
An aerial photograph of Lausanne, Switzerland, showing the city's layout, the lake, and the surrounding mountains. The image is used as a background for the slide.

Machine Learning for Predictive Maintenance Applications: Introduction to Predictive Maintenance Feature Extraction

Prof. Dr. Olga Fink



- Combination of all technical administrative and managerial actions during the life cycle of an item intended to retain it in, or restore it to, a state in which it can perform the required function.



Maintenance: Not too late!

Langnau i.E.	7	unb.
St. Gallen	1 AB	ca. 25 Min später
Belp	6	ca. 25 Min später
Ost	4	unb. Verspätung
Brig	3 A-C	unb. Verspätung
Zweisimmen	3 C	unb. Verspätung
Aarau	12 A	unb. Verspätung
Bühl	4	unb. Verspätung
Westside	12 C	unb. Verspätung
		Ausfall
Dorf	24	
	12 A	unb. Verspätung
anshorn		Ausfall
kofen	22	


ICE	13.04	
ICE	13.04	
ICE	13.04	
ICE	13.06	
ICE 8	13.06	Europapl
S6	13.06	Burgdorf
RE	13.07	Worblau
S8	13.07	Bümpliz
S5	13.08	Bümpli
S5	13.08	Fribou
RE	13.09	Fribou
Totalunterbruch Bah		
Dauer noch unbesti		

Maintenance: Not too early!






Source: based on EN 13306
/EN 60300 / DIN 31051



Corrective Maintenance




Scheduled Maintenance



Condition-Based Maintenance



Predictive Maintenance



What is artificial intelligence?

Artificial intelligence (AI) refers to the ability of a computer or machine to perform tasks that would normally require human intelligence, such as learning, problem solving, decision making, and language understanding.

What is machine learning?

Machine learning allows computers to act and make data-driven decisions without being directly programmed to carry out a specific task.


What is the difference between AI and machine learning?

Machine learning is considered as a part of AI

Other methods for achieving AI include

- rule-based systems,
- evolutionary computation
- expert systems...

Examples of what AI can do today



Solve homework



Explain complex topics in a simple way



Debug code



Translate between languages




Trade stocks




Write music / film scripts



Generate ideas for party themes, decoration, presents...





Write blog posts / articles



Machine learning research is part of research on artificial intelligence, seeking to provide knowledge to computers through data, observations and interacting with the world. That acquired knowledge allows computers to correctly generalize to new settings.

Defintion by Yoshua Bengio, Université de Montréal




**Supervised
Learning**

VS




**Unsupervised
Learning**

VS




**Reinforcement
Learning**




- **Definition:** Generative AI refers to artificial intelligence systems capable of creating new content such as images, text, music, and more.
- **Key Techniques:**
 - Generative Adversarial Networks (GANs): Uses two neural networks, a generator and a discriminator, that compete to produce realistic outputs.
 - Transformer-Based Models (e.g., GPT): Uses deep learning to generate coherent and contextually relevant text.
- **Applications:**
 - Art and Entertainment: Creating artwork, music, and creative writing.
 - Data Augmentation: Generating synthetic data for training other AI models.
 - Content Creation: Developing marketing copy, news articles, and social media posts.
- **Benefits:**
 - Innovation: Enables the creation of novel and diverse content.
 - Efficiency: Automates content creation, saving time and resources.
- **Challenges:**
 - Ethical Concerns: Issues related to copyright, plagiarism, and misinformation.
 - Quality Control: Ensuring the generated content is accurate and high-quality.

What are Large Language Models (LLMs)





- Large Language Models are a type of artificial intelligence designed to understand and generate human-like text. Think of them as advanced chatbots that can hold conversations, answer questions, write essays, and even create poetry. They are typically trained on vast amounts of text data.
- **How do they work?**
- **Training Data:**
 - LLMs are trained on massive datasets containing text from books, articles, websites, and other written sources. This helps them learn the structure, grammar, and nuances of human language.
- **Learning Patterns:**
 - During training, the model learns to predict the next word in a sentence. For example, if given the phrase "The cat sat on the...", it might predict "mat" as the next word. By repeating this process billions of times, the model learns to generate coherent and contextually appropriate text.
- **Parameters:**
 - The "large" in LLM refers to the number of parameters (adjustable components) the model has. For instance, GPT-3, one of the well-known LLMs by OpenAI, has 175 billion parameters. More parameters generally mean the model can understand and generate more complex and accurate text.
- **What can they do?**
 - **Conversation:** LLMs can chat with users in a natural and engaging way.
 - **Content Creation:** They can write articles, stories, and even code.
 - **Translation:** LLMs can translate text between different languages.
 - **Summarization:** They can summarize long documents into shorter, easy-to-understand versions.




- Foundational models are large-scale machine learning models that are pre-trained on vast amounts of data and can be fine-tuned for various downstream tasks. These models form the "foundation" upon which more specific applications are built, allowing for efficient transfer learning and rapid deployment across different use cases.
- **Large-Scale Pre-Training:**
 - **Data:** Foundational models are trained on extensive datasets, often comprising diverse and unstructured data such as text, images, and audio.
 - **Architecture:** These models typically use complex architectures like transformers, which can capture intricate patterns and dependencies in the data.
- **Transfer Learning:**
 - **Fine-Tuning:** Once a foundational model is pre-trained, it can be fine-tuned on smaller, task-specific datasets to perform well on particular applications.
 - **Versatility:** This approach allows the model to adapt to a wide range of tasks without requiring training from scratch for each new task.
- **Scalability:**
 - **Parameter Count:** Foundational models often have billions of parameters, making them highly capable but also resource-intensive.
 - **Compute Resources:** Training and deploying these models require significant computational power, often leveraging specialized hardware like GPUs and TPUs.
- **Different types of foundational models:**
 - Education
 - Sciences
 - Health
 - Sustainability / climate
 - Robotics


Key success factors for breakthroughs in Artificial Intelligence



Advancement of hardware





Emergence of “big” data



Advancements in machine learning and deep learning


Key success factors for breakthroughs in Artificial Intelligence



Advancement of hardware



Emergence of “big” data



Advancements in machine learning and deep learning





11.09.24

Source: Joel Casutt, SBB


Olga Fink

2/31


Key success factors for breakthroughs in Artificial Intelligence




Advancement of hardware




Emergence of “big” data





Advancements in machine learning and deep learning




Top-5 Error Rate on ImageNet (%)









Andrew Ng


Elements of Intelligent Maintenance



Source: Deloitte








- Prognosis: Prediction of the Remaining Useful Life of a component
- Remaining Useful Life (RUL) – The amount of time a component can be expected to continue operating within its stated specifications.
- RUL is dependent on future operating conditions
 - Input commands
 - Environment
 - Loads
- Time of Failure (TOF): the time a component is expected to fail (no longer meet its design specifications)



Detect, Diagnose and Predict Faulty System Conditions → monitor the health evolution



Shelling




Non-roundness





From data acquisition to health management:


Prognostics & Health Management (PHM)



Source: P1856™/D31 Standard



Predictive vs. Prescriptive Maintenance







Predict the remaining useful life
Anticipate the failure
Reduce the impact of the failure
Determine the optimal point in time
for maintenance intervention

What can we do to prolong the
remaining useful life?
How can we proactively adjust the
operating conditions?
How can we control the process
parameters?




Five levels of condition-based and predictive maintenance




The background of the slide is a photograph of the famous bronze statue 'The Thinker' by Auguste Rodin. The statue is shown from the waist up, with the figure sitting on a rock and resting their chin on their hand in a contemplative pose. The background is a clear blue sky with some light clouds and green foliage visible on the left side.

“The greatest challenge to any thinker is stating a problem in a way that will allow a solution”

Bertrand Russell




- Define the learning problem in a way that allows its solution based on existing constraints such as lack of fault samples
- Design data-driven predictive maintenance applications for complex engineered systems from raw condition monitoring data
- Assess / Evaluate the performance of the applied algorithms
- Choose machine learning algorithms for fault detection, diagnostics and prognostics
- Interpret the results of the algorithms




- Performance will be assessed during the semester based on
 - 5 exercises, requiring the students to perform defined sub-tasks for designing a predictive maintenance system (60% of the final grade in total)
 - Feature engineering (12.5%)
 - Fault detection (12.5 %)
 - Prognostics (12.5 %)
 - Domain adaptation (12.5 %)
 - XAI (10%)
 - Final project: Report (including the implementation) and presentation of a real case study of designing a predictive maintenance system based on raw condition monitoring signals of a complex engineered system (30%) (→ presentation on 12.12.24 /16.12.24)
 - 2 quizzes on the content of the course → 10% of the grade

(Preliminary) Course Schedule



Week	Date lecture	Date exercise	Lecture	exercise / graded Exercise	Quizzie / Project
1	9/9/2024	12/9/2024	Introduction to predictive maintenance / Basic Principles of PM, problem formulations / Pre-processing of raw condition monitoring data / Feature extraction;	Introduction to sciPy + Pytorch	
2	16/9/2024	19/9/2024	How to approach data science projects? Problem definition / Feature selection / dimensionality reduction methodology	Introduction to Exercise 1 / introduction to good practice data science	
3	23/9/2024	26/9/2024		Repetition signal processing, support slot exercise 1	
4	30/9/2024	3/10/2024	Fault diagnostics / Residual-based detection / Feature Learning	Submission Graded exercise 1, Introduction to graded exercise 2	
5	7/10/2024	10/10/2024	One-Class Classifiers / Health indicators / Self-supervised learning / performance evaluation	Debriefing graded exercise 1, Support slot graded exercise 2	
6	14/10/2024	17/10/2024	Prognostics / Hyperparameter tuning	Submission graded exercise 2, Introduction to graded exercise 3	Quizzie1/ Intro to Final Project (Data Final Project)
7	21/10/2024	24/10/2024	Semester break	Semester break	
8	28/10/2024	31/10/2024	Domain adaptation approaches for PHM	Excursion ?	
9	4/11/2024	7/11/2024	Federated learning in PHM/ Explainable ML algorithms	Submission graded exercise 3, Introduction to graded exercise 4	
10	11/11/2024	14/11/2024	Physics-informed ML Models / GNN	Debriefing graded exercise 3, Support slot graded exercise 4, Exercise Federated learning	Milestone Final Project
11	18/11/2024	21/11/2024	Semi-supervised learning for fault diagnostics and prognostics / Generative models	Submission Graded exercise 4, Introduction graded Exercise XAI	
12	25/11/2024	28/11/2024	Fleet approaches / Prescriptive maintenance concepts	Debriefing graded exercise 4, Submission graded exercise 5; Exercise GNN	
13	2/12/2024	5/12/2024	Decision support	Support slot final project	Quizzie2
14	9/12/2024	12/12/2024	Benefits and costs of PM	Presentation final project	
14	16/12/2024	19/12/2024	Presentation final project		




Dr. Florent
Forest



Dr. Zhan
Ma




Ismail
Nejjar



Vinay
Sharma



Raffael
Theiler



Mengjie
Zhao



Amaury
Wei



Leandro von
Kannichfeldt



Han
Sun



Keivan
Fagih Niresi



Sergei
Garmaev




Chenghao
Xu




Zepeng
Zhang

Typical steps to follow in a machine learning (ML) project (1/3)



1. Problem Definition and Understanding
 - Clearly articulate the problem you aim to solve.
 - Identify the objectives and success criteria.
 - Ask the right questions: Identify the questions your data can answer to solve the business problem.
 - Understand the data: Learn the context, limitations, and opportunities within the available data.
2. Identify the Value:
 - Determine the potential value and impact of solving the problem (both from business and from the scientific perspective)
 - Assess how the solution will benefit stakeholders and align with business goals.
3. Collect Data:
 - Gather relevant data from various sources.
 - Make sure that your data is representative for the test data (application data) → be aware of the variability of the operating conditions (→ domain shift)
 - Ensure data quality and completeness.
 - Acquire labels for your data if possible (ensure the quality of labels)
4. Explore and Understand the Data:
 - Perform exploratory data analysis.
 - Visualize data to understand patterns and relationships.
 - Understand the data distribution (Explore the statistical properties, distributions, and trends in the data)
 - Identify data quality issues (Look for missing values, outliers, and inconsistencies that may affect analysis).
 - Generate hypotheses (Form hypotheses based on patterns observed during data exploration).

Typical steps to follow in a machine learning (ML) project (2/3)



5. Prepare Data:

- Clean the data (handle missing values, outliers).
- Feature engineering (create meaningful features, transform existing ones).
- Split data into training, validation, and test sets.
- Scaling and normalization (Prepare the data for modeling by applying appropriate scaling techniques)
- Address class imbalances: If needed, use techniques like oversampling, undersampling, or synthetic data generation to balance the dataset.


6. Select and Train Models:

- Select appropriate models (Choose models based on the problem type (regression, classification, clustering, etc.) and data characteristics)
- Split the dataset (Use train-test splits (or cross-validation) to ensure your model's generalization to unseen data)
- Baseline model (Start with a simple model as a baseline to compare more complex models).
- Tune Hyperparameters (Optimize model hyperparameters for better performance).
- Iterate and improve (Experiment with different models and tuning hyperparameters)
- Avoid overfitting (Implement regularization techniques or apply cross-validation to avoid overfitting the training data).

7. Evaluate Models:

- Choose appropriate metrics (Select evaluation metrics that are relevant to the problem (e.g., accuracy, precision, recall, F1-score, RMSE))
- Validate on unseen data (Always validate your model on a test set or through cross-validation to assess its performance)
- Compare models: Compare different models based on performance metrics and business value.

Typical steps to follow in a machine learning (ML) project (3/3)



8. Model Interpretation

- What do the results mean?
- Explainability (Ensure that you understand how your model works, and use XAI techniques such as SHAP values or feature importance to explain the results).
- Check assumptions (Ensure that the model's assumptions hold and are consistent with the problem and data context)

9. Test Model:

- Evaluate the final model on the test dataset to gauge its real-world performance.

10. Deploy Model:

- Integrate the model into a production environment.


11. Monitor and Maintain:

- Continuously monitor model performance.
- Update the model as needed based on new data and changing conditions.
- Monitor how your model is actually used



12. Documentation and Communication

- Document thoroughly (Keep detailed notes on every step, including data sources, preprocessing steps, model choices, evaluation, and deployment procedures)
- Communicate results (Present findings to stakeholders in a clear, actionable manner. Tailor communication based on the audience (e.g., technical team vs. business executives).)


Most importantly:



- **KEEP IT AS SIMPLE AS POSSIBLE!**
- If a linear regression model solves your problem, don't opt for more complex models.
- Evaluate whether applying rules or thresholds is enough to make the decision.
- Keep the decision maker in mind, and focus on how to present model outputs clearly.



Possible means of condition monitoring




- Thermography
- Measurement of compressed air consumption
- Ultrasonic testing technology
- Oil quality monitoring
- Current consumption measurement
- Vibration diagnosis
- Acoustic emission monitoring
- Computer vision

Video inspection: example SBB (Swiss Federal Railways)




Source: J. Bisang, SBB


Defect detection of track: example SBB (Swiss Federal Railways)




Welding




Plastic Particle




Surface Defect




Chewing Gum






Squat





Wheel Slip




Defect detection of track: example SBB





Source: J. Casutt, SBB



Source: Anybotics





Infrared thermography for condition monitoring





Various condition monitoring applications of infrared thermography: (a) Monitoring of machineries where abnormal surface temperature distribution is an indication of a probable flaw. (b) Inspection of liquid levels in industrial components. (c) Inspection of printed circuit boards. Localized defects like short circuits or current leakages produce hot-spots which can be easily detected by infrared thermography. (d) Typical thermal images of a transformer circuit breaker where the faulty regions can be clearly seen as hot-spots. (e) Inspection of shaft belt where the thermal anomaly is due to over-tightening of a belt. (f) Condition monitoring of three phase electrical panel where local hot spots are developed due to load imbalance.

Source: Infrared thermography for condition monitoring – A review




Distributed Acoustic Sensing (fibre optic cable sensing) for detecting defects in switching mechanisms




Vidovic, I., & Landgraf, M. (2019). Fibre Optic Sensing as Innovative Tool for Evaluating Railway Track Condition?. In *International Conference on Smart Infrastructure and Construction 2019 (ICSIC) Driving data-informed decision-making* (pp. 107-114). ICE Publishing.





Source: Semiotic labs
Kistler





Source: SPIE digital library
Dynamics of Civil Structures, Volume 2
Polytec





Collection method




Point-cloud data preprocessing



Two-dimensional wavelet decomposition







Aggregate 3D reconstruction





Source: Yishun Li

Detection of Subsurface Distress via Ground Penetrating Radar





Source: Yishun Li


Wayside monitoring devices at SBB



11.09.24



Accelerometers



New NVH Microphones


Versatile NVH
Microphone



Engine Noise
Microphone




Wheelhouse Brake
Noise Microphone




Source: GRAS

Olga Fink

Maintenance reports / manuals / event recordings









Natural Language Processing (NLP) in maintenance


- **NLP** can process maintenance logs, manuals, and other unstructured text data.
- **Predictive Maintenance Insights:**
 - Analyzing unstructured text from maintenance logs and work orders to predict potential equipment failures.
 - Detecting patterns in historical records to forecast breakdowns.
- **Automated Reporting & Documentation:**
 - Automatically generating maintenance reports by summarizing complex work orders or logs using NLP.
 - NLP-based transcription for voice notes or technician inputs.
- **Fault Detection from Logs:**
 - Scanning equipment error messages and sensor logs to identify and classify common faults.
 - NLP can interpret error codes, anomaly reports, and operational data in natural language.
- **Knowledge Management:**
 - Extracting, organizing, and summarizing information from vast amounts of maintenance manuals, technical documentation, and historical repair records.
 - Providing quick searchability of solutions through conversational AI or chatbot assistance.
- **Chatbots and Virtual Assistants:**
 - NLP-powered chatbots to assist maintenance technicians with real-time troubleshooting, guiding them through steps or providing manual references.

Source: H. Canaday







Source: Wang, 2012



- Transaction/event (data are “pushed” by data originator)
 - Data records occur only at the specified event / transaction / time stamp
 - Data between the time stamps / events are undefined.
- Sensor (data are “pulled” from data originator)
 - Data samples are acquired only at the specified time stamp
 - Data between the time stamps are just not observed.
 - Sampling rate
 - Evenly sampled – controlled (e.g. 100 Hz)
 - Unevenly sampled - triggered




Examples:

Operating mode,
Event code, asset ID

Performance
level; severity
level; friction level
(Low-Medium-High →
ranked levels)


Number of
occurring
diagnostic events,
number of
occurring faults /
interruptions

Temperature,
pressure,
acceleration (most
sensors)




- A type of categorical data in which there are only two categories
- Binary data can either be nominal or ordinal
- Examples: event status, on/off sensor


Numerical Data: Discrete vs. Continuous




Discrete Data



Continuous Data




Categorical Data: Representations



Frequency Tables


<i>Cause of Error</i>	<i>Number of Occurrence</i>	<i>Percentage</i>
Friction	28	25.2%
Obstacle	7	6.3%
Screws Loose	27	24.3%
High Temperature	37	33.3%
Low Temperature	12	10.8%


Pie Charts



Bar Charts


Failure by Categories








- Replacing values
 - Freely assign numbers to the categories according to the use case / expert knowledge
- Encoding labels
 - convert each categorical value in a column to a number between 0 and $n_categories - 1$
- One-hot encoding
 - convert each category value into a new column and assign a 1 or 0 (True/False) value to the column
- Binary encoding
 - first the categories are encoded as ordinal, then those integers are converted into binary code, then the digits from that binary string are split into separate columns
- ...


Example of one-hot encoding



Color		Red	Yellow	Green
Red		1	0	0
Red		1	0	0
Yellow		0	1	0
Green		0	0	1
Yellow		0	0	1








- Some categorical indicators can be used to split the problem in sub-problems (e.g. indicator of the operating conditions for base and part load → developing two models for the two types of operating conditions)

Signal dynamics (relative to sampling)



Stationary (constant + white noise)

- Power, speed, temperature in steady state of, gas turbines, etc.

Stochastic (non-cyclic)


- Power, torque, speed

Cyclic (consider each period individually)

- Power, speed, pressure in manufacturing process, gas turbine startup, take-off of airplanes


Waveform (consider multiple periods together)

- Vibration sensors, acoustic sensors




Source: Wang, 2012

Why data pre-processing?




- Data in the real world is “dirty”
 - incomplete: lacking attribute values, lacking certain attributes of interest, or containing only aggregate data
 - noisy: containing errors or outliers
 - inconsistent: containing discrepancies in codes or names



Main tasks in data pre-processing




- Data profiling
 - examining, analyzing and reviewing data
 - collect statistics about its quality.
- Data cleaning
 - Fill in missing values, smooth noisy data,
 - identify or remove outliers, and resolve inconsistencies
- Data integration (if needed)
 - Integration of multiple databases, data cubes, or files
- Data transformation
 - Normalization and aggregation
 - Structuring unstructured data
- Data reduction
 - Obtains reduced representation in volume but produces the same or similar analytical results
- Data discretization (if required)
- Data enrichment
 - Feature engineering
- Data validation
 - Assessing the dataset for quality assurance




- Times series observed with 15% missing data




Missing, noisy, inconsistent data



- Missing data
 - Data imputation approaches (next slide)
- Noisy data
 - Binning
 - Filtering
 - Clustering
 - Remove manually
 - Apply denoising algorithms
- Inconsistent data
 - External references
 - Knowledge engineering tools




- Complete case analysis: Delete any record that has missing values from the data set.
- Nearest neighbors: to impute variable x average the value x of the k closest data points with no missing values.
- Average method: Average the value of x for the non-missing values.
- Hot deck: pick a “similar” record at random and use its value of x .
- Predictive: Fit a model to the data with variable x as the target and use it to predict the value (e.g. kernel regression)
- Single imputation: Draw a value at random from the conditional distribution of x given the other variables
- Multiple imputation: Repeatedly draw values at random from the conditional distribution of x given the other variables (e.g. as above), creating new data sets. Make the predictions with these now complete datasets and average the predictions.



Caution with the different imputation approaches

- Complete case analysis: can result in a bias
- Nearest neighbors: The definition of the “close points” and the value of k required
- Average method: Easy to implement but crude
- Hot deck: A definition of “similar” is required
- Predictive: Better suitable but understates the uncertainty in the imputation process.
- Single imputation: Better suitable, respects the uncertainty. However, just a single value is sampled.
- Multiple imputation: generally regarded as the best method (a sample is better than a single observation)




- Feature scaling
(→ also referred to as min-max
Normalization)

$$X' = \frac{X - X_{min}}{X_{max} - X_{min}}$$


- Standard score
(particularly suitable for normally
distributed data)
(→ also referred to as Z-Score
Standardization)

$$X' = \frac{X - \mu}{\sigma}$$

Alternatives (particularly for data with outliers)




- Quantile Transformation
- Power Transformation (non-linear transformation)







- Create bins (e.g. equal depth, equal size) → e.g. different categories of part-load conditions of a gas turbine
- Additional smoothing possible → replace the values in a bin by their mean


What are Features and Why do we need them?




- Transform raw signals into more informative signatures (or fingerprints) of a system
- Reduce size / complexity of the dataset
- Provide a physical description / representation
- Reduce resources necessary for further processing
- Achieve intended objectives
- Features are often to be the most crucial point for PHM applications

Univariate versus multivariate feature engineering







Option 1: good feature engineering



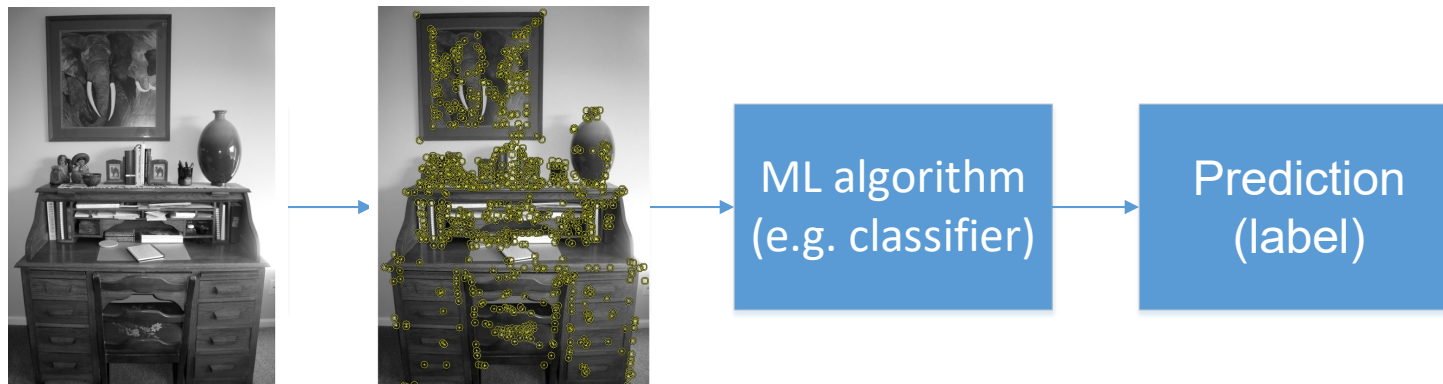

Features





$$\Rightarrow X_1 X_2$$




Example feature extraction in images



Example




FT Transform




Features: Magnitude of 3 carefully chosen frequency bands

In the future, we can compare these features for any measurements




Feature Extraction Process:



- **Raw Data:** Can be of different type and nature and size
- **Generation:** exhaustive, ad-hoc,
 - prior knowledge (domain, physics), modify variable type (cat. to numeric)
- **Feature Set:** Various representation and
 - size (dimensionality change)
- **Dimensionality reduction:**
 - select best subset
 - transform in space of lower dimensions

What matters for the process



- Kind of data
 - available type of features
- Aim of the application and domain
 - (eg. monitoring, detecting, predicting)
 - (vibration, turbines, electrical,...)
- Kind of methods that will be used
 - To extract feature,
 - To use the features for the application

Different Technologies

- Statistical analysis
- Signal processing
- Image processing
- Time-series analysis
- Control theory
- Information theory

Different PHM applications

- Vibration analysis
- Turbine machines
- Electrical systems
- Electronic devices
- Batteries
- SHM

Different data types


- Continuous
- Categorical
- Binary
- ...

Time dependency

- Time independent (stationary)
- Time dependent (non-stationary)

Univariate vs.
multivariate

Different data
sampling rate





Aims and properties of features


- Explainability
- Parsimony
- Robustness (eg. to missing data)
- Uncertainty handling
- Link to knowledge and physics...

In fact, many features come from domain know-how

- Domain: Mechanical, structural, thermal, electrical,...
- Kind of system: vehicle, turbine, machinery, ...
- Components gearbox, motor, pump, battery,...




Feature extraction



- For fault detection:
 - Features should be dependent on the failure (and failure type)


- Feature tuned to operating conditions
 - feature relevance depends on the OC

- Feature designed based on knowledge:
 - Known relationships
 - known behavior of components



Some feature extraction approaches

- Descriptive statistical features:
 - For regularly sampled data: moments, correlation, RMS, ...
 - For event based data: count, rate, duration, delay, ...
- Descriptive models:
 - Distribution/histogram
 - Information based (mutual information)
 - Regression models, curve fitting
 - Classification, clustering (class label as a feature), sequence matching
- Mathematical Transformation:
 - derivative, cumulative sum, power, log, (eg. log if noise depends on amplitude, log normal)
 - Principal/Independent Component Analysis, etc...



Some examples of features

Power / Energy

- Root mean squared (RMS)

- Signal to noise ratio.

- Power Loss Efficiency

- (P_{out}/P_{in})

AC signatures

- Frequency

- Wavelength

- Crest Factor

Curve shape

- Curve length

- Curvature

- Eccentricity

$$\text{Crest Factor} = \frac{0.5 * (x_{max} - x_{min})}{RMS}$$

Magnitude

- Peak to peak

- Root mean squared (RMS)

- Peak to average ratio (PAR)

Statistical measures (moments)

- Mean

- Variance Skewness

- Kurtosis

- Hyperskewness


- Hyperflatness

Multi-variables


- Correlation

- Principal Components

- Distance to centroid (Silhouette, Mahalanobis)




Example: Feature extraction for vibration analysis



feature engineering is a research field

Yan, W. et al, 2008




Some feature extraction Approaches

domain know-how

- Physics based :
 - expected input-output relations, etc.
 - comparison to expected output (model)

- Special procedures for data processing:
 - operational regime segmentations, envelop analysis, etc...
 - Time synchronous averaging

Statistical Features: Moments





$$\begin{aligned}\mu_n &= \int_{-\infty}^{\infty} (x - c)^n f(x) dx = \int_{-\infty}^{\infty} (x - c)^n dF(x) \\ &= E[(X - c)^n]\end{aligned}$$

- $c = 0$: raw moment
- $c = \text{average}(X)$: central moment
- Normalized (central) moment:


$$\frac{\mu_n}{\sigma^n} = \frac{E[(X - \mu)^n]}{\sigma^n}$$

n	Raw Moment	Central M	Normalised M
1	Mean	0	0
2		Variance	1
3			Skewness
4			kurtosis
5			Hyperskewness
6			Hypertailedness
7			

Skewness and Kurtosis




negative skew




positive skew

Kurtosis






Take away



- Procedure:
 - feature extraction + dimension reduction
- What to extract:
 - data property vs. application domain vs. algorithm requirements
- Feature extraction vs. signal processing
- Feature goodness:
 - Relevance and redundancy
- Feature selection:
 - wrapper approach vs. filter approach vs. embedding
- Feature consistency and sensitivity issues

From feature extraction to end-to-end / deep learning







Option 2: feature learning

Deep Learning Algorithm

Input Data


$$Y = f(X)$$

Output model



leading performance with abundant data

Feature engineering vs. Feature learning



Feature Engineering

Feature extraction

Feature dimensionality reduction

- Feature selection
- Feature low-dim. Projection

Knowledge based
Manuel, labor intensive
Domain/problem specific
Not scalable

Feature Learning

Shallow feature learning

- Supervised
- Unsupervised

Deep feature learning

- Supervised
- Unsupervised

Data driven
Automated
Generic
Scalable

