FAKULTÄT FÜR !NFORMATIK
Faculty of Informatics

FWF Der Wissenschaftsfonds.

electronic commerce

RUCON

Consistency of the Fittest:
Towards Dynamic Staleness Control for Edge Data Analytics

Atakan Aral, Ivona Brandić
Institute of Information Systems Engineering
Vienna University of Technology, Vienna, Austria

{atakan.aral, ivona.brandic}@tuwien.ac.at
http://rucon.ec.tuwien.ac.at/

A Paradigm Shift from Core Analytics to Edge Analytics

Two main driving forces:

- Faster response to detected events
  - Edge nodes are in close proximity.
  - Network latency is low.
- Less data transfer / storage
  - Edge nodes can filter out raw data.
  - Only valuable information is transferred.

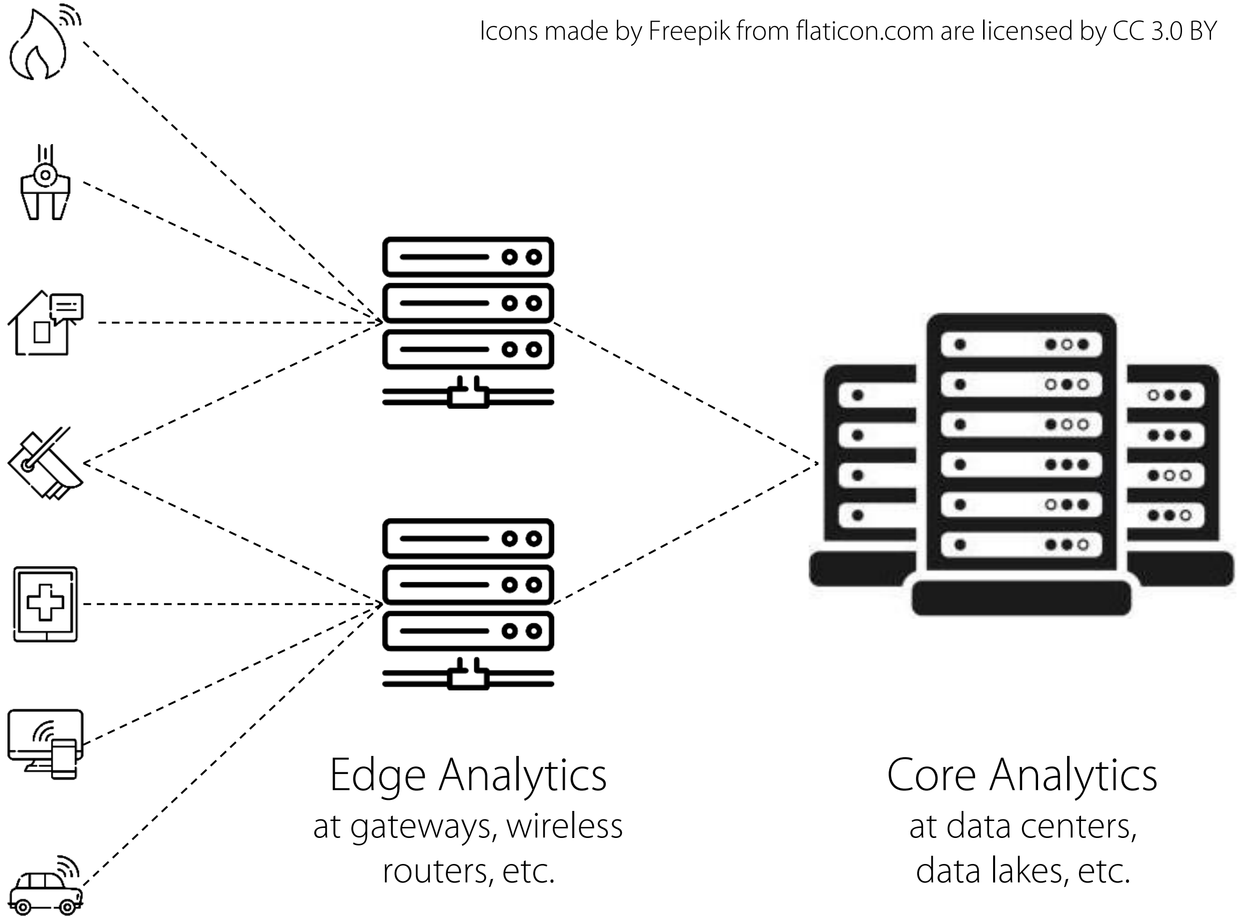

Typical Applications

Common Features:

- Time-sensitive
- Data-intensive
- Distributed
- **Non-stationary**

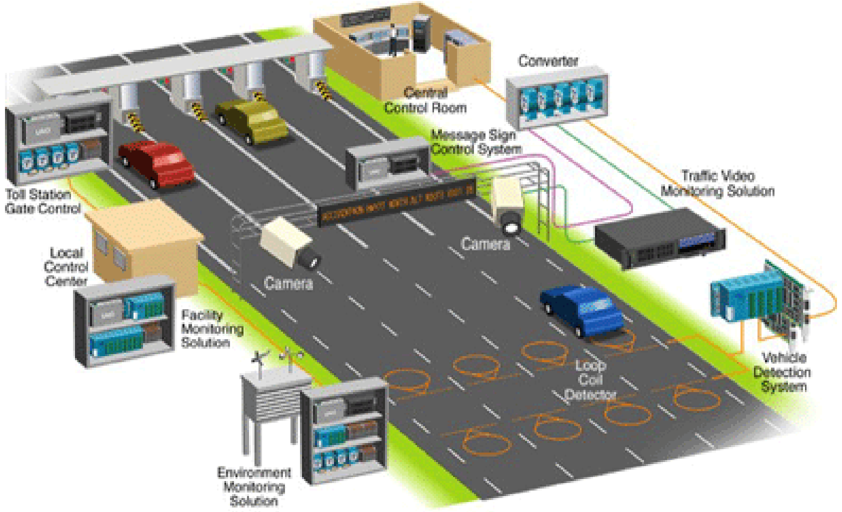

Intelligent Traffic Control

Spam or Fraud Detection

Computational Advertising

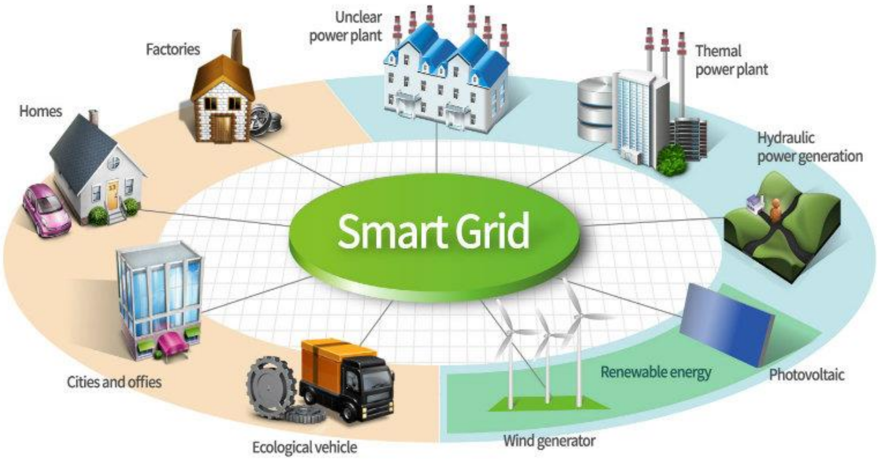

Transactive Energy Control

i. Data from distributed streams are collected in nearby edge nodes.

ii. Decisions based on local ML model and current data are sent back to the actuators **[inference]**.

iii. Local model is updated online with current data **[training]** and sent to a parameter server.

iv. A global model is formed and it replaces stale or incomplete models at edge nodes.

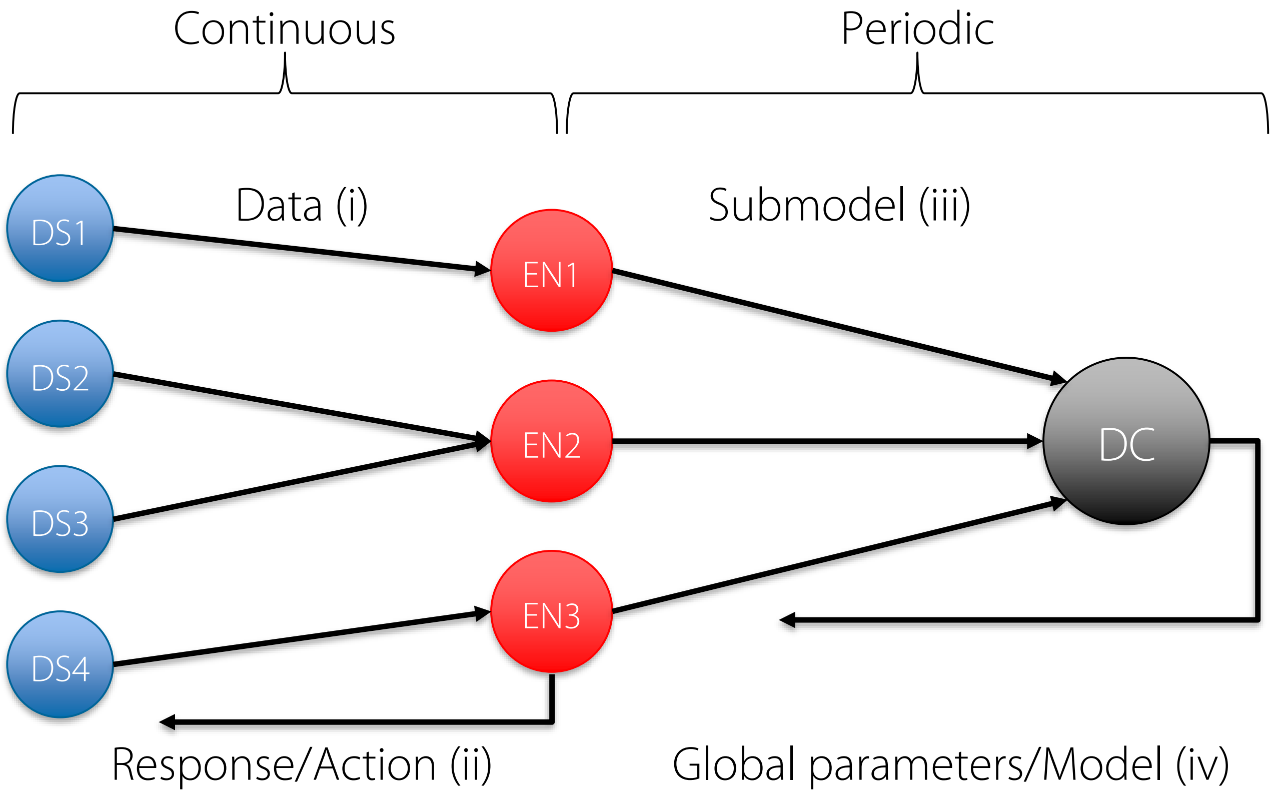

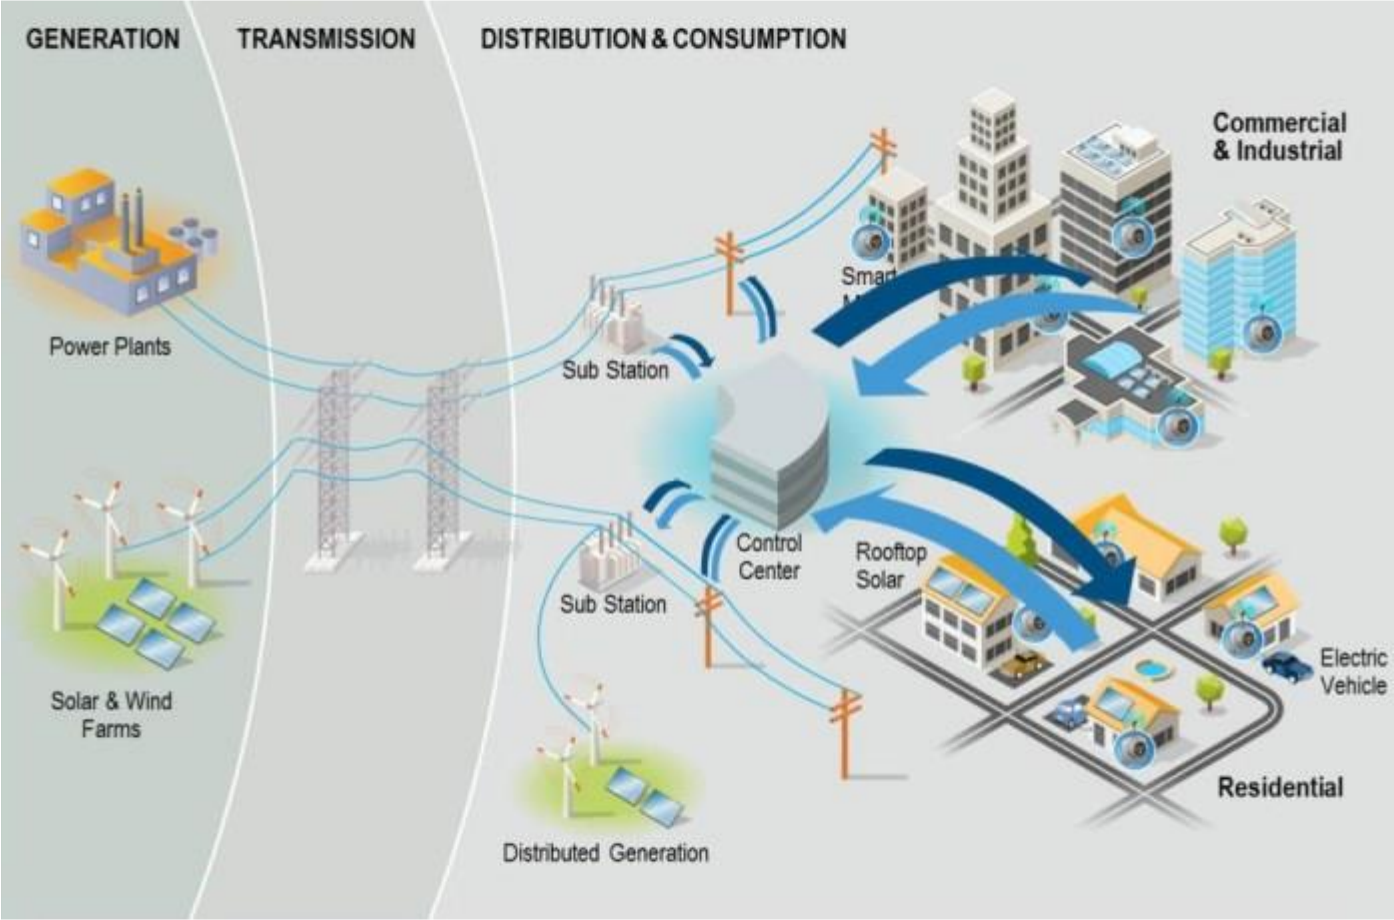

Smart Grid architecture. Image source: ecoideaz.com

- *"Dynamic balance of supply and demand across the entire electrical infrastructure"*
- Supply/Demand forecasting is a critical issue
- DOML & edge is a natural fit
  - coping with spikes in demand,
  - dynamically controlling voltage to reduce losses,
  - increasing the utilization of generators.

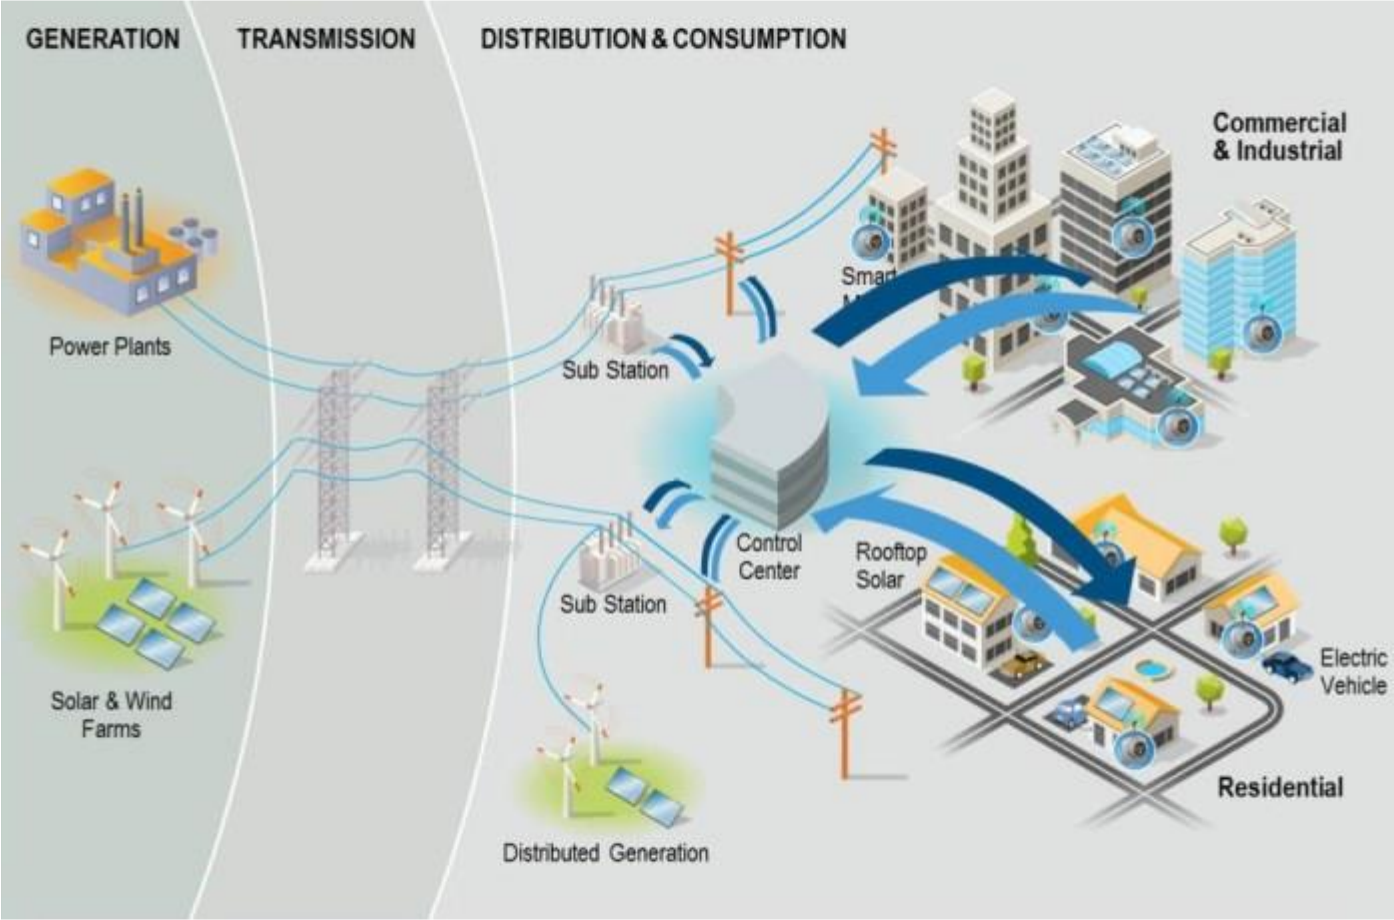

Smart Grid architecture. Image source: ecoideaz.com

- Non-stationarity arises from:
  - **Structural changes:** joining or leaving producers, forming or dissolving links, technological advances, etc.
  - **Quantitative changes:** evolving consumption habits, unexpected events, seasonality, etc.
- Updates from edge nodes arrive **asynchronously**.
- Not all updates have the same **significance**.

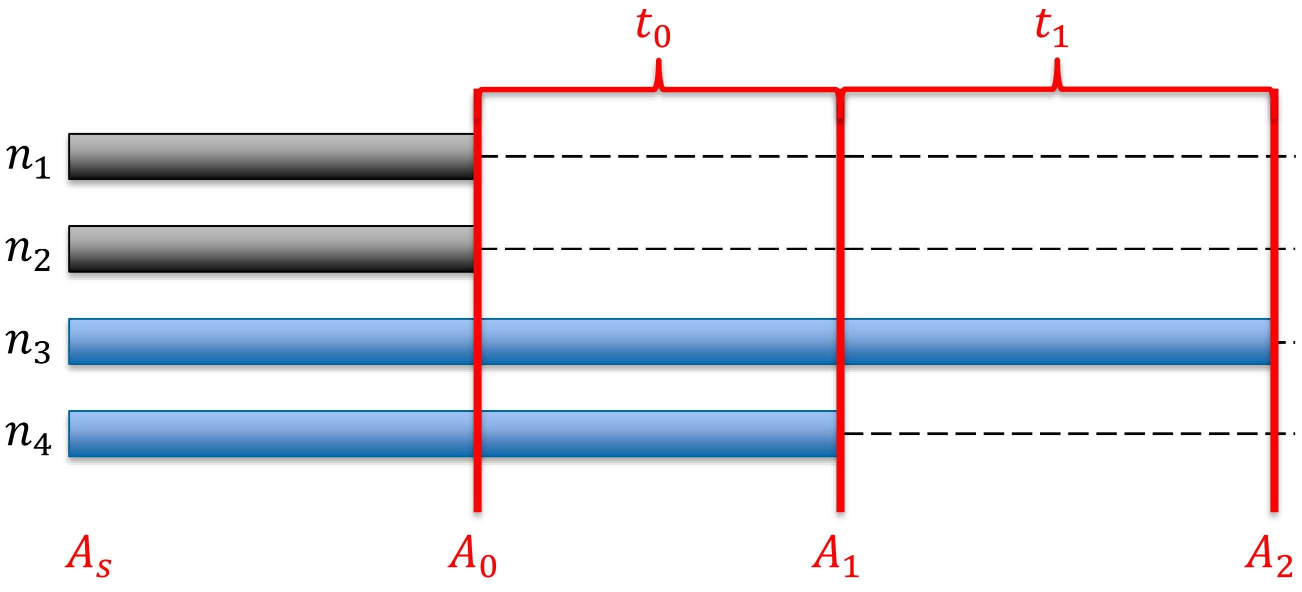

How many responses to wait for before updating global model and sending it to edge nodes?

- Quorum size
- Staleness bound
- Dynamic periodicity

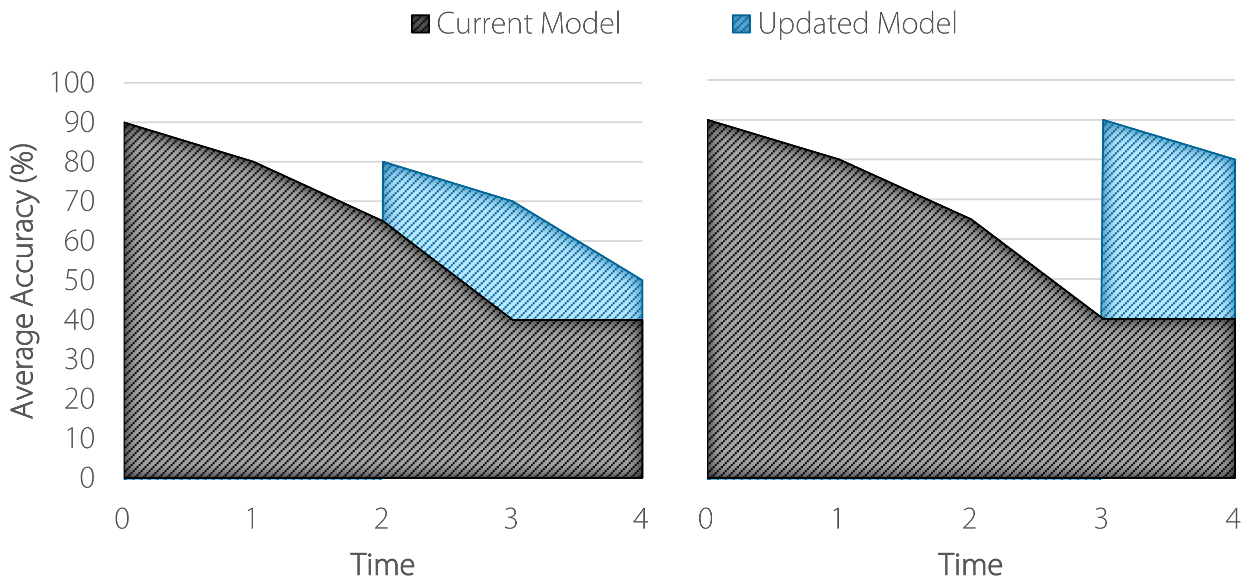

$$
\begin{aligned}
\underset{i}{\text{maximize}} \quad & \int_0^{\tau_i} A_s(x)\,dx + \int_{\tau_i}^{\epsilon} A_i(x)\,dx \\
\text{where} \quad & \tau_i = \sum_{j=0}^{i-1} t_j \\
\text{subject to} \quad & i \in \mathbb{Z},\ 0 \le i \le k.
\end{aligned}
$$

$$A(x) = C(x)A$$

$$\begin{aligned}
\underset{i}{\text{maximize}} \quad & \int_0^{\tau_i} A_s(x)\,dx + \int_{\tau_i}^{\epsilon} A_i(x)\,dx \\
\text{where} \quad & \tau_i = \sum_{j=0}^{i-1} t_j \\
\text{subject to} \quad & i \in \mathbb{Z},\ 0 \leq i \leq k.
\end{aligned}$$

Characteristic Function

- Depends on ML algorithm
- One time curve fitting with sample data
- Four parameter logistic regression (Sigmoidal)
- Range normalized to [0,1]

$$y = a + \frac{b-a}{1+\left(\frac{x}{c}\right)^d}$$

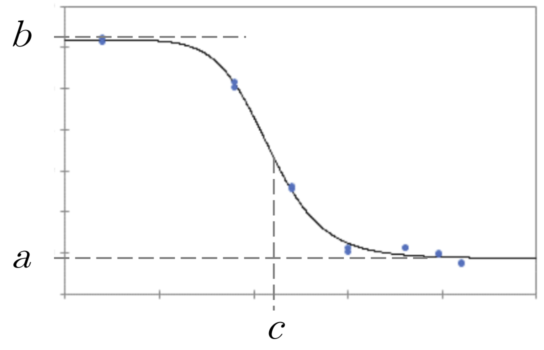

Initial Accuracy

- Depends on local conditions
- Time series prediction (ARIMA)
- Based on previous improvement rates

$$R_i = \frac{A_i - A_{i-1}}{A_{i-1}}$$

Response Time

- Depends on local conditions
- Support Vector Machine regression
- Based on previous response times
- Works well for time-to-failure

Metrics for Non-Stationarity

Non-Stationarity

Data Obsolescence

Velocity that old data becomes irrelevant to the model

Data Horizon

Urgency of including new evidence and updating the model

Depreciation by Stale Data (DSD)

The average rate of deterio-ration (in percentage) per old data points included

Time of Inflection (TOI)

The time it takes to observe a significant drop in the performance of a model (c)

Slope of the Characteristic Function (SCF)

The slope of the character-istic function at the point of inflection (d)

Contribution of Updates (COU)

The magnitude of improve-ment in the performance when the model is updated

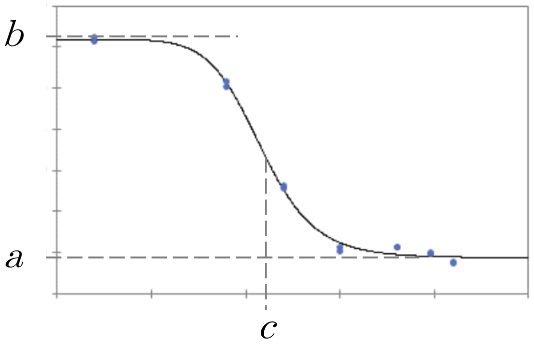

- Elec2 data set
  - Electricity market of New South Wales
  - Forecasting electricity demand
  - SVM regression model
  - 100 half-hourly data points

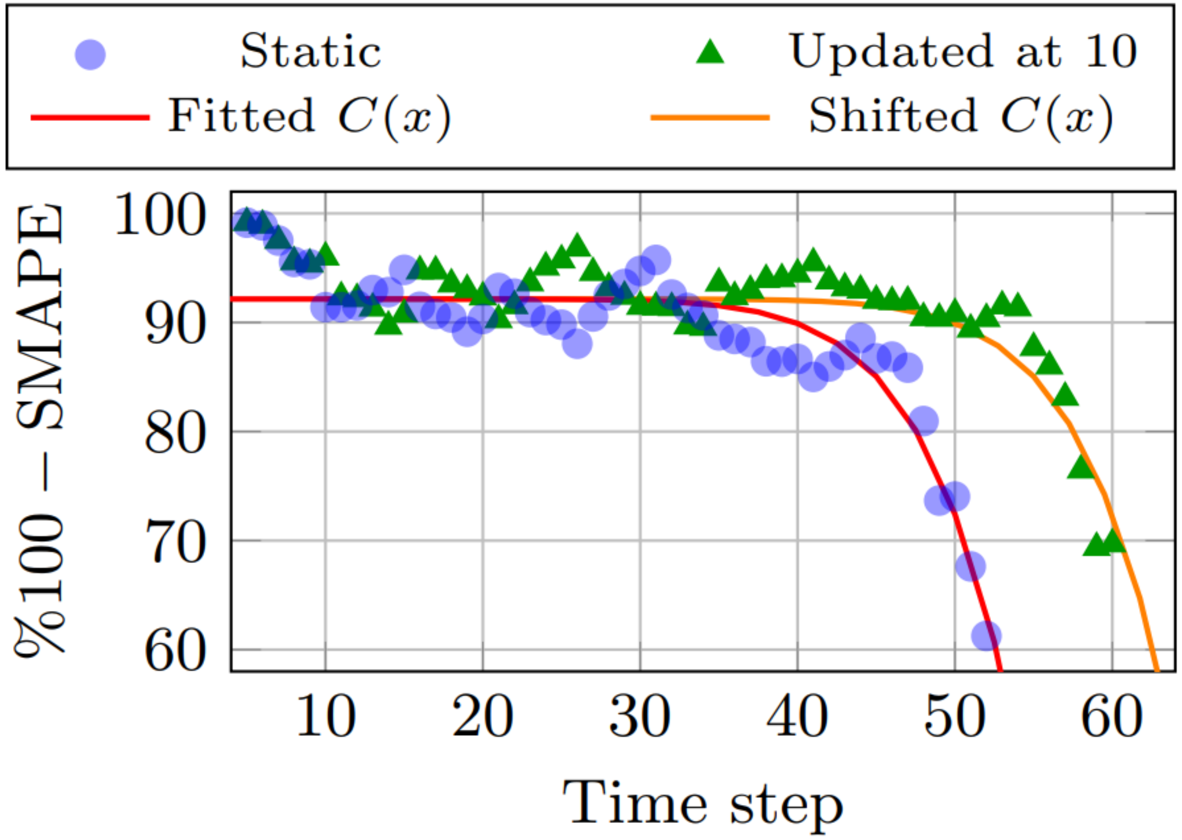

- Further experiments:
  - DBN availability prediction of client computers with SETI@Home data set
  - Correlation of non-stationarity metrics

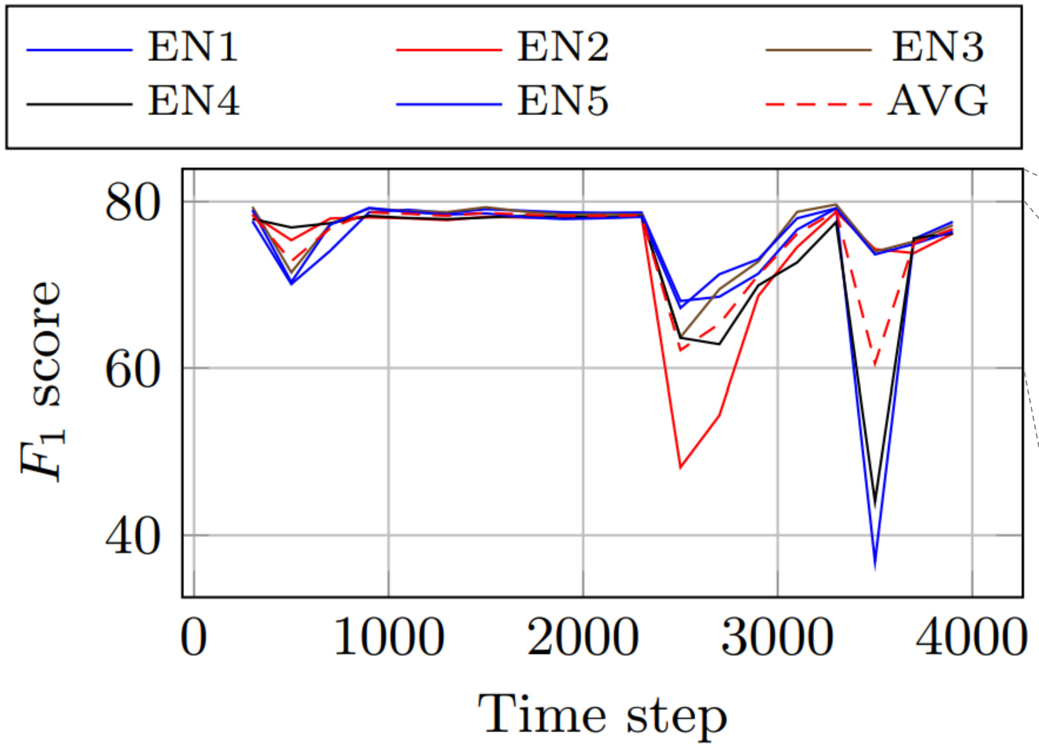

EN1
EN2
EN3
EN4
EN5
AVG
80
60
40
F_1 score
0
1000
2000
3000
4000
Time step

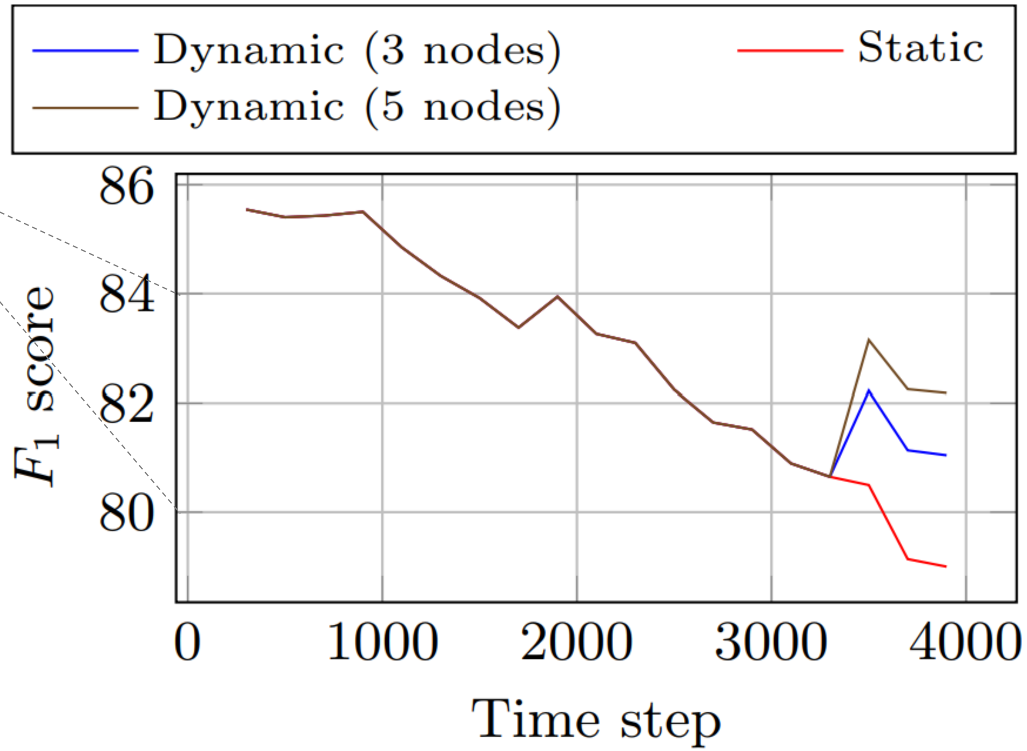

Dynamic (3 nodes)
Static
Dynamic (5 nodes)
86
84
82
80
F_1 score
0
1000
2000
3000
4000
Time step

Conclusion

- Takeaways:
  - Edge computing is promising for implementing DOML.
  - Non-stationary is a major challenge. It is measurable.
  - Dynamic periodicity can be an effective solution.

- Future Work:
  - Extended evaluation on a real DOML system
  - Source of the problem: avoiding (or at least predicting) stragglers
    - Resource selection
    - Load distribution

FAKULTÄT FÜR !NFORMATIK

Faculty of Informatics

Der Wissenschaftsfonds.

electronic commerce

RUCON

Thank you for your attention!

Atakan Aral, Ivona Brandić

Institute of Information Systems Engineering

Vienna University of Technology, Vienna, Austria

{atakan.aral, ivona.brandic}@tuwien.ac.at

http://rucon.ec.tuwien.ac.at/

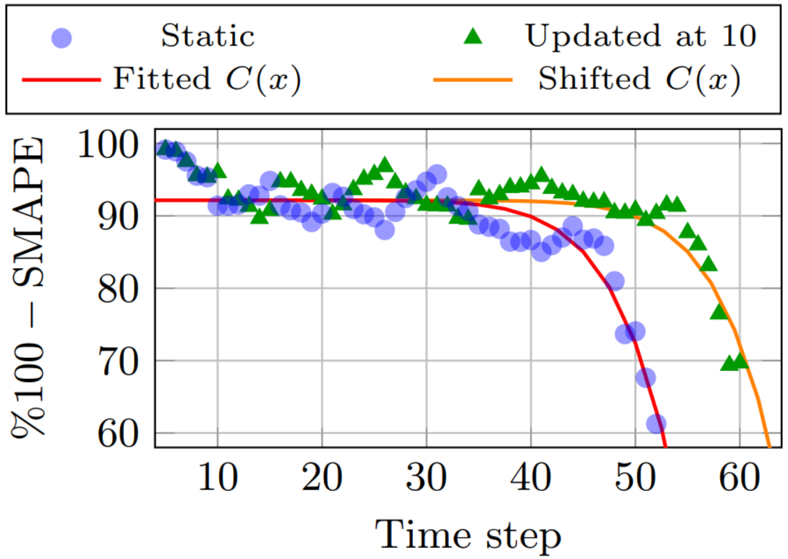

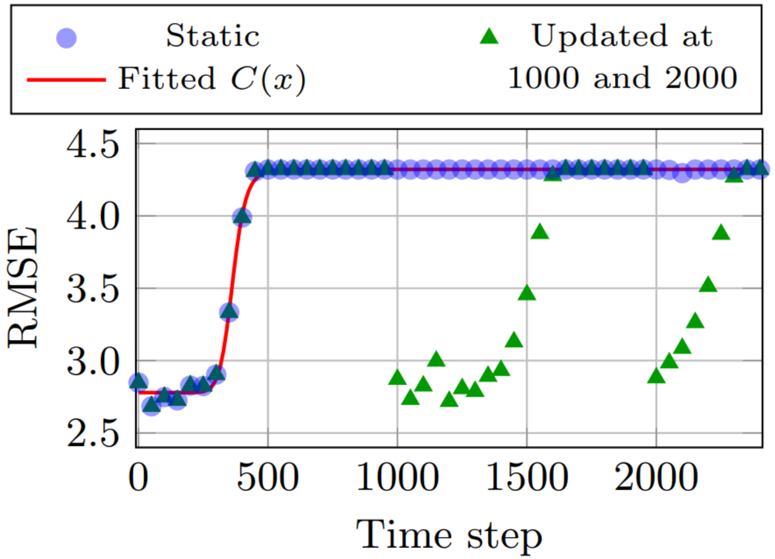

DS	ML (#P)	SCF	TOI	COU	DSD
[11]	SVM (10)	6.702	91.682	8.244	22.49
[11]	SVM (15)	7.408	84.037	3.302	23.71
[11]	SVM (20)	8.264	78.034	10.32	24.81
[11]	SVM (40)	8.394	67.497	5.636	23.54
[11]	SVM (50)	10.97	37.941	3.738	25.38
[13]	DBN	14.27	363.83	32.00	2.430

	SCF	TOI	COU	DSD
TOI	**-0.9879**			
COU	-0.3849	0.4654		
DSD	**0.8571**	-0.7731	-0.1505	
#P	**0.8212**	**-0.9489**	-0.4506	0.6325