Pattern Localization in Time Series through Signal-To-Model Alignment in Latent Space

UNIVERSIDAD DE CANTABRIA

Steven Van Vaerenbergh*, Ignacio Santamaría*, Víctor Elvira†, Matteo Salvatori∗

* University of Cantabria, Santander, Spain; † IMT Lille Douai CRISTAL (UMR 9189), Lille, France; ∗ Tecnatom S.A., Madrid, Spain

IMT Lille Douai
École Mines-Télécom
IMT-Université de Lille

Problem

We are given a time series that contains a sequence of patterns, and a theoretical model (the "blueprint") that lists the expected locations of these patterns (start/end in mm, seconds, ...).

The localization problem consists in estimating the precise (start and end) locations of each of the patterns in the time series.

- Example 1: Audio-to-score alignment [1].
- Example 2: Aligning a sensor measurement of a 1D structure to its known blueprint (e.g. in industrial environments). Illustration:

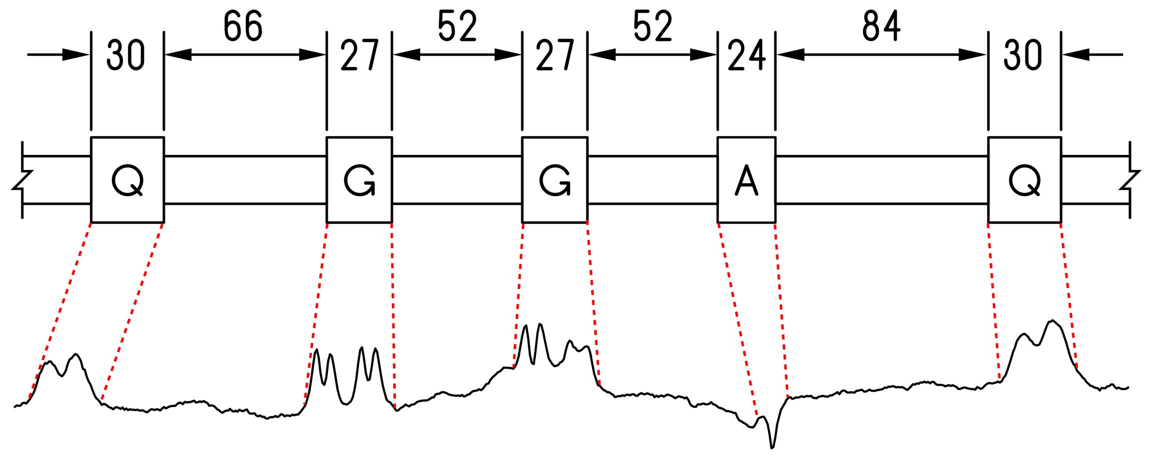

- In the literature, this localization problem is solved through alignment, typically by using Dynamic Time Warping (see bottom box).
- But alignment techniques require 2 time series → the second time series is obtained by *synthesis* based on the model.
- An accurate synthesis technique is key! → requires *domain knowledge*.
- True patterns may include *variations* that worsen alignment (e.g. "G").

Main idea

Use supervised machine learning to improve the alignment by learning transformations for the true and the synthesized time series into a space in which they are more similar.

- The optimal transformation is learned by Canonical Correlation Analysis (CCA) (see box to the right).
- CCA can compensate for the shortcomings of a generic synthesis.

Dynamic Time Warping

Dynamic Time Warping (DTW) is the de-facto standard technique for aligning time series [2]. Given two time series, DTW seeks the warping path that optimally aligns them.

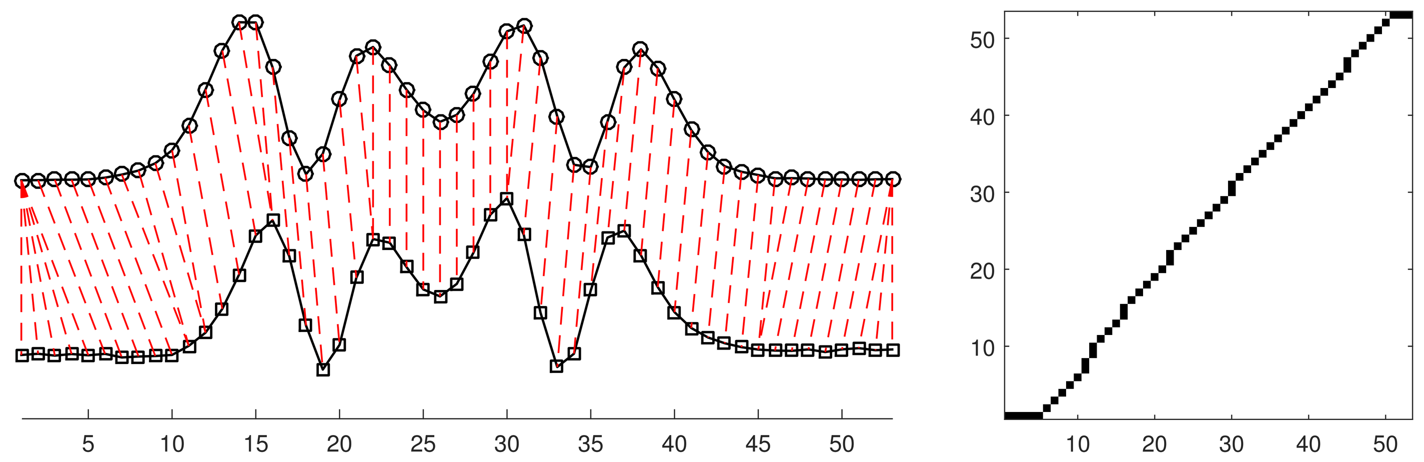

Proposed Machine Learning Framework

Training stage:
Input: A time series and the corresponding *aligned* model (labeled by a domain expert).

1. The model is synthesized into a time series;
2. Both time series may optionally undergo a fixed transformation;
3. CCA is performed.

Output: The coefficients of the 2 CCA transformations.

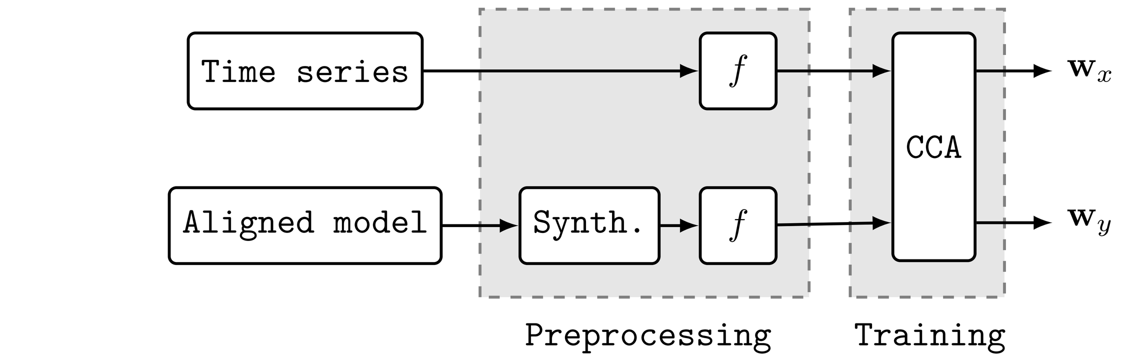

Testing stage:
Input: A time series and the corresponding *unaligned* model.

1. Same preprocessing as in the training stage;
2. Both time series are mapped into the latent space.
3. DTW is applied.

Output: Alignment solution for the time series and the model.

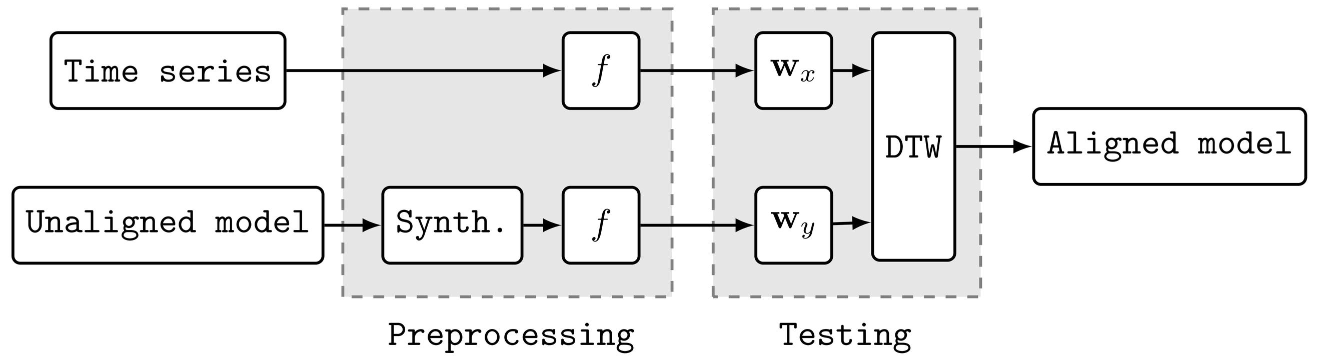

Canonical Correlation Analysis

Given two multidimensional random variables x and y, canonical correlation analysis (CCA) seeks a pair of optimal linear transformations such that the transformed variables are maximally correlated [3].

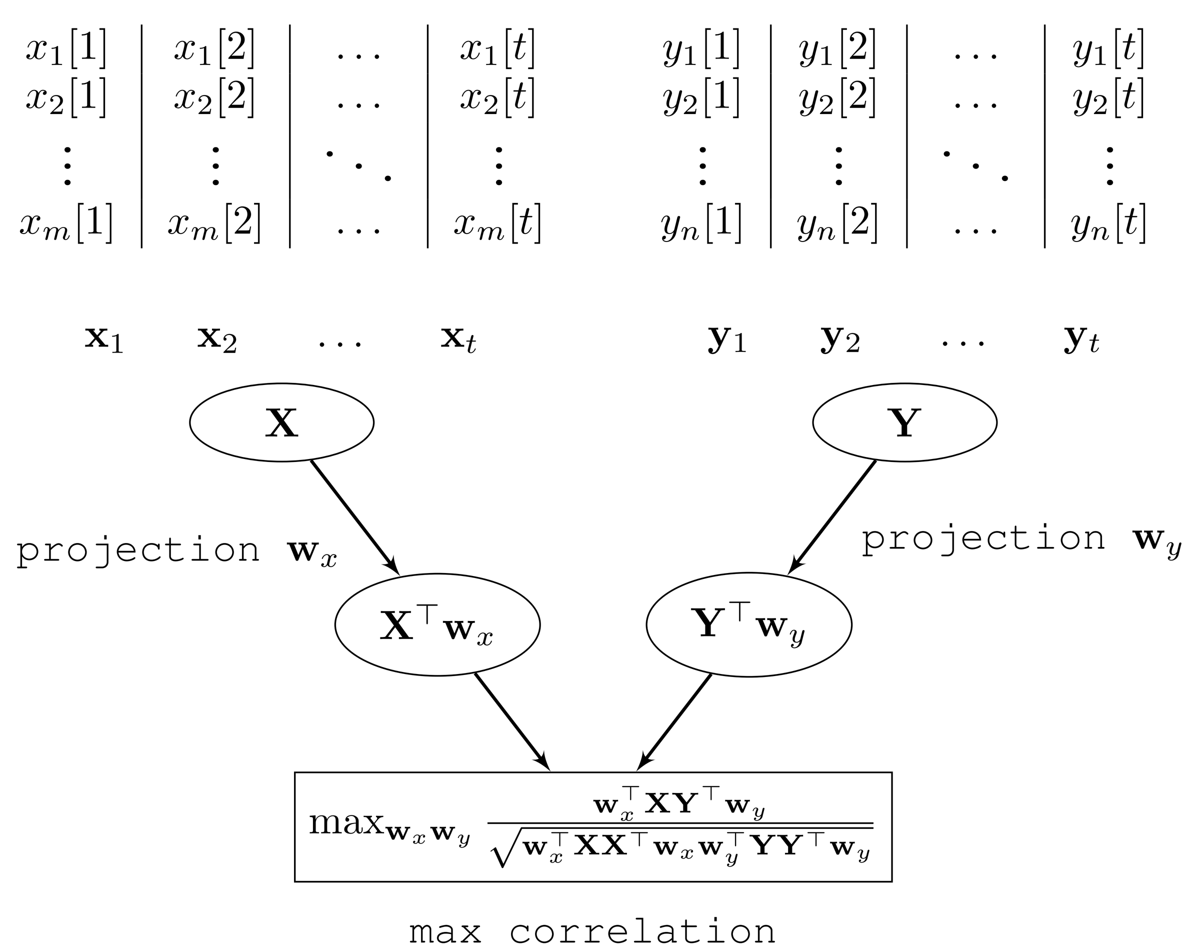

Experiments

Pattern localization for non-destructive testing of heat generator tubes in nuclear power plants (cf. initial illustration). Blueprint (model) is available.

Experiment 1: Training with a single time series (7 patterns).
Synthesis through replication.
Function f: time-embedding of 20 past and 20 future samples.

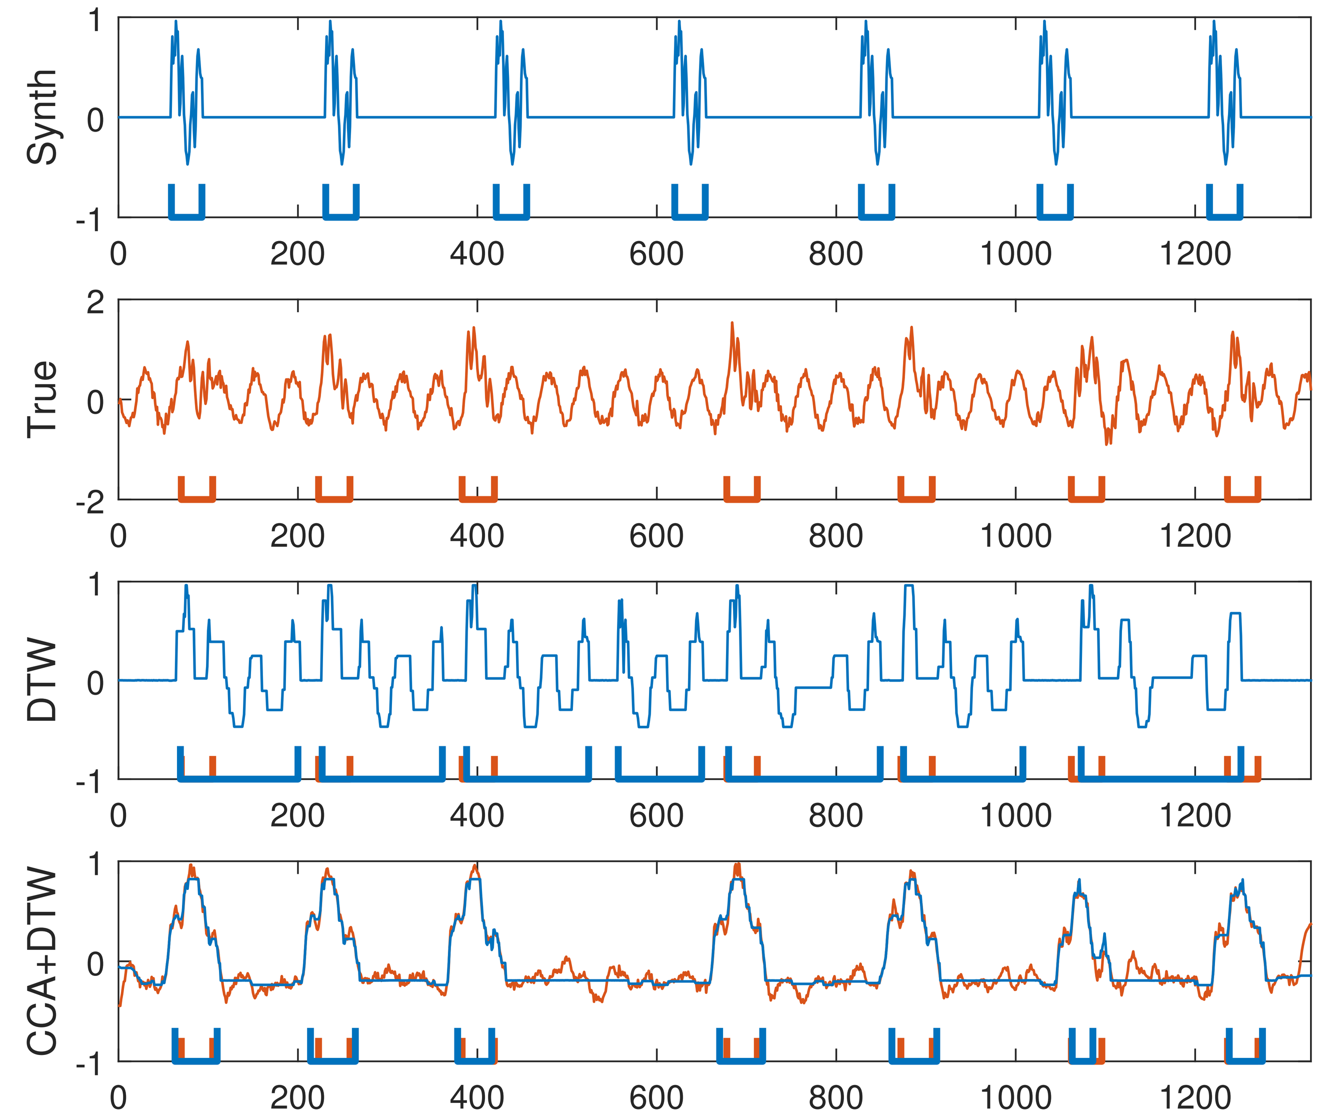

Experiment 2: Influence of different rates of pilgrim noise (simulated).

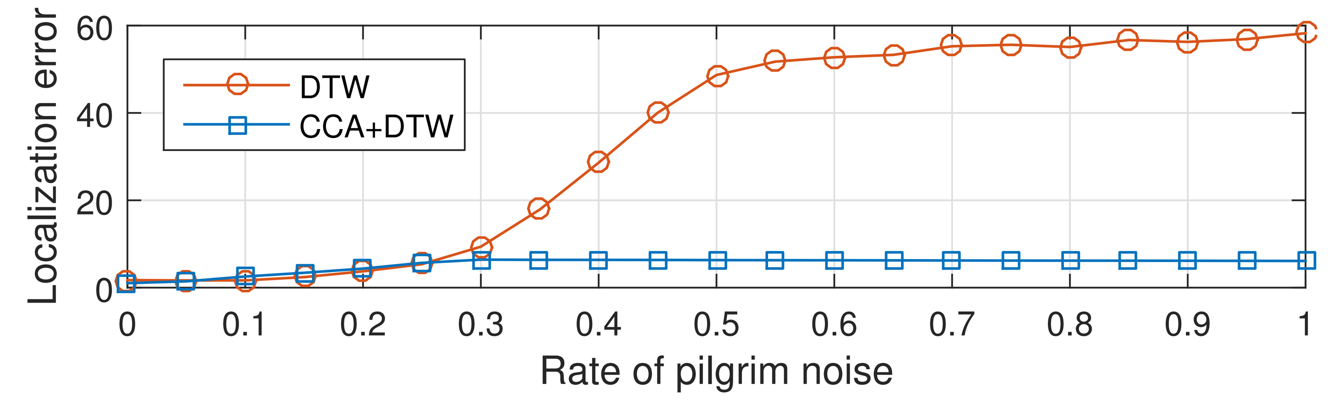

Conclusions

- Pattern localization in a time series for which a model is available.
- True and synthesized time series are mapped into an optimal latent space (learned by CCA), then aligned by DTW.
- Results are better and more robust compared to DTW only.
- Uses little domain-specific knowledge – applicable in many contexts.

References

[1] John Thickstun, Zaid Harchaoui, and Sham Kakade. Learning features of music from scratch. In *ICLR*, 2017.

[2] Taras K. Vintsyuk. Speech discrimination by dynamic programming. *Cybernetics*, 4(1):52–57, 1968.

[3] Harold Hotelling. Relations between two sets of variates. *Biometrika*, 28(3/4):321–377, 1936.

2018 IEEE International Conference on Acoustics, Speech and Signal Processing, 15-20 April 2018, Calgary, Alberta, Canada