Queen Mary University of London

Telefonica

Deep Massage: Processing Mobile Sensor Data for Online Deep Learning Predictions

Kleomenis Katevas, Ilias Leontiadis, Martin Pielot, Joan Serrà

Data Natives 2017

April 28, City, University of London

NFC

GPS

Ambient Light Sensor

Camera(s)

WiFi

Magnetometer

Water Sensor

Bluetooth

Barometer

Accelerometer

Microphone(s)

Gyroscope

Proximity Sensor

Motion Coprocessor

Machine Learning

- Accelerometer
- Battery level
- Network data activity
- Noise level
- Call logs
- Location
- Screen Status
- …

- Mean and Max Acceleration
- Battery drain levels in the last hour
- Data received during the day
- Average ambient noise level during the last five minutes
- Time since the last outgoing call
- Current distance from home
- Number of device unlocks
- …

- Random Forests
- XGBoost
- …

Machine Learning

- Accelerometer
- Battery level
- Network data activity
- Noise level
- Call logs
- Location
- Screen Status
- …

- Mean and Max Acceleration
- Battery drain levels in the last hour
- Data received during the day
- Average ambient noise level during the last five minutes
- Time since the last outgoing call
- Current distance from home
- Number of device unlocks
- …

- Random Forests
- XGBoost
- …

Who needs Feature Engineering?

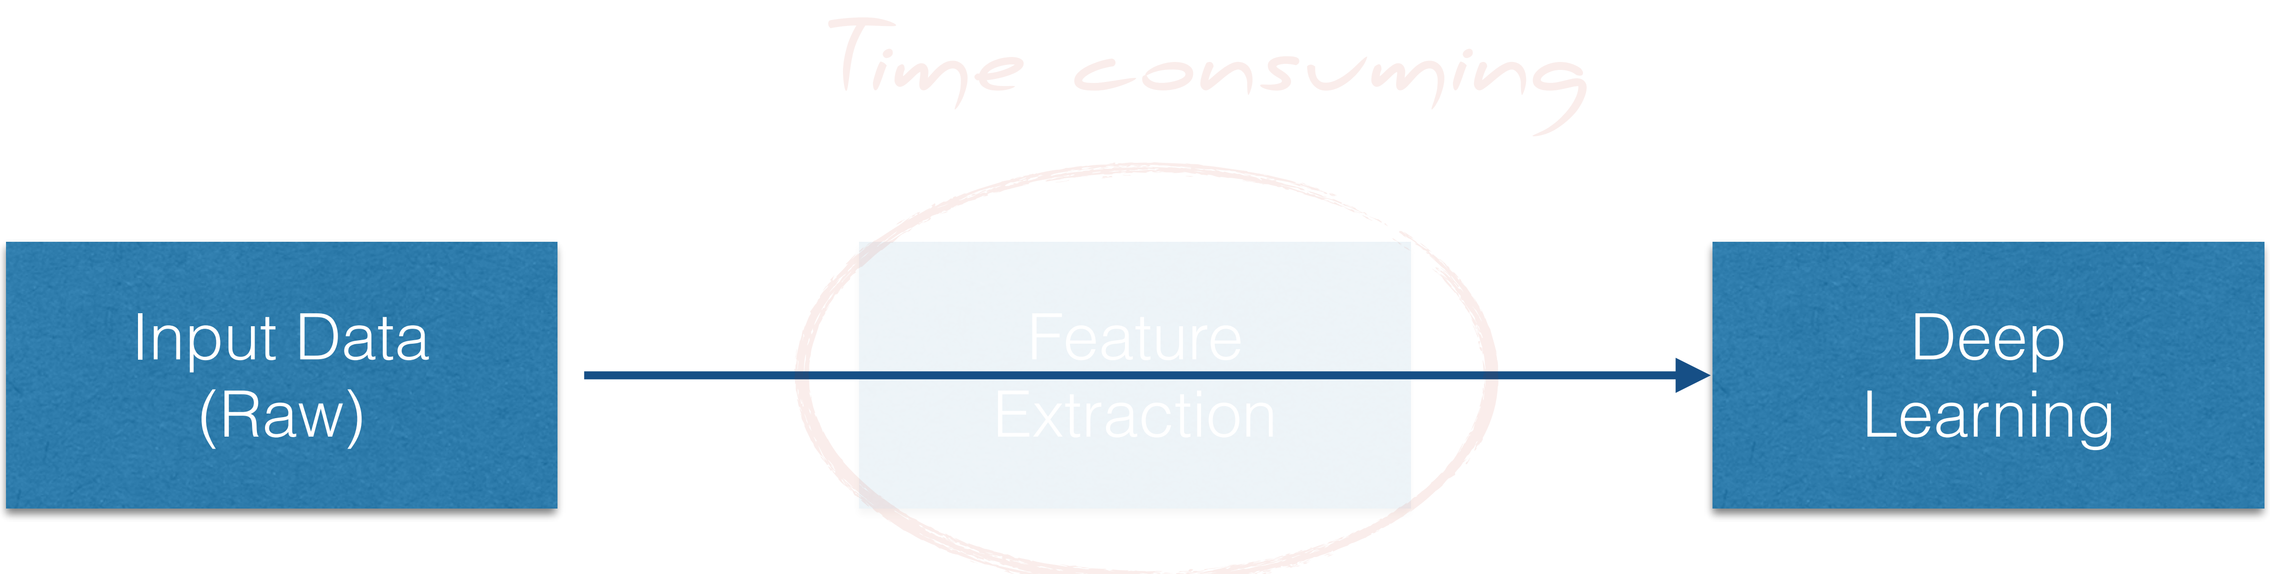

- Accelerometer
- Battery level
- Network data activity
- Noise level
- Call logs
- Location
- Screen Status
- …

- Mean and Max Acceleration
- Battery drain levels in the last hour
- Data received during the day
- Average ambient noise level during the last five minutes
- Time since the last outgoing call
- Current distance from home
- Number of device unlocks
- …

Deep Neural Networks

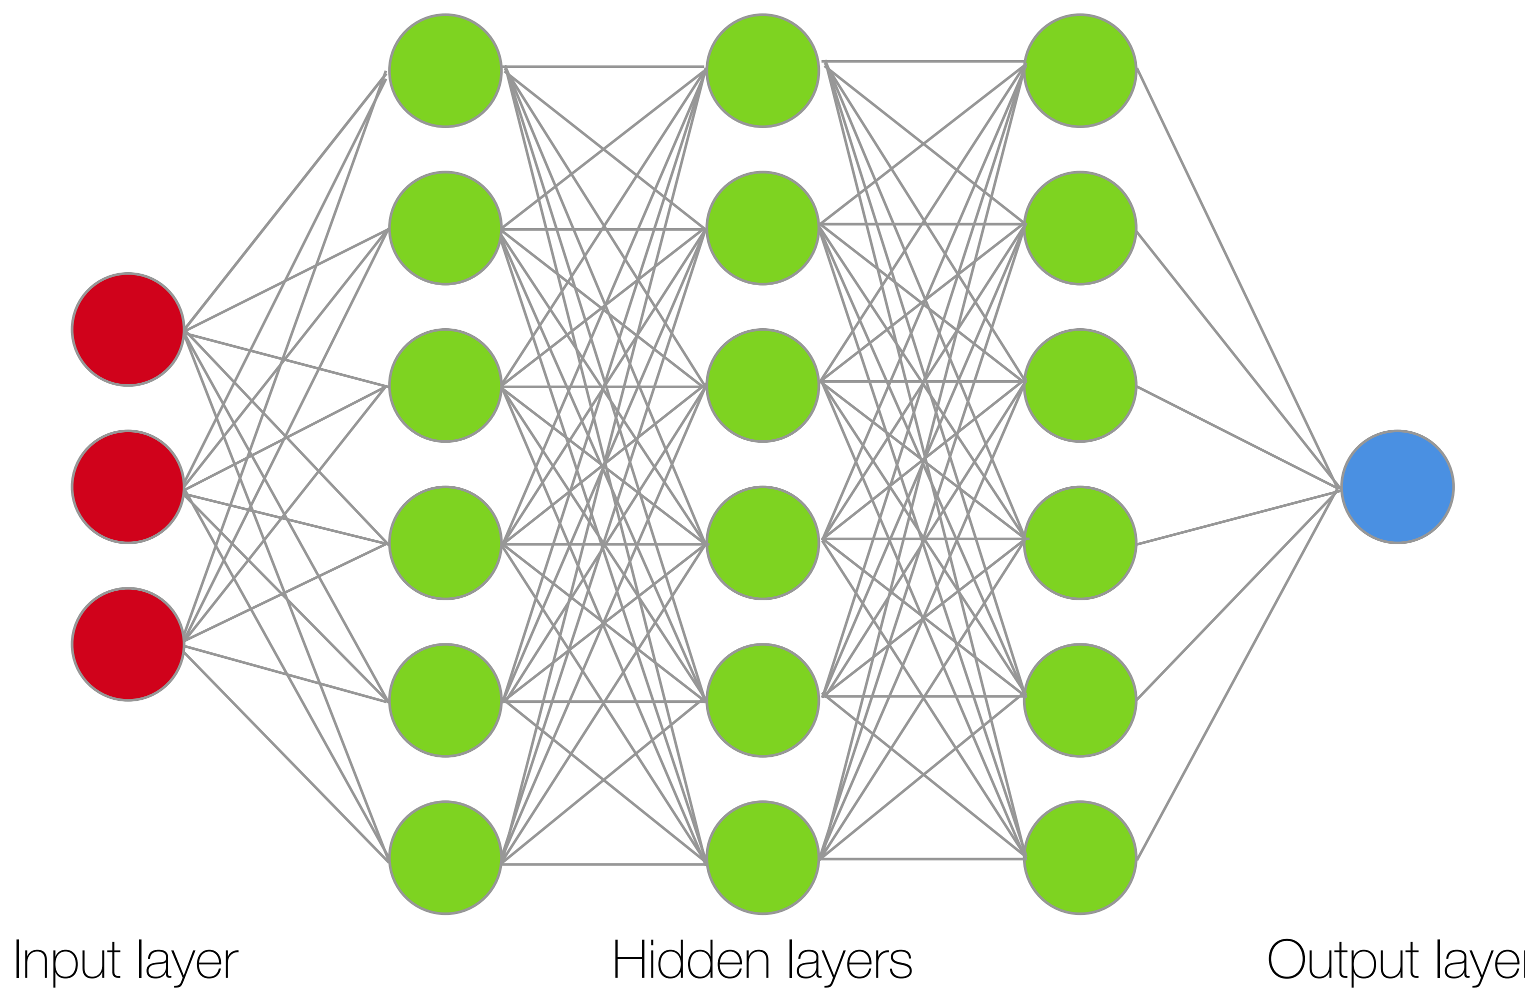

Recurrent Neural Networks (RNNs)

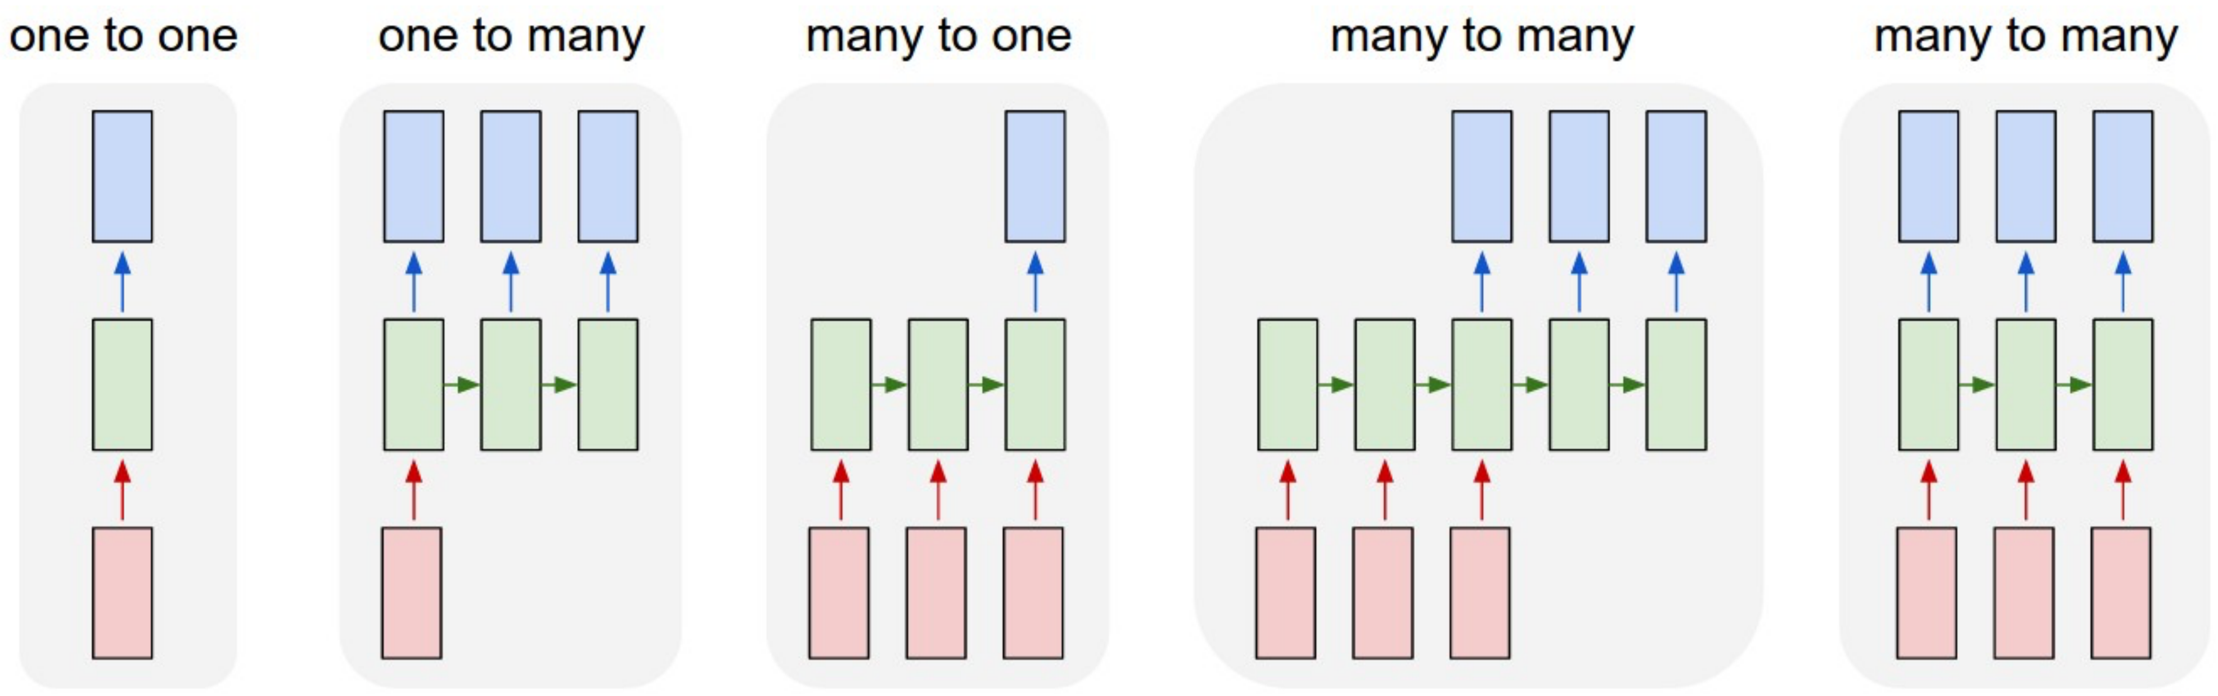

Source: Andrej Karpathy, 2015

Recurrent Neural Networks (RNNs)

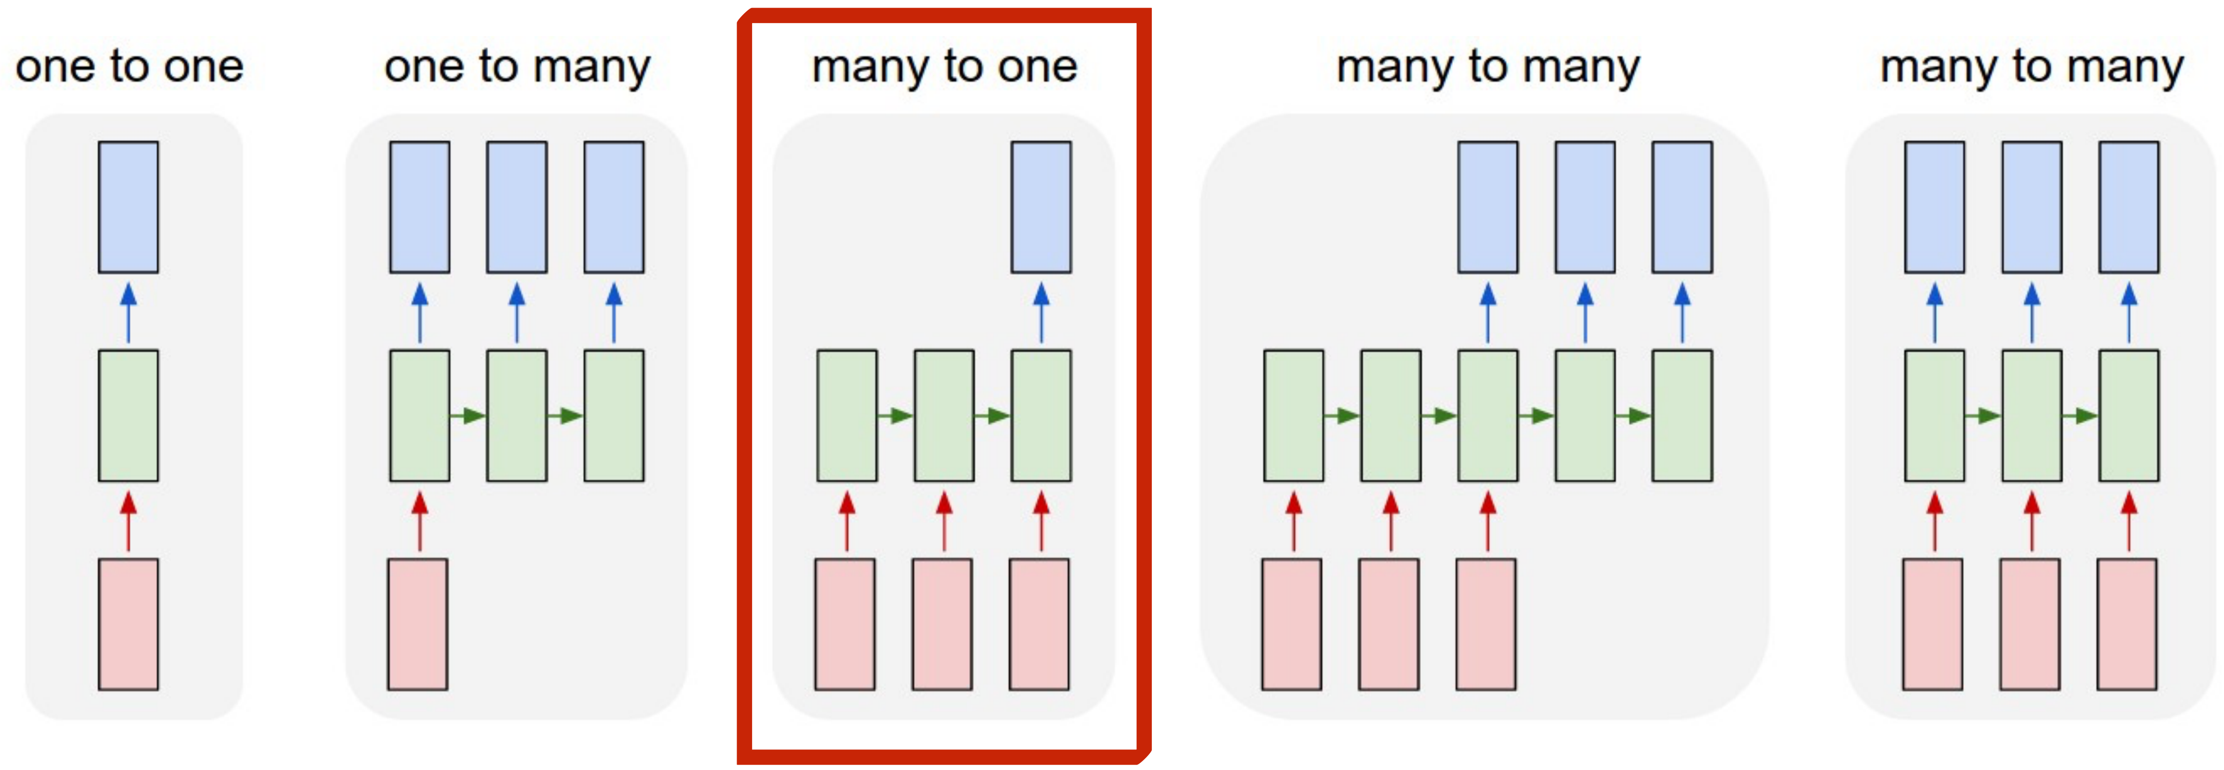

Source: Andrej Karpathy, 2015

Deep Learning Pipeline

		x				y	w
S_1	t_1	0	0	0.4	0	0	0
S_2	t_2	0.1	0.2	0	0	0	0
S_3	t_3	0	0	0	0.2	0	0
S_4	t_4	0	0	0.5	0	0	0
S_5	t_5	0.3	0.1	0	0	1	0.7
S_6	t_6	0	0	0	0.5	0	0

- S_i: Sensor event (one-hot encoded)
- t_i: Time delta
- x: Sensor values
- y: Ground truth
- w: Weight

Structure Sensor Data and GT | Normalisation | Capping | Time-based Compression | Structure Data for Training

Deep Learning Pipeline

		x				y	w
S_1	t_1	0	0	0.4	0	0	0
S_2	t_2	0.1	0.2	0	0	0	0
S_3	t_3	0	0	0	0.2	0	0
S_4	t_4	0	0	0.5	0	0	0
S_5	t_5	0.3	0.1	0	0	1	0.7
S_6	t_6	0	0	0	0.5	0	0

- S_i: Sensor event (one-hot encoded)
- t_i: Time delta
- *x*: Sensor values
- *y*: Ground truth
- *w*: Weight

Structure Sensor Data and GT | **Normalisation** | Capping | Time-based Compression | Structure Data for Training

Deep Learning Pipeline

		x				y	w
S_1	t_1	0	0	0.4	0	0	0
S_2	t_2	0.1	0.2	0	0	0	0
S_3	t_3	0	0	0	0.2	0	0
S_4	t_4	0	0	0.5	0	0	0
S_5	t_5	0.3	0.1	0	0	1	0.7
S_6	t_6	0	0	0	0.5	0	0

- S_i: Sensor event (one-hot encoded)
- t_i: Time delta
- *x*: Sensor values
- *y*: Ground truth
- *w*: Weight

Deep Learning Pipeline

		x				y	w
S_1	t_1	0	0	0.4	0	0	0
S_2	t_2	0.1	0.2	0	0	0	0
S_3	t_3	0	0	0	0.2	0	0
S_4	t_4	0	0	0.5	0	0	0
S_5	t_5	0.3	0.1	0	0	1	0.7
S_6	t_6	0	0	0	0.5	0	0

	x				y	w
t_{1+2+3}	0.1	0.2	0.4	0.2	0	0
t_{4+5}	0.3	0.1	0.5	0	1	0.7
t_6	0	0	0	0.5	0	0

Structure Sensor Data and GT | Normalisation | Capping | **Time-based Compression** | Structure Data for Training

Deep Learning Pipeline

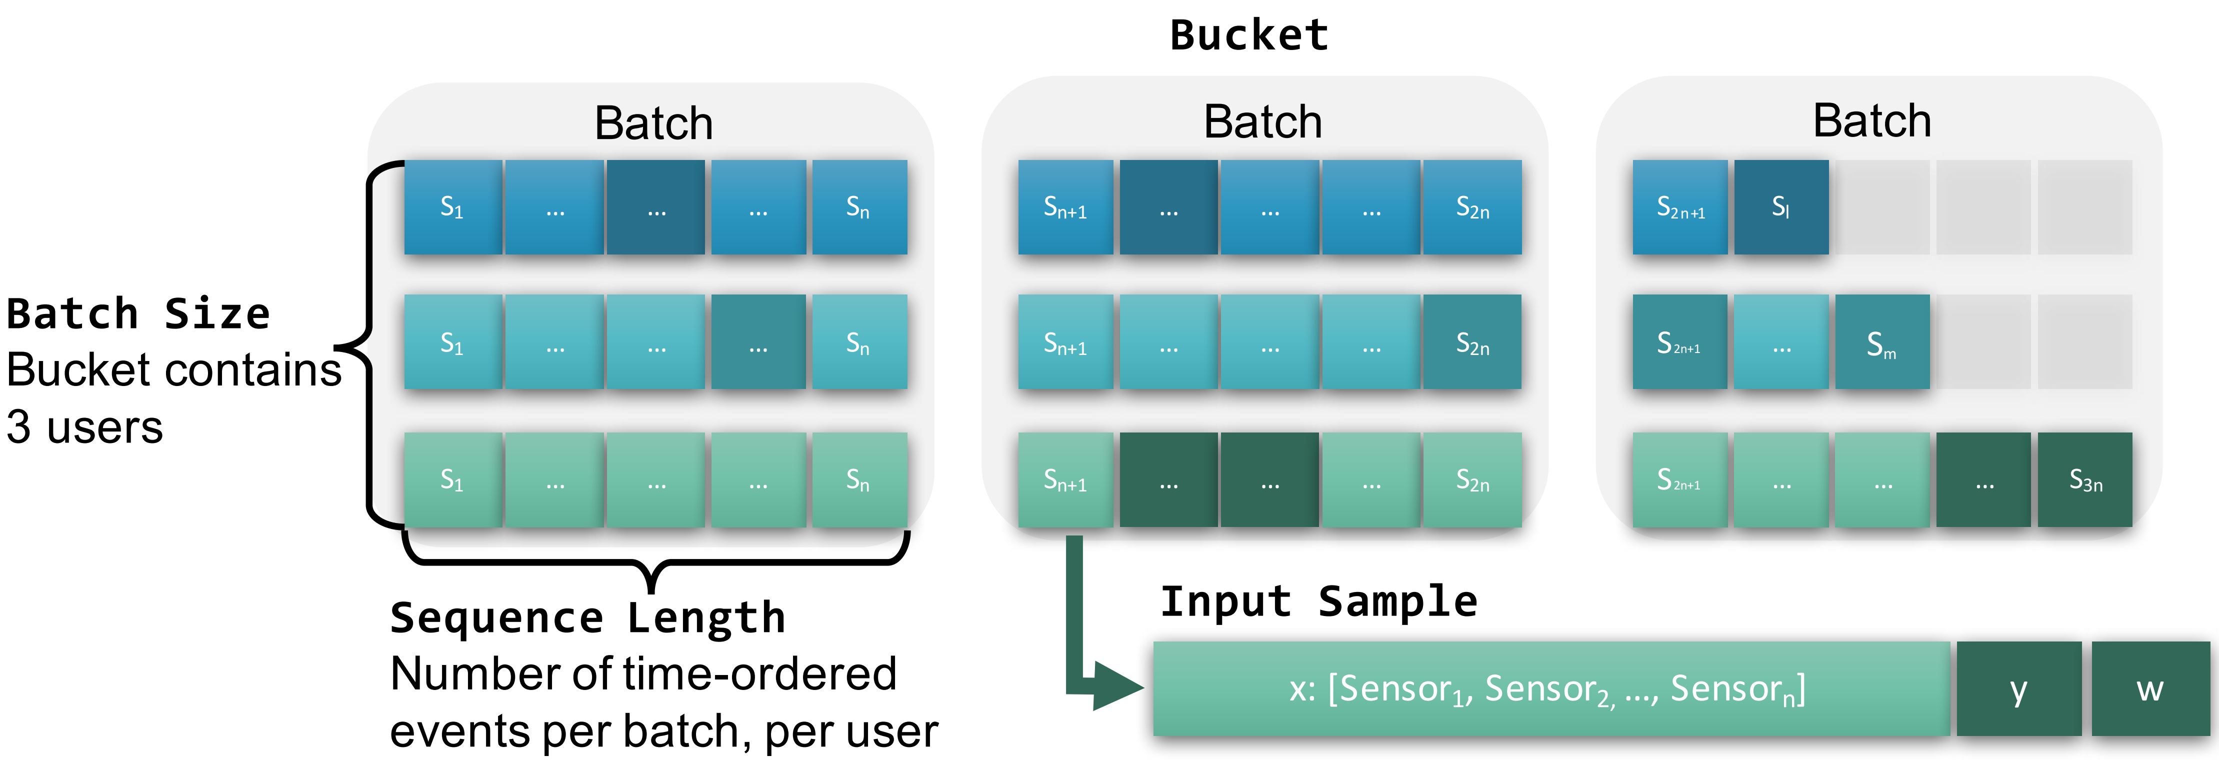

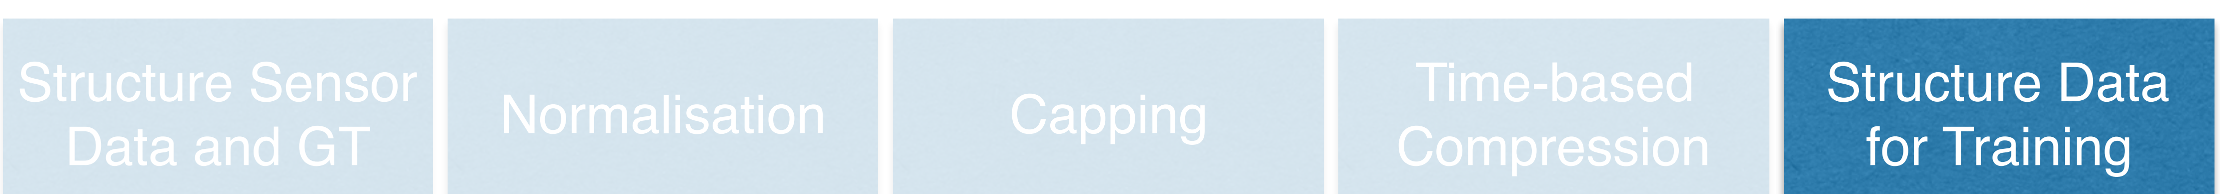

Case Study: Predicting Reactiveness to Notifications

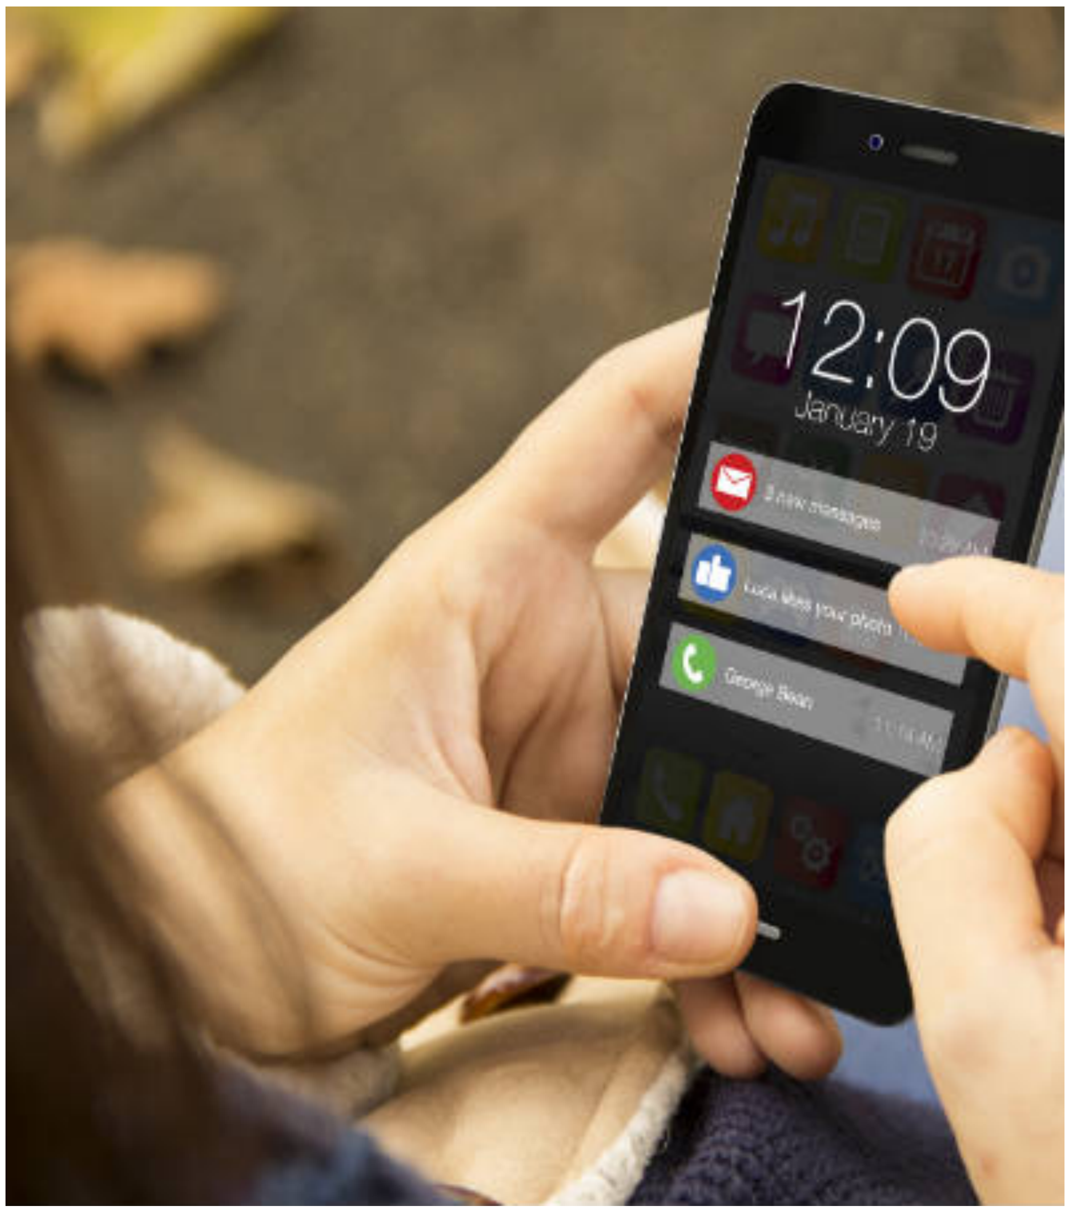

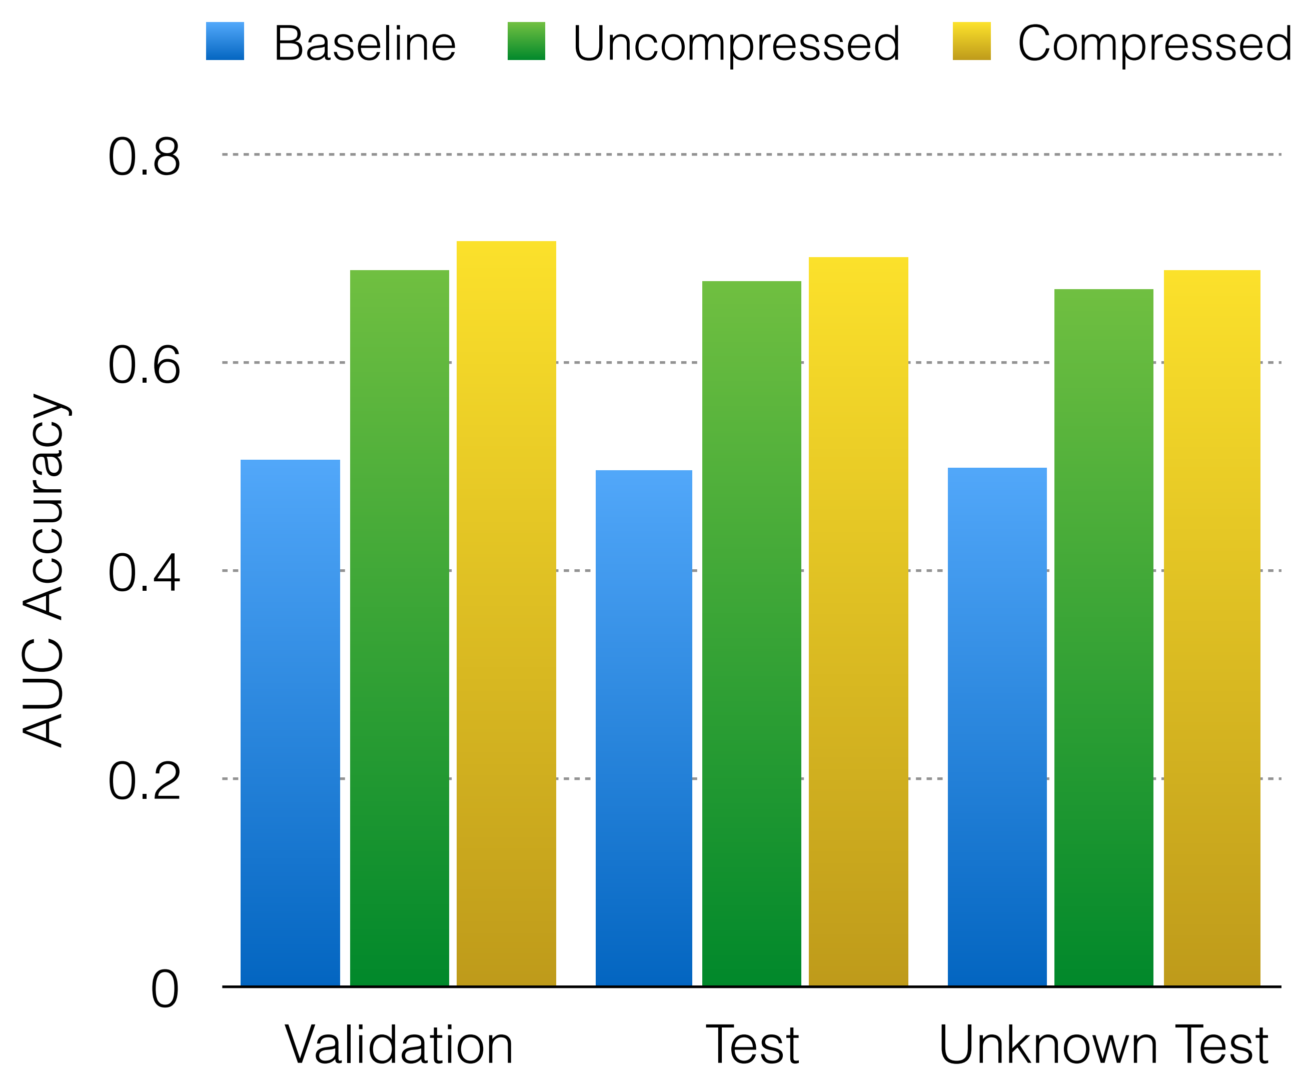

Conclusions

- Introduced an approach for preparing time series data for deep learning applications.
- Demonstrated the effectiveness in a case study.
- Achieved a 40% performance increase compared to a probabilistic random baseline.
- The model generalises to unknown users.

Future Work

- Comparison of the performance to canonical approaches.
- Improve the compression strategy.
- Explore more sophisticated deep learning techniques (e.g. transfer learning, generative adversarial networks).

Ilias Leontiadis

Martin Pielot

Joan Serrà