-826 Copyright (c) 2019 C. Faloutsos 180

180

Carnegie Mellon

LV equations

Interaction between multiple (‘*d*’) species/products/viruses

 … 

 

$$P_i(t+1) = P_i(t)\left[1 + r_i\left(1 - \frac{\sum_{j=1}^{d} a_{ij} P_j(t)}{K_i}\right)\right], \quad (i = 1, \cdots, d),$$

a_{ij} - effect of species j on species i

- (positive: hurts)

15-826 Copyright (c) 2019 C. Faloutsos 181

181

Carnegie Mellon

EcoWeb at work - forecasting

Train: 2/3 sequences	Forecast: 1/3 following years

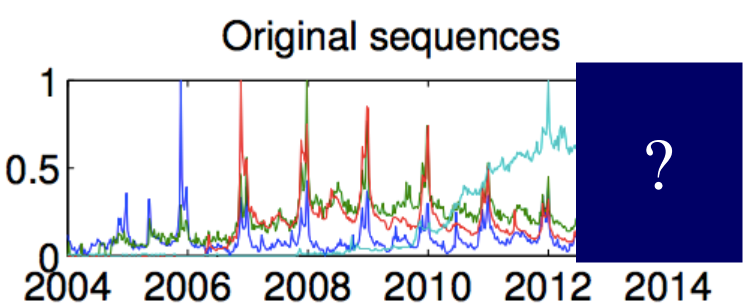

15-826 Copyright (c) 2019 C. Faloutsos 182

182

Carnegie Mellon

EcoWeb at work - forecasting

Train:
2/3 sequences

Forecast:
1/3 following years

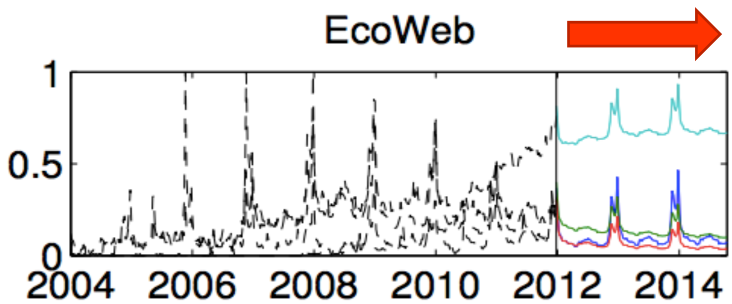

EcoWeb can capture future patterns

15-826 Copyright (c) 2019 C. Faloutsos 183

183

Carnegie Mellon

EcoWeb at work - forecasting

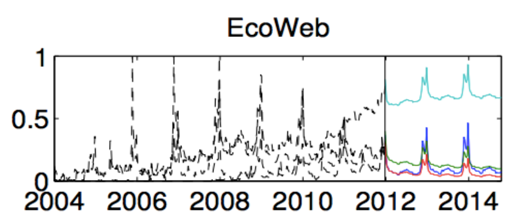

Open source code: here

15-826 Copyright (c) 2019 C. Faloutsos 184

184

Carnegie Mellon

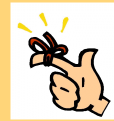

Answer

Q: How to model (competing) species/products/ideas?

A: Lotka-Volterra

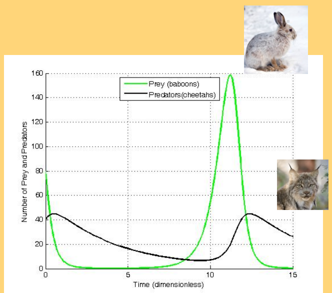

$$P_i(t+1) = P_i(t)\left[1 + r_i\left(1 - \frac{\sum_{j=1}^{d} a_{ij}P_j(t)}{K_i}\right)\right], \quad (i = 1, \cdots, d), \quad (3)$$

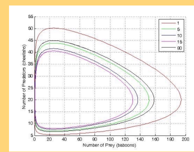

185

Carnegie Mellon

Outline

- Motivation
- ...
- Linear Forecasting
- Bursty traffic - fractals and multifractals
- Non-linear forecasting
- Gray box modeling – Lotka Volterra eq's
- Conclusions

186

Carnegie Mellon

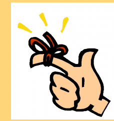

Overall conclusions

- Similarity search: **Euclidean**/time-warping; **feature extraction** and **SAMs**
- Linear Forecasting: **AR** (Box-Jenkins)
- Non-linear forecasting: **lag-plots**
- Gray-box modeling: **Lotka-Volterra**

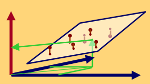

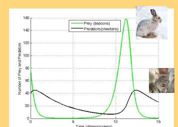

187

Carnegie Mellon

References

- Deepay Chakrabarti and Christos Faloutsos *F4: Large-Scale Automated Forecasting using Fractals* CIKM 2002, Washington DC, Nov. 2002.
- Sauer, T. (1994). *Time series prediction using delay coordinate embedding*. (in book by Weigend and Gershenfeld, below) Addison-Wesley.
- Takens, F. (1981). *Detecting strange attractors in fluid turbulence*. Dynamical Systems and Turbulence. Berlin: Springer-Verlag.

188

Carnegie Mellon

References

- Weigend, A. S. and N. A. Gerschenfeld (1994). *Time Series Prediction: Forecasting the Future and Understanding the Past*, Addison Wesley. (Excellent collection of papers on chaotic/non-linear forecasting, describing the algorithms behind the winners of the Santa Fe competition.)

