

Mining Time Series: Tools and Applications

Christos Faloutsos

CMU SCS




<https://www.cs.cmu.edu/~christos/TALKS/21-09-exec-ed/>

こんにちは

Thanks to:



- Prof. Zachary Lipton



- Susan Caplan

Outline



- ➔ • Introduction - Motivation
- P1. Similarity Search and Indexing
- P2. DSP (Digital Signal Processing)
- P3. Linear Forecasting
- P4. Non-linear forecasting
- Conclusions

Problem definition

- Given: one or more sequences



$x_1, x_2, \dots, x_t, \dots$

$(y_1, y_2, \dots, y_b, \dots$

$\dots)$


- Find

- **Forecast**; similar sequences
- patterns; clusters; outliers



Background of audience

- Finance / bank



- Insurance company

- Manufacturer (chemical; electronic devices)




- Automotive



Background of audience

- Finance / bank
 - Forecasting; fraud/anomaly detection
- Insurance company
 - Same
- Manufacturer (chemical; electronic devices)
 - Sensor measurements; supply-chain forecasting
- Automotive
 - same



Questions from audience

- Are there any examples that have used DX/AI to innovate in BtoB sales, including new sales methods and the development of new markets or needs?

- **[working on it, with a financial institution; 'recommendation']**
- **everybody is using forecasting (traditional AR etc; and Deep Learning)**
- **Cloud services offer such functionality, eg,**
 - **AWS: <https://aws.amazon.com/forecast/>**
 - **MS/Azure: <https://docs.microsoft.com/en-us/azure/machine-learning/how-to-auto-train-forecast>**
 - **Google: <https://cloud.google.com/blog/products/data-analytics/get-started-with-data-analytics-demand-forecasting-with-ml-models>**

Outline



- ➔ • Introduction - Motivation
- P1. Similarity Search and Indexing
- P2. DSP (Digital Signal Processing)
- P3. Linear Forecasting
- P4. Non-linear forecasting
- Conclusions



Motivation - Applications

- Financial, sales, economic series



Alcoa Corp.

American express

Citi group



ECG - physionet.org



EEG - epilepsy



Motivation - Applications (cont'd)

- civil/automobile infrastructure
 - bridge vibrations [Oppenheim+02]



Tokyo Gate
Bridge

Motivation - Applications (cont'd)

- civil/automobile infrastructure
 - bridge vibrations [Oppenheim+02]



Tokyo Gate
Bridge

- <https://www.dm.sanken.osaka-u.ac.jp/~yasuko/TALKS/17-KDD-tut/>



Prof. Yasushi Sakurai



Prof. Yasuko Matsubara

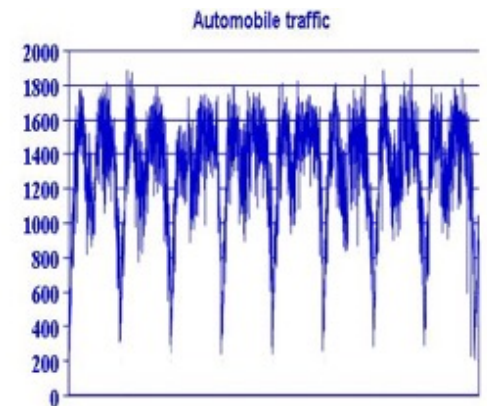


Motivation - Applications (cont'd)

- civil/automobile infrastructure
 - road conditions / traffic monitoring



From
www.transportation.gov




Overall Problem:

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

Find: patterns, periodicities, and/or compress



count




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

year


Problem#1: Similarity search

Eg., Find a 10-tick pattern, similar to the last one



Problem#1: Similarity search

Eg., Find a 10-tick pattern, similar to the last one




Problem #2: Patterns

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

Find: patterns, periodicities, and/or compress



count



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

year

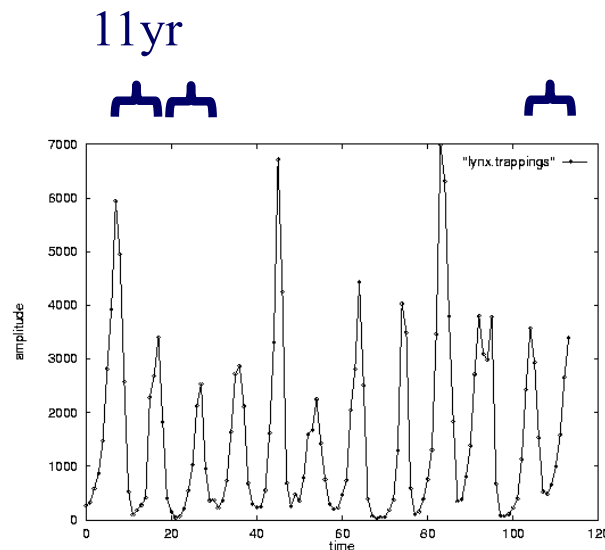
Problem #2: Patterns

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

Find: patterns, periodicities, and/or compress



count

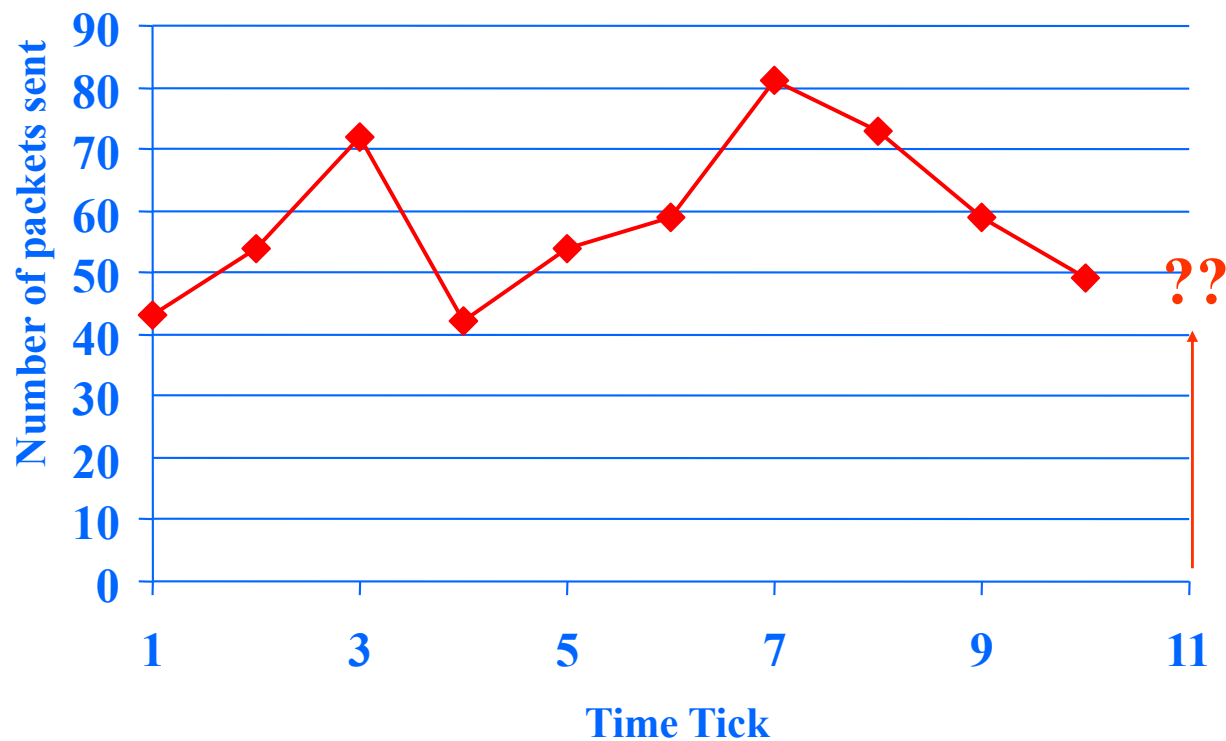


lynx caught per year
(packets per day;
temperature per day)

year

Problem #3: Forecast

Given x_t, x_{t-1}, \dots , forecast x_{t+1}



Outline



- Introduction - Motivation
 - ➔ – Administrative announcements
- P1. Similarity Search and Indexing
- P2. DSP (Digital Signal Processing)
- P3. Linear Forecasting
- P4. Non-linear forecasting
- Conclusions


Problem#1

Problem#2

} Problem#3



Important observations




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

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

Prob.#3


Prob.#2

Prob.#1



'Recipe' Structure:

- Problem definition
- Short answer/solution
- LONG answer – details
- Conclusion/short-answer



Check-point questions



With yellow back-ground

Outline



- Introduction - Motivation
 - ➔ – 2-slide summary of this tutorial
- P1. Similarity Search and Indexing
- P2. DSP (Digital Signal Processing)
- P3. Linear Forecasting
- P4. Non-linear forecasting
- Conclusions

Problem#1


Problem#2

} Problem#3



Problem:



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





Answers

- P1. Similarity search: **Euclidean/time-warping; feature extraction and SAMs**
- P2. Periodicities: **DFT/DWT**
- P3. Linear Forecasting: **AR (Box-Jenkins)**
- P4. Non-linear forecasting: **lag-plots**



Outline



- Introduction - Motivation
- ➔ • P1. Similarity Search and Indexing
- P2. DSP (Digital Signal Processing)
- P3. Linear Forecasting
- P4. Non-linear forecasting
- Conclusions

Outline





- Introduction - Motivation
- P1. Similarity Search and Indexing
 - ➔ – P1.1. distance functions: Euclidean; Time-warping
 - P1.2. indexing
 - P1.3. feature extraction
- P2. DSP (Digital Signal Processing)
- ...



Problem:

Q: How similar are two sequences?







Answer:

Q: How similar are two sequences?

A: Euclidean distance (\leftrightarrow cosine similarity)

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




Importance of distance functions

Subtle, but **absolutely necessary**:


- A ‘must’ for similarity indexing (-> forecasting)
- A ‘must’ for clustering

Two major families

- Euclidean and L_p norms
- Time warping and variations



A1) 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_∞

A2) Time Warping


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



Time Warping



‘stutters’:

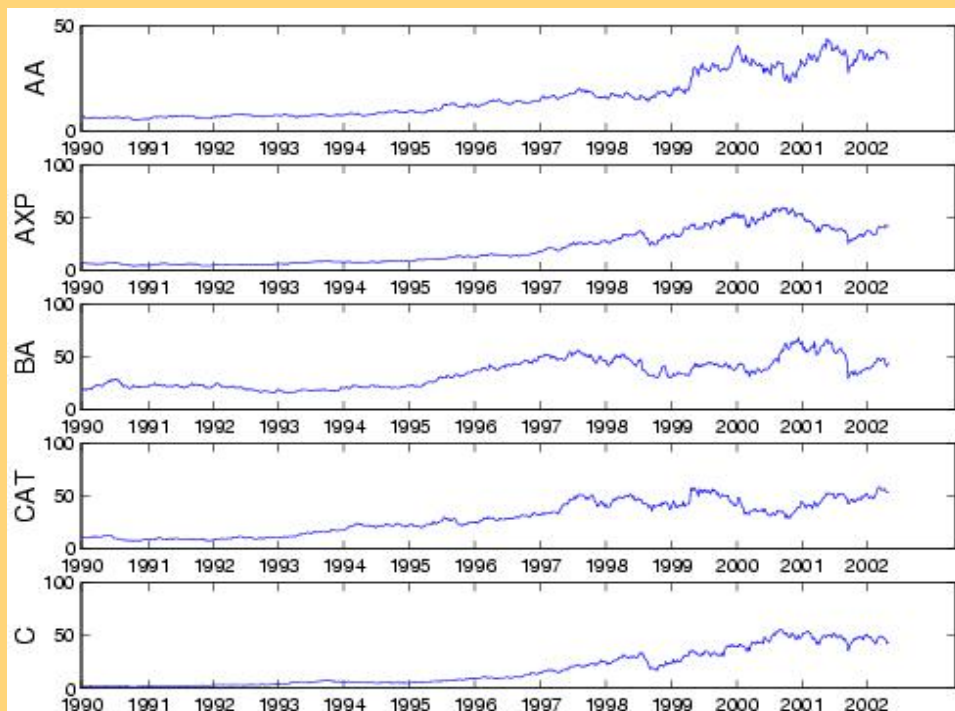


Answer:




Q: How similar are two sequences?

A: Euclidean distance



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



Outline





- Motivation
- P1. Similarity Search and Indexing
 - P1.1. distance functions: Euclidean; Time-warping
 - ➔ – P1.2. indexing
 - P1.3. feature extraction
- P2. DSP (Digital Signal Processing)
- ...



Problem:

Q: find quickly stocks like 'C' (or customers like 'smith')









Answer:

Q: find quickly stocks like 'C' (or customers like 'smith')

A: summarize seq. to a few numbers/features (eg., avg, stdv, Fourier coeff.)




\$price




⋮

\$price



\$price



distance function: by expert

Idea: 'GEMINI'


Eg., *'find stocks similar to MSFT'*

Seq. scanning: too slow


How to accelerate the search?

[Faloutsos96]

'GEMINI' - Pictorially




GEMINI



Solution: Quick-and-dirty' filter:

- extract n features (numbers, eg., avg., etc.)
- map into a point in n -d feature space
- organize points with off-the-shelf spatial access method ('SAM')
- discard false alarms

GEMINI



Solution: Quick-and-dirty' filter:

- extract n features (numbers, eg., avg., etc.)
- map into a point in n -d feature space
 - = feature extraction
 - = feature engineering
 - = dimensionality reduction
 - Singular value decomposition (SVD)
 - Independent Component Analysis (ICA)
 - = embedding (eg, with Deep Learning)

Examples of GEMINI

- Time sequences: DFT (up to 100 times faster) [Faloutsos, 1996]

Conclusions




- Fast indexing: through GEMINI
 - feature extraction and
 - (off the shelf) Spatial Access Methods [Gaede+'98], or nearest-neighbor methods (Locality Sensitive Hashing LSH [Andoni+'08])



Answer:

Q: find quickly stocks like 'C' (or customers like 'smith')

A: summarize seq. to a few numbers/features (eg., avg, stdv, Fourier coeff.)



Books

- William H. Press, Saul A. Teukolsky, William T. Vetterling and Brian P. Flannery: *Numerical Recipes in C*, Cambridge University Press, 1992, 2nd Edition. (Great description, intuition and code for SVD)
- ★ • C. Faloutsos: *Searching Multimedia Databases by Content*, Kluwer Academic Press, 1996 (introduction to SVD, Fourier, Wavelets, and GEMINI)

References

- Gaede, V. and O. Guenther (1998). “Multidimensional Access Methods.” *Computing Surveys* 30(2): 170-231.
- Alexandr Andoni, Piotr Indyk (2008). “Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions”, *CACM* 51, 1, 2008
<https://doi.org/10.1145/1327452.1327494>

References

- Oppenheim, I. J., A. Jain, et al. (March 2002). A MEMS Ultrasonic Transducer for Resident Monitoring of Steel Structures. SPIE Smart Structures Conference SS05, San Diego.



Part 2:

DSP (Digital

Signal Processing)

Outline



- Motivation
- P1. Similarity Search and Indexing
 - P1.1. distance functions: Euclidean; Time-warping
 - P1.2. indexing
 - P1.3. feature extraction
- P2. DSP (Digital Signal Processing)
- ...

Outline



- Motivation
- P1. Similarity Search and Indexing
- ➔ • P2. DSP (Digital Signal Processing)
 - P2.1. Discrete Fourier Transform (DFT)
 - P2.2. Discrete Wavelet Transform (DWT)
- P3. Linear Forecasting
- P4. Non-linear forecasting
- Conclusions

Detailed Outline





- P1. Indexing etc
- P2. DSP (Digital Signal Processing)
 - P2.1. DFT
 - Definition of DFT and properties
 - how to read the DFT spectrum
 - P2.2. DWT
 - Definition of DWT and properties
 - how to read the DWT scalogram



Problem



Q: How to summarize / extract few features







Answer:

Q: How to summarize / extract few features

A: Fourier; Wavelets


(A': SVD Singular Value Decomposition,
ICA = Independent Component Analysis)



Introduction - Problem#2

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

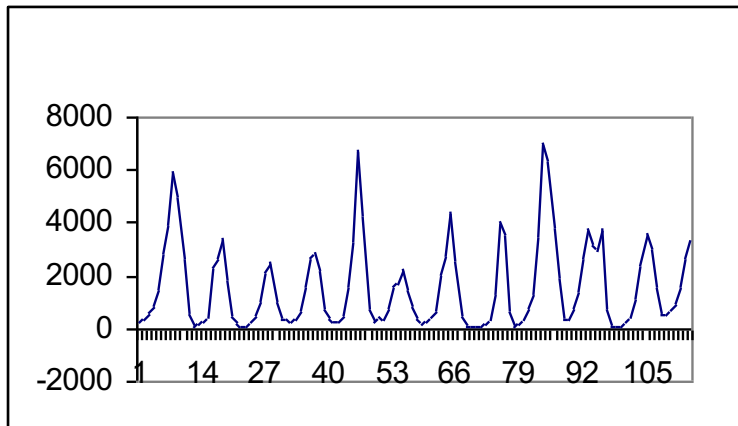
Find: patterns and/or compress



lynx caught per year
(packets per day;
automobiles per hour)

Problem #2: patterns?


Q: How to automatically find patterns?



Problem #2: patterns?




Q: How to automatically find patterns?

A: Discrete Fourier Transform (DFT)




DFT: finds sinusoids


Given a signal,
DFT decomposes it
to sinusoids



+



+



DFT: definition

- For a sequence x_0, x_1, \dots, x_{n-1}
- the (**n-point**) Discrete Fourier Transform is
- X_0, X_1, \dots, X_{n-1} :

$$X_f = 1/\sqrt{n} \sum_{t=0}^{n-1} x_t * \exp(-j2\pi tf/n) \quad f = 0, \dots, n-1$$

$$(j = \sqrt{-1})$$

$$x_t = 1/\sqrt{n} \sum_{f=0}^{n-1} X_f * \exp(+j2\pi tf/n)$$

inverse DFT




DFT: examples

flat


Amplitude: $A_f^2 = \text{Re}^2(X_f) + \text{Im}^2(X_f)$

Plot[Abs[Fourier[x]]];

Amplitude




time




freq

DFT: examples

Low frequency sinusoid




time




freq

DFT: examples

- Sinusoid - symmetry property: $X_f = X_{n-f}^*$




time




freq

DFT: examples

- Higher freq. sinusoid



time






freq


DFT: check-point




Spectrum??



+







+




DFT: check-point



Spectrum:



Ampl.




Freq.

DFT: Amplitude spectrum


Amplitude: $A_f^2 = \text{Re}^2(X_f) + \text{Im}^2(X_f)$

count



year


Ampl.



Freq.


DFT: Amplitude spectrum

count



year


Ampl.



Freq.


DFT: Amplitude spectrum

count



year



Ampl.



Freq.


DFT: Amplitude spectrum

- excellent approximation, with only 2 frequencies!
- so what?




DFT: Amplitude spectrum

- excellent approximation, with only 2 frequencies!
- so what?
- A1: **(lossy) compression**
- A2: pattern discovery




DFT: Amplitude spectrum

- excellent approximation, with only 2 frequencies!
- so what?
- A1: (lossy) compression
- A2: **pattern discovery**



DFT - Conclusions

- It spots periodicities (with the ‘**amplitude spectrum**’)
- can be quickly computed ($O(n \log n)$), thanks to the FFT algorithm.
- **standard** tool in signal processing (speech, image etc signals)
- (closely related to DCT and JPEG)



Detailed Outline



- P1. Indexing etc
- P2. DSP (Digital Signal Processing)
 - P2.1. DFT
 - Definition of DFT and properties
 - how to read the DFT spectrum
 - P2.2. DWT
 - Definition of DWT and properties
 - how to read the DWT scalogram




Problem #2:

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

Find: patterns, periodicities, and/or compress



count




lynx caught per year
(packets per day;
virus infections per month)

year

Wavelets - DWT

- DFT suffers on short-duration waves (eg., baritone, silence, soprano)

value




time

Wavelets - DWT

- DFT suffers on short-duration waves (eg., baritone, silence, soprano)

value




time

Wavelets - DWT

- DFT suffers on short-duration waves (eg., baritone, silence, soprano)

value




time

Wavelets - DWT

- DFT suffers on short-duration waves (eg., baritone, silence, soprano)

value




time

Wavelets - DWT

- DFT suffers on short-duration waves (eg., baritone, silence, soprano)

value




time

Wavelets - DWT

- DFT suffers on short-duration waves (eg., baritone, silence, soprano)


value



time

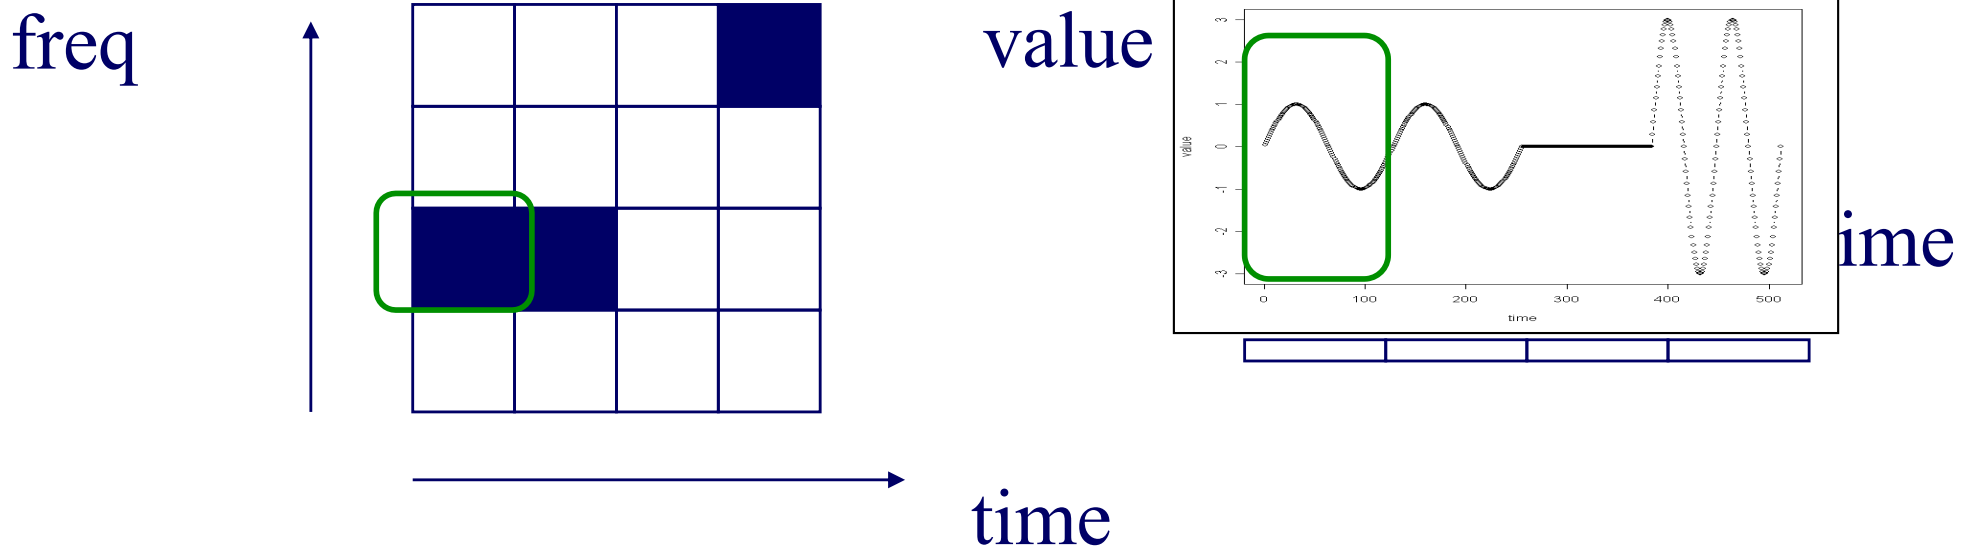
Short Window F.T.

- Solution#1: Short window Fourier transform (SWFT)



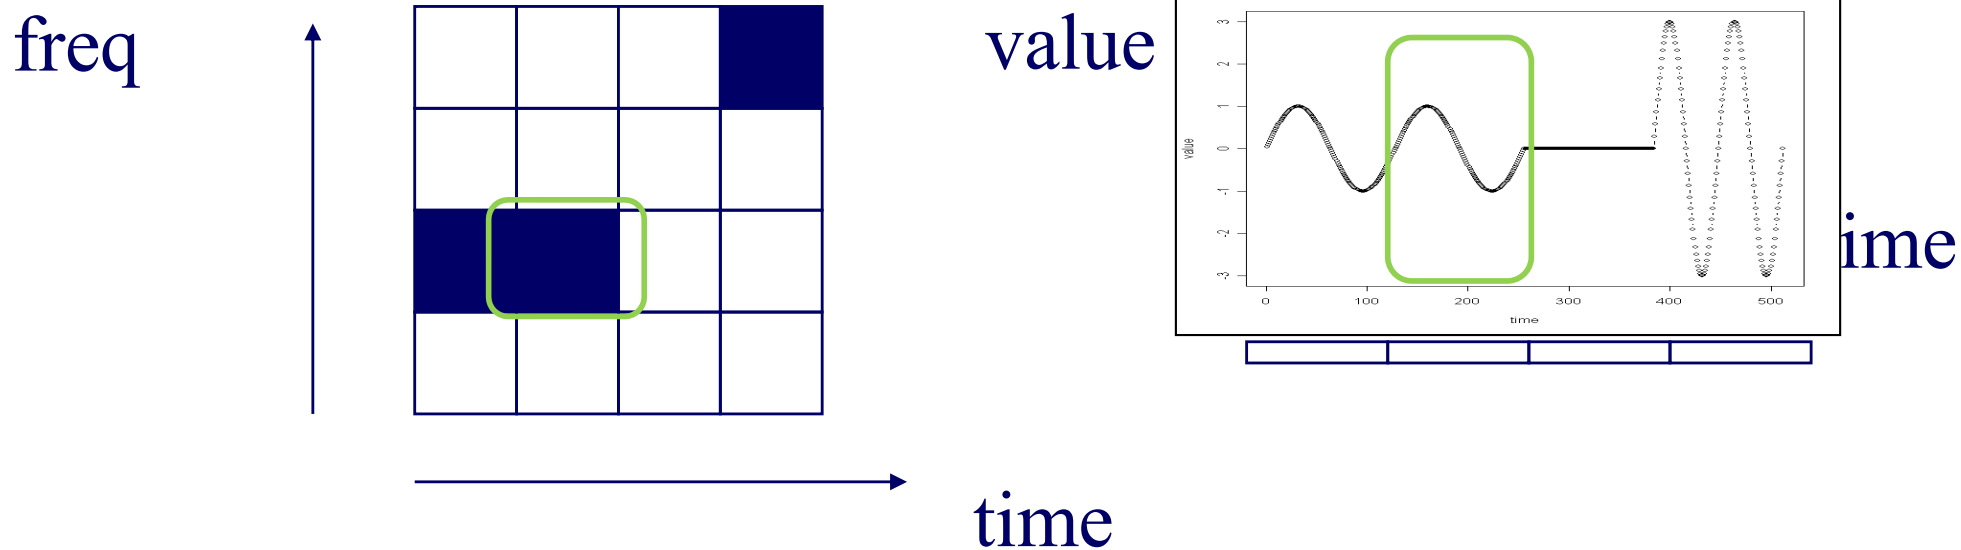
Short Window F.T.

- Solution#1: Short window Fourier transform (SWFT)
- But: how short should be the window?




Short Window F.T.

- Solution#1: Short window Fourier transform (SWFT)



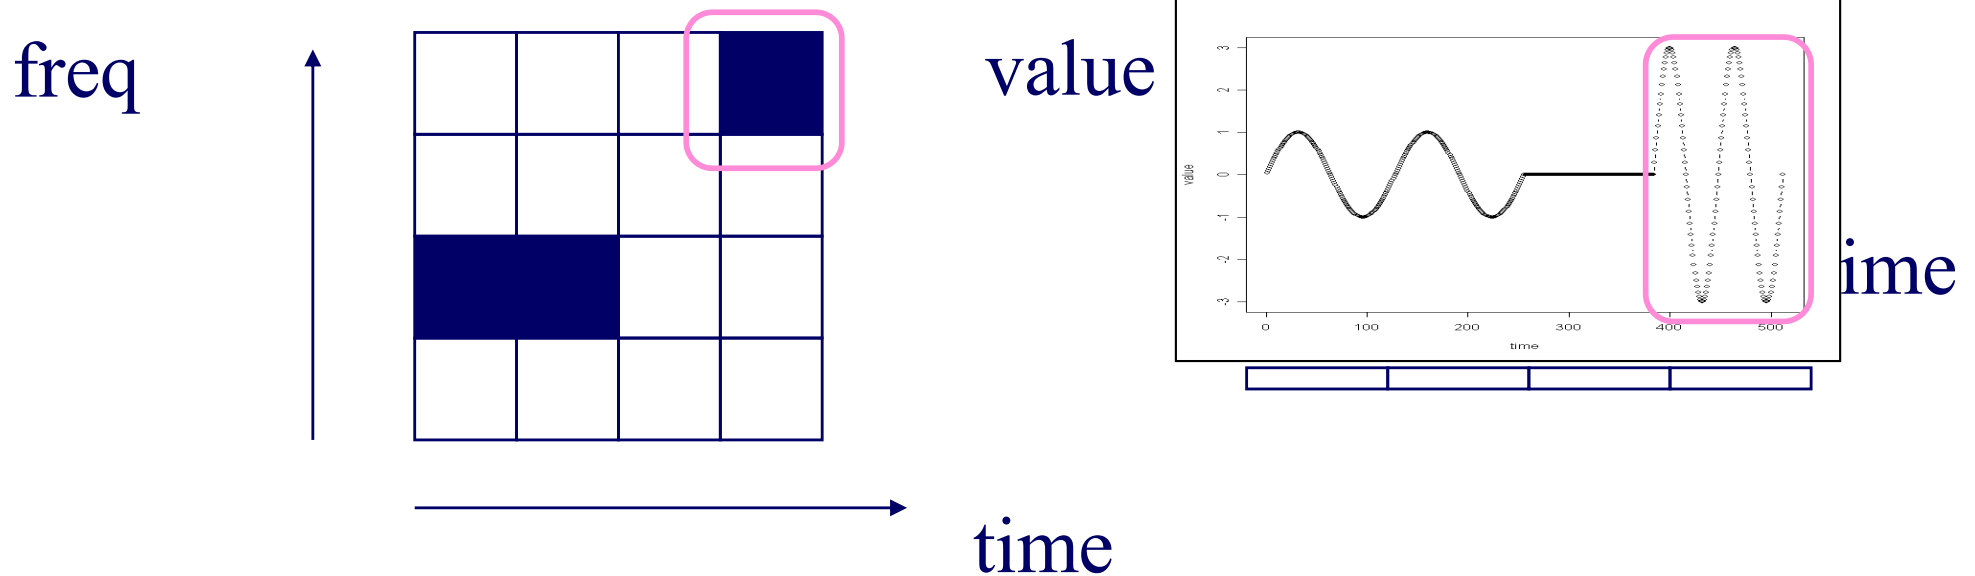
Short Window F.T.

- Solution#1: Short window Fourier transform (SWFT)



Short Window F.T.


- Solution#1: Short window Fourier transform (SWFT)





Short window FT: check-point

Chirp




$$x(t) = \sin (2 \pi t^2 / N)$$

(usual sinusoids: ↓)

$$x(t) = \sin (2 \pi f * t / N)$$


freq



time

Short window FT: check-point

Chirp




$$x(t) = \sin (2 \pi t * t / N)$$

(usual sinusoids: ↓)

$$x(t) = \sin (2 \pi f * t / N)$$

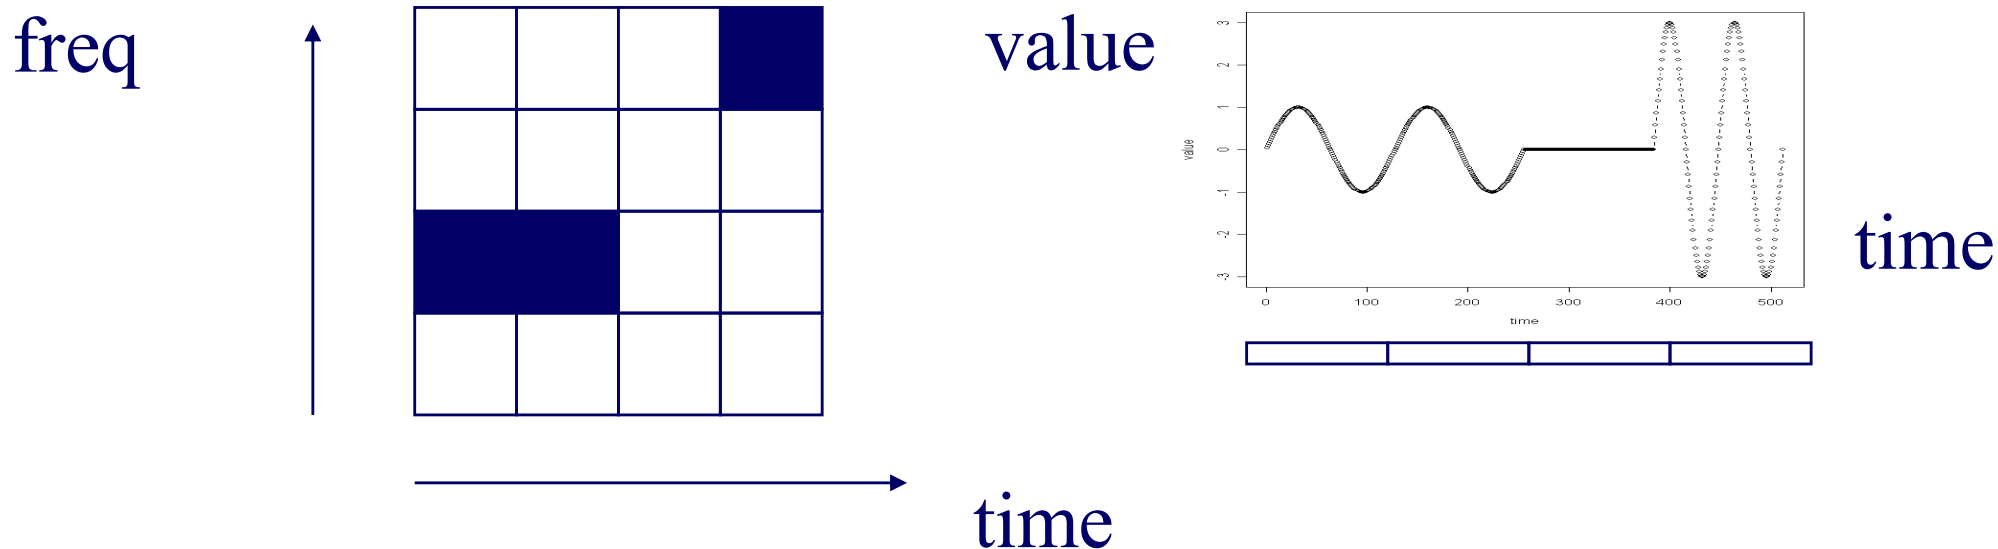
freq




time

Wavelets - DWT


- Solution#1: Short window Fourier transform (SWFT)
- **But: how short should be the window?**




Haar Wavelets




- Answer: **multiple** window sizes! -> DWT



Haar Wavelets



- Basis functions: waves (of length 1, $\frac{1}{2}$, $\frac{1}{4}$, etc)
- (vs DFT: sinusoids of full length)



Detailed Outline




- P1. Indexing etc
- P2. DSP (Digital Signal Processing)
 - P2.1. DFT
 - Definition of DFT and properties
 - how to read the DFT spectrum
 - P2.2. DWT
 - Definition of DWT and properties
 - how to read the DWT scalogram




Wavelets in action

- # of correction packets vs time




Case study: BGP updates




BGP-lens: Patterns and Anomalies in Internet Routing Updates B. Aditya Prakash et al, SIGKDD 2009

Case study: BGP updates




BGP-lens: Patterns and Anomalies in Internet Routing Updates B. Aditya Prakash et al, SIGKDD 2009

Case study: BGP updates



BGP-lens: Patterns and Anomalies in Internet Routing Updates B. Aditya Prakash et al, SIGKDD 2009

Case study: BGP updates



BGP-lens: Patterns and Anomalies in Internet Routing Updates B. Aditya Prakash et al, SIGKDD 2009


Case study: BGP updates

Given a time series: 1) PLOT IT
2) DFT / DWT

Case study: BGP updates

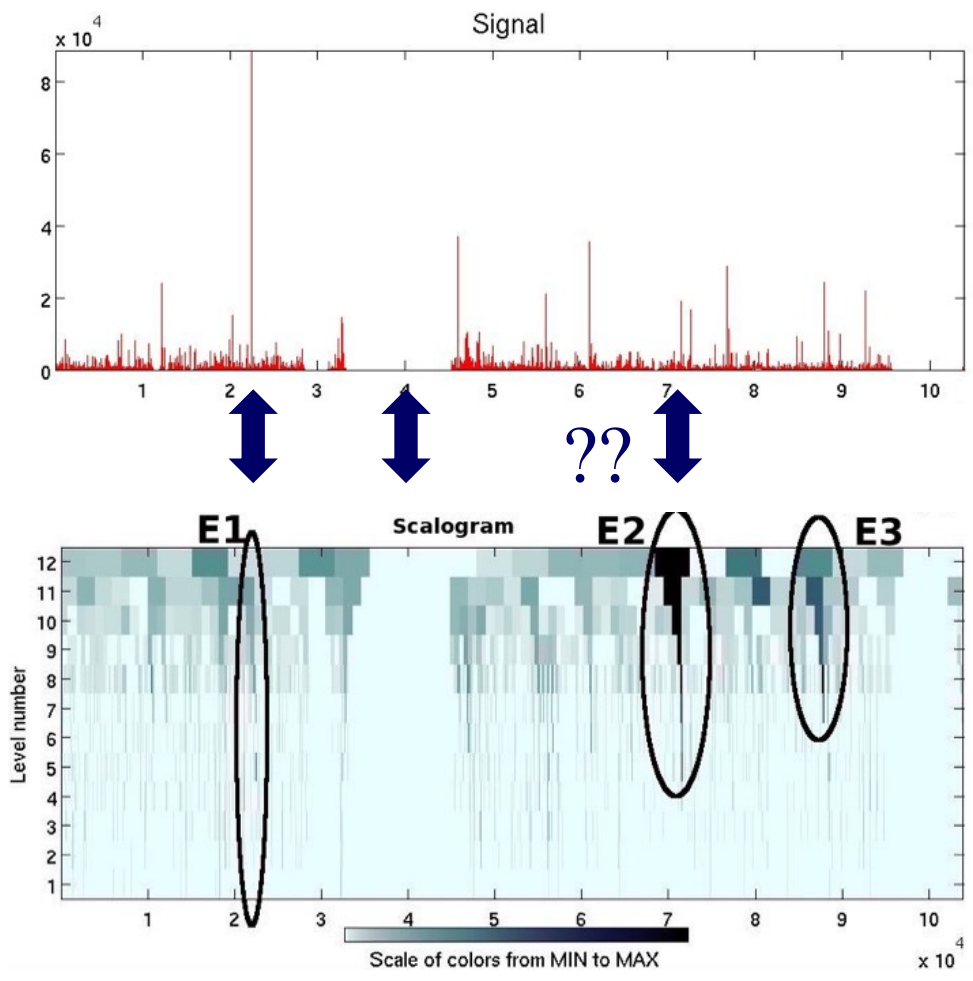
Low freq.:
omitted

freq.

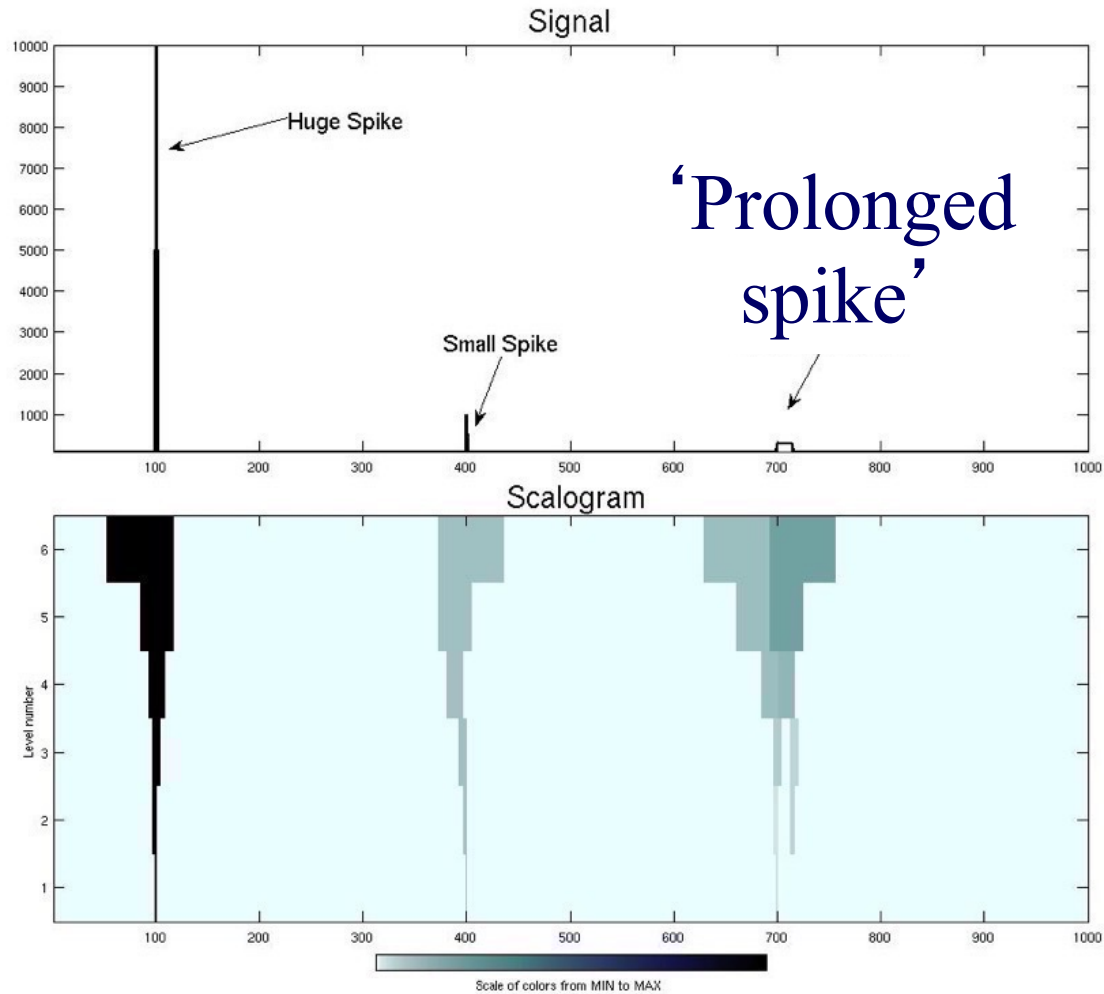


Case study: BGP updates

freq. ↓




Case study: BGP updates



freq. ↓


Case study: BGP updates

15K msgs, for several hours: 6pm-4am



Case study: BGP updates



15K msgs, for several hours: 6pm-4am



Conclusions for DSP





- DFT spots periodicities
- **DWT** : multi-resolution - matches processing of mammalian ear/eye better
- Both: powerful tools for **compression**, **pattern detection** in real signals



Conclusions for DSP



- DFT spots periodicities
- **DWT** : multi-resolution - match the resolution to the processing of the signal
- Plotting the results of the processing of a time series:
 - 1) PLOT IT
 - 2) DFT / DWT



Resources: software and urls

- *xwpl*: open source wavelet package from Yale, with excellent GUI
- *pywavelets* (python library)
- Included in matlab™

Books

- William H. Press, Saul A. Teukolsky, William T. Vetterling and Brian P. Flannery: *Numerical Recipes in C*, Cambridge University Press, 1992, 2nd Edition. (Great description, intuition and code for DFT, DWT)
- C. Faloutsos: *Searching Multimedia Databases by Content*, Kluwer Academic Press, 1996 (introduction to DFT, DWT)



Part 3:

Linear Forecasting

Outline



- Motivation
- P1. Similarity Search and Indexing
- P2. DSP (Digital Signal Processing)
- ➔ • P3. Linear Forecasting
- P4. Non-linear forecasting
- Conclusions

Forecasting

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


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





Problem#3: Forecast

- given $x_{t-1}, x_{t-2}, \dots,$
- Q: forecast x_t







Solution: AR(IMA)

- given x_{t-1}, x_{t-2}, \dots ,
- Q: forecast x_t
- A: AR(IMA) = Box-Jenkins, Kalman



Auto **R**egressive
Integrated
Moving **A**verage






Solution: AR(IMA)

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



Problem#3: Forecast


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



Forecasting: Preprocessing

MANUALLY:


remove trends



time

spot periodicities

7 days



time

Problem#3: Forecast


- Solution: try to express

x_t

as a linear function of the past: x_{t-2}, x_{t-2}, \dots ,
(up to a window of w)

Formally:

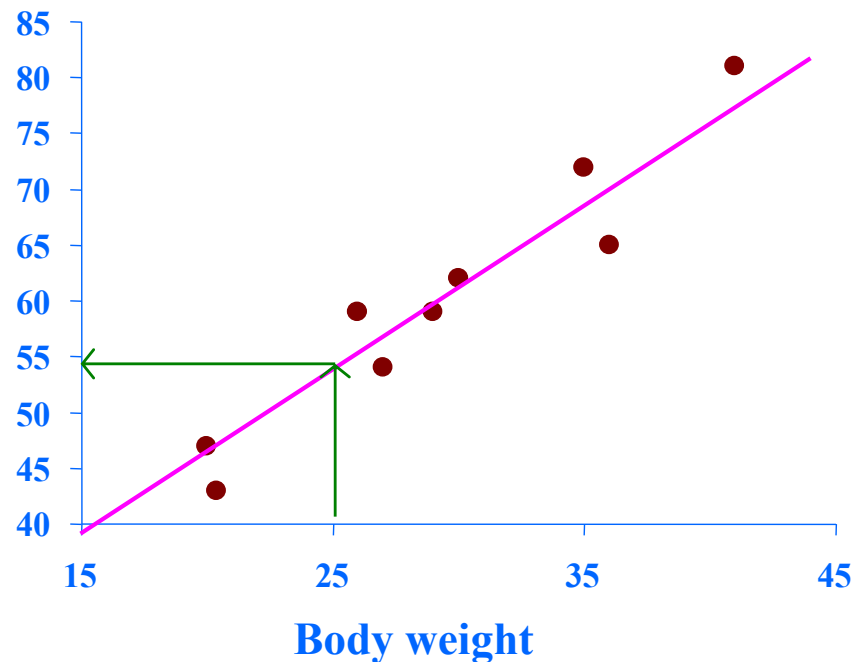
$$x_t \approx a_1 x_{t-1} + \dots + a_w x_{t-w} + \textit{noise}$$



Linear Regression: idea

<i>patient</i>	<i>weight</i>	<i>height</i>
1	27	43
2	43	54
3	54	72
...
N	25	??

Body height




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

Linear Auto Regression:

<i>Time</i>	<i>Packets Sent(t)</i>
1	43
2	54
3	72
...	...
N	??

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 ($S[t-1]$)

Deep Learning:


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 ($S[t-1]$)

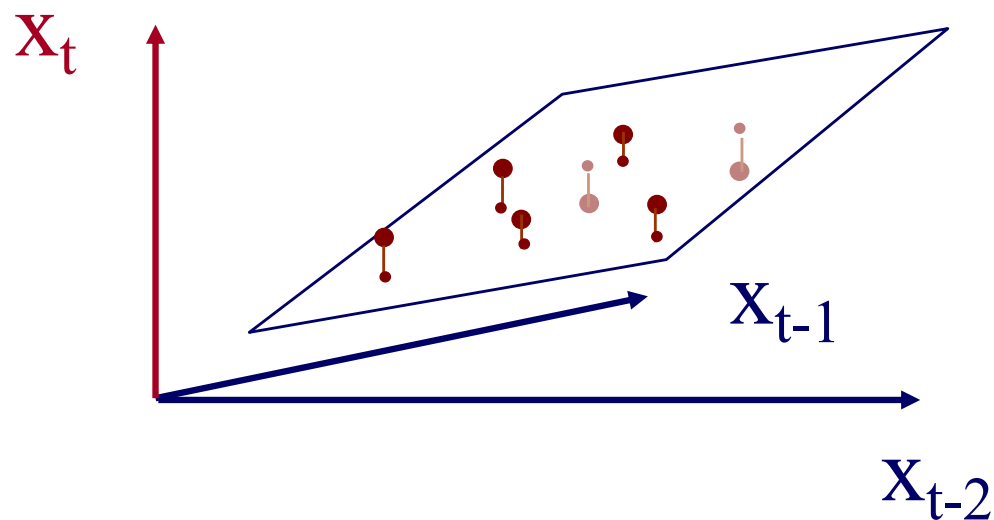
More details:

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




More details:

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



More details:

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



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**
 - \mathbf{a} is the vector of the regression coefficients
 - \mathbf{X} has the N values of the w indep. variables
 - \mathbf{y} has the N values of the dependent variable



Solution: AR(IMA)

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



Resources: software and urls

- <http://cran.r-project.org/> ('R' system)
- <https://www.statsmodels.org/> (python)
- ...<many more>

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.

Additional Reading

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



Part 4: chaos and non-linear forecasting

Outline



- Motivation
- P1. Similarity Search and Indexing
- P2. DSP (Digital Signal Processing)
- P3. Linear Forecasting
- ➔ • P4. Non-linear forecasting
- Conclusions

Detailed Outline




- ...
- P3. Linear forecasting
- P4. Non-linear forecasting
 - Problem
 - Idea
 - How-to
 - Experiments





Problem: Forecast



- given x_{t-1}, x_{t-2}, \dots , ('chaotic'/non-linear)
- Q: forecast x_t






Solution

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



Recall: Problem #3




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




How to forecast?

- what could go wrong with ARIMA?




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






How to forecast?

- what could go wrong with ARIMA?
- A: **linearity** assumption might not hold!



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



ARIMA pitfall

Example: logistic parabola

Models population of flies [R. May/1976]




#flies(t+1)

Reproductive
rate


Max# flies

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

Time-series plot



x_{t+1}



x_t 139

ARIMA pitfall

Example: logistic parabola


Models population of flies [R. May/1976]




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

Logistic
map

Time-series plot



x_{t+1}



x_t 140

ARIMA pitfall

Example: logistic parabola


Models population of flies [R. May/1976]



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


Logistic
map

- = SI virus prop. model (covid: S.I.R.)
- ~ Bass equation (market penetration) + 1
- Special case of Lotka-Volterra
 - Competing species / products



ARIMA pitfall

Linear equations, e.g., AR, ARIMA, ...




ARIMA pitfall

Linear equations, e.g., AR, ARIMA, ...

e.g., AR(1)

$$x_{t+1} = ax_t + \epsilon$$




ARIMA pitfall

Linear equations, e.g., AR, ARIMA, ...

e.g., AR(1)

$$x_{t+1} = ax_t + \epsilon$$




Solution?




“Delayed Coordinate Embedding” = Lag Plots
[Sauer94]

k-nearest neighbor search



General Intuition (Lag Plot)



Q: How to interpolate?


How do we interpolate between the k nearest neighbors?

A1: Average


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

Q: How to interpolate?


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



Detailed Outline

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


Datasets



Logistic Parabola:

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

Models population of flies [R. May/1976]




Lag-plot

Logistic Parabola

Our Prediction from here

Value





Timesteps

Logistic Parabola

Comparison of prediction to correct values

Value



Timesteps



Edward Lorenz

Datasets



LORENZ: Models convection currents in the air

$$dx / dt = a (y - x)$$

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



$$dz / dt = xy - c z$$

- **Deterministic chaos**
- **‘butterfly effect’**



LORENZ

Comparison of prediction to correct values




Value

Timesteps

Value

Datasets

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




Time


Laser

Comparison of prediction to correct values

Value





Timesteps






Solution

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



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.

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.)

Almost Done

- Take-home messages:
- M1: equivalence of research problems
- M2: superset, 3h tutorial
- M3: summary of ‘recipes’




M1: 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)



M2: Extended, 3 hour tutorial, KDD17

<https://www.dm.sanken.osaka-u.ac.jp/~yasuko/TALKS/17-KDD-tut/>



Prof. Yasushi Sakurai





Prof. Yasuko Matsubara



M3: Time series - Answers



- P1. Similarity search: **Euclidean/time-warping; feature extraction and SAMs**
- P2. Periodicities: **DFT/DWT**
- P3. Linear Forecasting: **AR (Box-Jenkins)**
- P4. Non-linear forecasting: **lag-plots**



Time series - Answers



- P1. Similarity search



Thank you!
ありがとうございました
christos@cs.cmu.edu

