Discussion of *Visualization of complex seasonal patterns in time series* by Prof. Hyndman

Matteo Pelagatti

Department of Economics, Management and Statistics
University of Milano-Bicocca

General consideration about Prof. Rob Hyndman's work

"Statistics as one of the methods to solve real world problems" (from Ilaria Prosdocimi's slides, yesterday)

- RH solves real problems through statistics.
- RH's approach: first formalize the problem in a very general setting and then propose an optimal solution (so was today's presentation, but same approach for *forecasting reconciliation*, for example).
- RH makes research and teaching very *open*: check his book on forecasting `https://otexts.com/fpp3/` and his website. `https://robjhyndman.com/` to get an idea.
- RH creates great R packages (co-authored some 40 pkgs).

The forecast pacakge

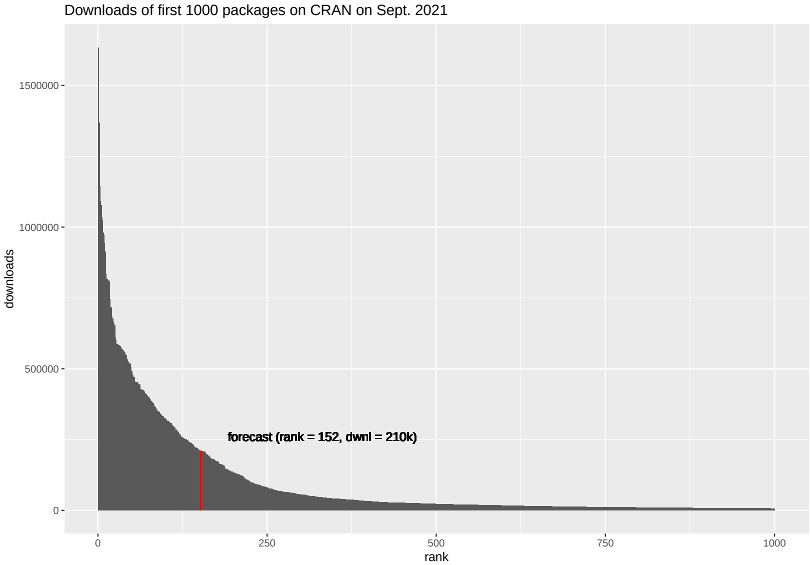

The papers on which today's presentation is based

1. "Visualizing probability distributions across bivariate cyclic temporal granularities" (J. Comp. and Graph. Stat.). DEFINITIONS AND MAPPING INTO PLOTS
2. "Detecting distributional differences between temporal granularities for exploratory time series analysis" (WP). AUTOMOTIC DETECTION OF INTERESTING PLOTS
3. "STR: Seasonal-trend decomposition using regression" (Informs J. Data Science). NEW METHOD TO DECOMPOSE TIME SERIES WITH COMPLEX SEASONAL PATTERNS
4. "MSTL: A seasonal-trend decomposition algorithm for time series with multiple seasonal patterns" (Int. J. Op. Research). ALTERNATIVE TO STR

My discussions follows the ordering of the above papers.

1. Visualizing probability distributions ...

- This paper sets the playing field.
- Defines for discrete-time time series
  - Time domain
  - Index set
  - Linear granularity
  - Ordering and grouping of granularities
  - Periodic granularities
  - Order of a granularity
  - Circular, quasi-circular and aperiodic granularities
  - Clashes and harmonies of cyclic granularities
  - Mapping of granularities in the *grammar of graphics*

Not much for me to discuss

2. Detecting distributional differences between temporal granularities ...

- Given a time series, you want to find out what granularities are interesting for data exploration through graphical representation.
- You want to find the grouping of data induced by granularities that produces the most dissimilar distributions in order to produce informative plots.
- The method
  - use *Jensen-Shannon distance* between two distributions
  - within and between groups synthesis of all bivariate comparisons of distributions
  - give more weight to within group distances than to between group (1/3 vs 2/3)
  - neutralize the effect of the number of comparisons (three alternatives)
  - use permutation to determine a threshold (similar to a critical value) for interesting harmonies to be considered

Simulated example of ranked plots

From less to more interesting plots

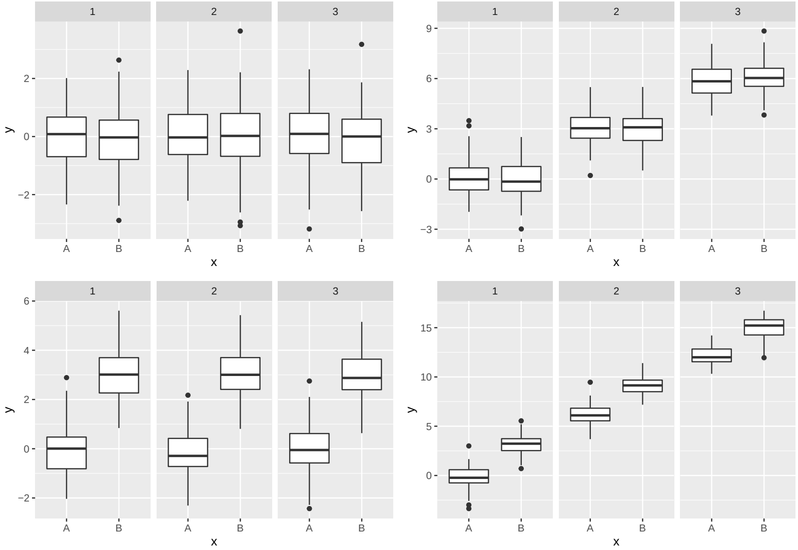

My comment

You want to find dissimilar distributions among a set of predefined groupings of the data.

What is the reason to design your procedure instead of just using distributional equality tests and their p-values for selecting and ranking?

For example, using the Anderson-Darling test on the data used in the above plots, we get the following p-values and strength $= -\log(\text{p-value})$ transforms, for ranking the grouping.

	p-value (w)	p-value (b)	strength (w)	strength (b)
plot 1	0.361	0.560	1.0	0.6
plot 2	0.088	0.000	2.4	204.8
plot 3	0.000	0.422	95.5	0.9
plot 4	0.000	0.000	93.7	214.7

3. STR: Seasonal-trend decomposition ...

- Non-parametric way to extract trend, seasonal components, ... based on smoothness-penalized likelihood.
- Given the penalty coefficients, you show how to estimate parameters and signals by least-squares.
- Penalization constants determined by grid-search and CV.

My comment

- Your method can be shown to be the optimal solution to a **parametric estimation problem**.
- We can find **stochastic processes for the trend and seasonality** such that your procedure produces the minimum MSE estimate of the signals based on the whole time series.
- Stating those stochastic components in **state-space** form you can use
  - the Kalman filter to **estimate the penalization constants** (instead of using CV),
  - the **smoother** to carry out inference for the components.

Stochastic process for the trend

The trend component you are using is the *smooth trend* or *integrated random walk* commonly used in **UCM**.

$$\mu_{t+1} = \mu_t + \beta_t$$
$$\beta_{t+1} = \beta_t + \mathrm{NID}(0, \sigma^2)$$

Stochastic process for the seasonal component

- Your specification is **different from any form used in the UCM** literature (see Proietti 2000 IJF).
- You work on **second differences**, while seasonal components are generally specified using first difference.
- You introduce **smoothness w.r. to time and w.r. to seasonal variation**.
- While the former is common in all UCM seasonal models, the latter is generally obtained using **low-frequency harmonics** of stochastic sinusoids.
- I will definitely try to find the process for your seasonal component and cast it in state-space form.
- Why preferring parametric UCM?
  - Missing observations are not a problem
  - Forecasting and signal extraction coherent
  - No cross-validation needed

Empirical comparison with UCM

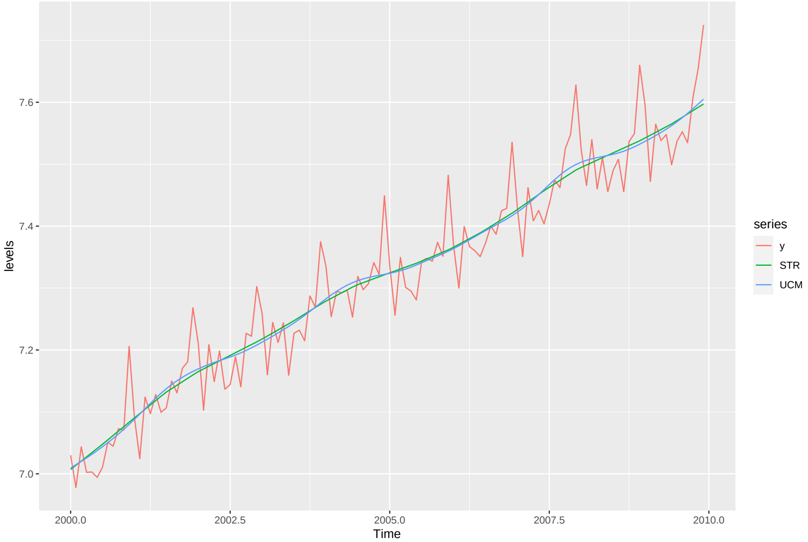

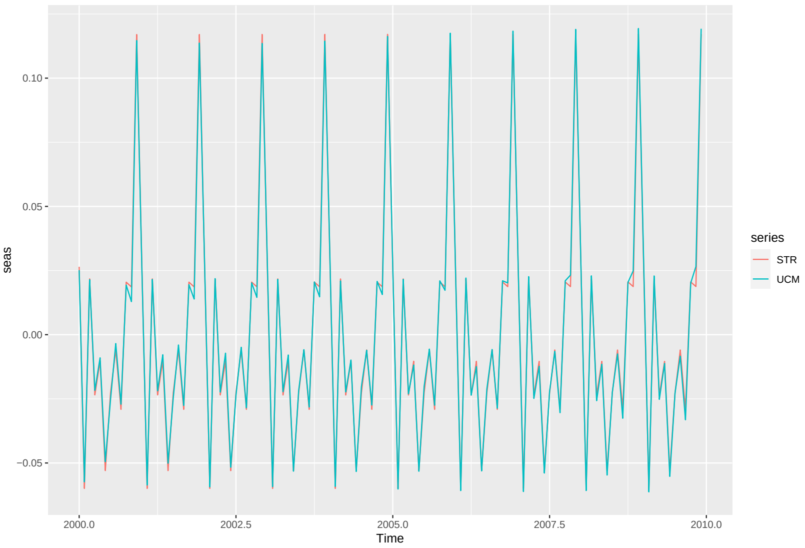
seas
0.10
0.05
0.00
−0.05
2000.0
2002.5
2005.0
2007.5
2010.0
Time
series
STR
UCM

To do list for modelling seasonality

- Automatic identification of seasonal cycles
- Seasonal patterns that change as functions of other seasonal patterns: for example, consider hourly data (as electricity demand)
  - the shape of 24-hour cycle changes (abruptly) according to weekday
  - the shape of the 24-hour (and 168-hour) cycle depends (smoothly) on the time of the year