

VIBE PROJECT

Virtual Biomedical and STEM/STEAM Education

2021-1-HU01-KA220-HED-000032251



Funded by
the European Union



Erasmus+

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.



PÉCSI TUDOMÁNYEGYETEM
UNIVERSITY OF PÉCS

U.PORTO




Politechnika
Śląska



The logo for the VIBE Project. The word "VIBE" is in a large, white, stylized font. The letter "E" is composed of three horizontal bars, with the middle bar being teal. Below "VIBE" is the word "PROJECT" in a smaller, white, sans-serif font. A teal line starts from the bottom left of the "V" and extends diagonally down to the left, then horizontally to the right, ending under the "P".

VIBE PROJECT

Advanced Techniques in Pattern Recognition: From Features to Recognition Systems

A short, solid yellow horizontal line is positioned below the main title text.

Overview of Advanced Pattern Recognition

Pattern recognition is the backbone of intelligent systems, enabling machines to interpret and classify data in various forms such as images, speech, and text. This presentation explores advanced techniques in recognition systems, covering key areas in:






Image Recognition: how advanced algorithms extract meaningful patterns from visual data.




Speech Recognition: from raw audio signals to accurate transcriptions.




Text Recognition: understanding text from documents to complex natural language tasks.

Importance in Modern AI Applications


Pattern recognition plays a crucial role in numerous applications, pushing the boundaries of AI. Some of the most impactful areas include:



Healthcare: automated diagnosis through image and signal recognition (e.g., MRI analysis).



Autonomous Systems: real-time image, speech, and text processing in self-driving cars.



Natural Language Processing (NLP): text recognition and language understanding, empowering applications like virtual assistants.

Understanding these techniques is essential for advancing real-world AI systems that require precise and robust decision-making.

Key Mathematical Foundations



Politechnika
Śląska




UCZELNIA
BADAWCZA
INICJATYWA INŻYNIERSKA


VIBE
PROJECT

Mathematical Foundations in Pattern Recognition


Pattern recognition relies heavily on advanced mathematical tools to model and classify complex data. Key foundations include:



Probabilistic Models: understanding the likelihood of observed data.



Dimensionality Reduction: efficiently representing data in lower dimensions.



Feature Space Transformations: transforming raw data into a feature space where patterns can be recognized effectively.

In this section, we will briefly review these critical concepts and their relevance to modern pattern recognition tasks.

Probabilistic Models - Bayes' Theorem

Bayes' theorem is a fundamental concept in probability and statistics that describes the probability of an event based on prior knowledge of conditions related to that event. It provides a way to revise existing predictions or theories (update probabilities) given new or additional evidence. The theorem was named after 18th-century British mathematician Thomas Bayes. It is expressed mathematically as:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

Where:

- $P(A)$ is the prior probability of A,
- $P(B|A)$ is the conditional probability of B given A,
- $P(A|B)$ is the posterior probability of A given B,
- $P(B)$ is the prior probability of B.

Bayes' theorem has a wide range of applications in various fields, for example in machine learning - employed in spam filtering, natural language processing, or in medicine - determines the accuracy of medical test results. It plays a crucial role in classifiers such as Naive Bayes, which assumes feature independence for simplicity.



Erasmus+

Probabilistic Models - Hidden Markov Models (HMMs)

An HMM is defined by:

- S : a set of hidden states.
- A : a state transition matrix, where A_{ij} is the probability of transitioning from state i to state j .
- B : an observation probability distribution, where is the probability of observing o_t in state j ,
- π : the initial state distribution.

State Transition Probability (A):
state at time t , and s_j is the next

$$A_{ij} = P(q_{t+1} = s_j | q_t = s_i)$$

$$B_j(o_t) = P(o_t | q_t = s_j)$$

where q is the hidden

Observation Probability (B):
hidden state s_j .

$$\pi_i = P(q_1 = s_i)$$

where o_t is the observed data at time t , conditioned on being in

Initial State Distribution (π):

the probability of starting in state s_i at time $t=1$.




Politechnika
Śląska




Probabilistic Models - Hidden Markov Models (HMMs)


HMMs have been applied in various fields, including:




Bioinformatics: used extensively for biological sequence analysis, such as gene prediction and multiple sequence alignment. Variants like profile-HMMs and pair-HMMs help in modeling biological sequences effectively.



Speech Recognition: HMMs model phonemes and words in speech signals, enabling systems to recognize spoken language accurately.



Natural Language Processing: applications include part-of-speech tagging and named entity recognition, where sequences of words are analyzed for structure and meaning.



Computer Vision: HMMs are employed for action recognition in video sequences by modeling temporal patterns in human movements



Erasmus+

Probabilistic Models - Maximum A Posteriori (MAP) Estimation

MAP estimation is a method used to estimate unknown parameters of a statistical model, balancing both the data (likelihood) and prior knowledge (prior distribution). It seeks to maximize the posterior probability $P(\theta|D)$, where θ represents the model parameters and D represents the observed data.

Bayes' Theorem for MAP:

$$P(\theta|D) = \frac{P(D|\theta)P(\theta)}{P(D)}$$

where:

- $P(\theta|D)$: posterior probability (what we want to maximize)
- $P(D|\theta)$: likelihood (how well the parameters explain the data)
- $P(\theta)$: Prior (prior beliefs about parameters)
- $P(D)$: Marginal likelihood (normalizing constant)



Erasmus+

Probabilistic Models - Maximum A Posteriori (MAP) Estimation

MAP Estimation:

$$\theta_{MAP} = \underset{\theta}{\operatorname{argmax}} P(\theta|D) = \underset{\theta}{\operatorname{argmax}} P(D|\theta)P(\theta)$$

This is equivalent to maximizing $P(D|\theta)$ (the likelihood) multiplied by $P(\theta)$ (the prior). It can be interpreted as a balance between the model fitting the observed data and adhering to prior knowledge.

Example:

In speech recognition, MAP can be used to estimate the parameters of a Hidden Markov Model (HMM) by incorporating prior knowledge about speech patterns (e.g., phoneme sequences) into the likelihood estimation from the acoustic data.

Feature Space Transformations - Kernel Methods

Kernel methods are a class of algorithms used in machine learning for pattern analysis and solving complex decision-making problems. They operate by mapping data into a higher-dimensional feature space, where linear patterns can be detected more easily. This mapping is achieved implicitly through the use of a kernel function, which calculates the inner product between the images of two data points in the feature space without explicitly computing the mapping.

Kernel methods have been successfully applied in various domains, including:

- Classification: Support Vector Machines (SVMs) use kernel methods for non-linear classification tasks.
- Regression: Kernel Ridge Regression extends linear ridge regression to non-linear problems.
- Dimensionality Reduction: Kernel Principal Component Analysis (KPCA) performs non-linear dimensionality reduction.
- Clustering: Kernel K-Means and spectral clustering algorithms utilize kernel methods for non-linear clustering.

Advantages: handles nonlinear patterns efficiently.

Applications: face recognition, image classification.



Erasmus+

Feature Space Transformations - Manifold Learning

Manifold learning is a subset of machine learning techniques focused on understanding and visualizing high-dimensional data by uncovering its underlying low-dimensional structure. It is particularly useful for tasks such as dimension reduction, data visualization, and anomaly detection. A manifold is a mathematical space that, on a small scale, resembles Euclidean space. In the context of data, it refers to the idea that high-dimensional data often lies on or near a lower-dimensional manifold.

Common Manifold Learning Techniques

- Isomap: this method constructs a graph based on the distances between neighboring points and estimates geodesic distances to preserve the manifold's structure. It effectively captures global geometric properties.
- Locally Linear Embedding (LLE): LLE focuses on preserving local relationships by reconstructing each data point as a linear combination of its nearest neighbors. This technique is particularly effective for complex shapes like folded or twisted manifolds.
- t-Distributed Stochastic Neighbor Embedding (t-SNE): t-SNE is widely used for visualizing high-dimensional data in two or three dimensions. It emphasizes preserving local distances and is effective in revealing clusters in the data.
- Variational Autoencoders (VAEs): VAEs are generative models that learn to encode data into a lower-dimensional space while maintaining essential information, allowing for effective dimension reduction.

Applications: visualization of high-dimensional image datasets (e.g., face images under different lighting conditions).

Source: Wang, C. (2010). A geometric framework for transfer learning using manifold alignment.



Erasmus+

Dimensionality Reduction

As datasets grow in complexity, reducing the number of features without losing critical information becomes crucial. Dimensionality reduction not only simplifies models but also improves efficiency.

1. Principal Component Analysis (PCA): PCA projects data onto the directions of maximum variance, reducing dimensionality while preserving the most important features. It is particularly effective for high-dimensional datasets like images.

Example: PCA in facial recognition to reduce pixel-based features into principal components representing critical patterns.

1. Linear Discriminant Analysis (LDA): while PCA maximizes variance, LDA maximizes the separation between classes. This is especially effective when classifying data that can be linearly separated.


Applications: face recognition (Fisherfaces), document classification.

1. 3. t-SNE and UMAP: For visualization of high-dimensional data, techniques like t-SNE (t-distributed Stochastic Neighbor Embedding) and UMAP (Uniform Manifold Approximation and Projection) reduce dimensionality while maintaining the structure of data, allowing us to visually explore clusters and relationships.


Applications: image clustering, word embeddings visualization.

Summary of Mathematical Foundations


The key mathematical concepts we've reviewed provide the underlying structure for advanced pattern recognition:



Probabilistic Models: capture uncertainty and update predictions with new data.



Feature Space Transformations: make complex data separable and more easily classifiable.



Dimensionality Reduction: simplifies data, reducing complexity without losing important information.

These foundations enable the development of robust and efficient recognition systems across image, speech, and text domains, ensuring better generalization and performance.

Feature Engineering in Image Recognition



Politechnika
Śląska




UCZELNIA
BADAWCZA
INICJATYWA INNOWACYJNA


VIBE
PROJECT

Feature Engineering in Image Recognition


Feature engineering is critical to image recognition systems, helping to transform raw pixel data into structured features that models can efficiently learn from. The goal is to extract and represent relevant patterns, edges, textures, or objects from images. This process has evolved from traditional handcrafted features to deep learning methods that automatically learn representations.



Handcrafted features (SIFT, HOG)



Deep learning-based feature extraction (CNNs)



Handling variations in scale, rotation, and noise



Erasmus+

Traditional Handcrafted Features - SIFT

Before the deep learning revolution, image recognition systems relied on handcrafted features—manually designed algorithms to extract key aspects of an image. Some widely-used methods include:

1. Scale-Invariant Feature Transform (SIFT):

- **Objective:** Extract distinctive, invariant features from images that remain stable across scale, rotation, and translation.
- **How it works:**
 - Detects key points (local extrema) using Difference of Gaussians (DoG).
 - Describes local gradients around each key point using orientation histograms.
 - SIFT features are robust against changes in viewpoint, illumination, and partial occlusion.
- **Applications:** object recognition, image stitching, panorama creation.



Erasmus+

Traditional Handcrafted Features - HOG

2. Histogram of Oriented Gradients (HOG):

- **Objective:** Capture the distribution of gradient directions in localized portions of an image.
- **How it works:**
 - The image is divided into small cells, and within each cell, a histogram of gradient orientations is computed.
 - These histograms are then concatenated to form a feature vector for image representation.
 - HOG features are particularly robust for tasks like object detection.
- **Applications:** Pedestrian detection, object recognition.

Handcrafted features like SIFT and HOG are designed to be invariant to transformations like scaling and rotation, but they have limitations in dealing with complex, real-world images.



Erasmus+

Deep Learning-Based Feature Extraction (CNNs)

Convolutional Neural Networks (CNNs) are a type of deep learning algorithm that is particularly well-suited for image recognition and processing tasks. They are inspired by the structure of the human visual cortex and have a hierarchical architecture that allows them to learn and extract features from images at different scales.

Key components of a CNN include:

- Convolutional Layers: apply convolutional operations to input images using filters to detect features like edges, textures, and patterns while preserving spatial relationships,
- Pooling Layers: downsample the spatial dimensions of the input, reducing computational complexity and the number of parameters,
- Activation Functions: introduce non-linearity to allow the model to learn complex relationships, such as ReLU,
- Fully Connected Layers: make predictions based on the high-level features learned by previous layers.

Source: <https://www.geeksforgeeks.org/convolutional-neural-network-cnn-in-machine-learning/>

Popular CNN Architectures for Feature Extraction

1. AlexNet:

- Architecture: consists of five convolutional layers followed by three fully connected layers.
- Breakthrough: first to show that deep CNNs trained on large datasets (like ImageNet) outperform traditional methods.
- Key Features: introduced ReLU activation and dropout, improving training efficiency and preventing overfitting.

2. VGGNet:

- Architecture: uses very deep networks (up to 19 layers) with small 3×3 convolution filters.
- Key Features: known for simplicity and uniformity in filter size, making it easier to understand and implement.
- Applications: transfer learning, where pretrained VGG features are reused for new tasks.

3. ResNet:

- Architecture: introduces skip connections to solve the vanishing gradient problem, allowing for extremely deep networks (e.g., 152 layers).
- Key Features: residual blocks that make training deep networks more efficient.
- Impact: ResNet won the ImageNet competition and remains widely used in feature extraction for complex tasks.

4. EfficientNet:

- Architecture: scales model depth, width, and resolution systematically, making it highly efficient.
- Key Features: balances performance and computational efficiency, often used in edge devices for image recognition.



Erasmus+

Transfer Learning and Feature Extraction

Transfer learning is a widely used technique for feature extraction in image recognition, where pretrained models are adapted to new tasks:

- Pre-trained Models: transfer learning typically begins with a model that has been pre-trained on a large dataset for a specific task. This model has already learned general features that can be useful for similar tasks.
- Fine-tuning: after selecting a pre-trained model, the next step involves fine-tuning it on the new task's dataset. This process adjusts the model's parameters to better align with the specific requirements of the new task while retaining the knowledge from the original training.
- Layers in Transfer Learning:
 - Frozen Layers: layers of the pre-trained model that remain unchanged during fine-tuning. They retain the general features learned during the initial training.
 - Trainable Layers: modified or retrained layers during the fine-tuning process to adapt to the specifics of the new task.

Advantages:

- Reduced Training Time: since the pretrained model has already learned general features, only the last few layers need to be retrained.
- Improved Accuracy: transfer learning can significantly improve accuracy when limited labeled data is available for the new task.

Applications: medical image analysis, satellite image classification, facial recognition.

Handling Variations and Robustness in Image Features

Robust feature extraction involves handling variations in data that can affect the performance of machine learning models. These variations can include:

- noise,
- illumination changes,
- scale and rotation,
- occlusion,
- deformation,
- background clutter.


Strategies for improving robustness:

- Data Augmentation: applying transformations such as rotations, scaling, and random noise during training makes the model more resilient to real-world variations.
- Adversarial Training: introducing small, imperceptible perturbations to the training data to increase the model's ability to handle adversarial attacks and distorted images.




Erasmus+


Summary of Feature Engineering in Image Recognition




Handcrafted Features like SIFT and HOG laid the foundation for early image recognition tasks, but they have limitations in flexibility and scalability.



Deep Learning (CNNs) revolutionized feature extraction by automatically learning representations from data, achieving state-of-the-art performance in various applications.



Transfer Learning enables efficient adaptation of pretrained models to new tasks.



Robustness to scale, rotation, noise, and occlusion is critical for real-world image recognition systems, and modern techniques handle these challenges effectively.

Speech Recognition: Features and Models



Politechnika
Śląska



UCZELNIA
BADAWCZA
INICJATYWA INŻYNIERSKA

VIBE
PROJECT




Erasmus+

Speech Recognition: Features and Models

The speech signal is complex and contains rich information about linguistic content, speaker identity, emotion, and even environment. Key properties include:

- Time-Varying Nature: speech signals change over time, making temporal modeling crucial.
- Frequency Components: different speech sounds have unique frequency signatures.
- Context Dependence: meaningful units of speech (phonemes, syllables) depend on their context in the signal.



Source: <https://medium.com/data-science/beginners-guide-to-speech-analysis-4690ca7a7c05>

Speech recognition systems aim to convert spoken language into text. This involves several key stages:

- Feature extraction: converting raw audio signals into meaningful representations.
- Modeling: mapping extracted features to phonemes, words, and sentences using probabilistic and deep learning models.



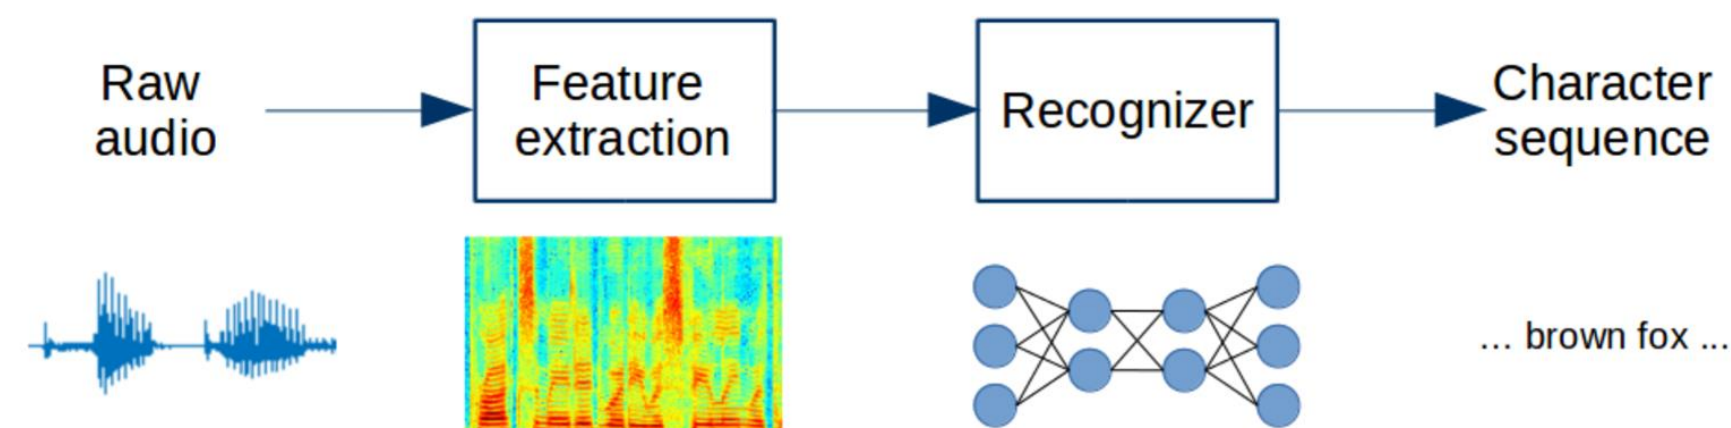
Erasmus+

Acoustic Features for Speech Recognition

Feature extraction is the process of converting raw audio into compact, informative representations. The goal is to capture the essential aspects of speech while being robust to variations like accent, speed, and noise.

1. Perceptual Linear Prediction (PLP):

- **Overview:** PLP is another common feature extraction method that incorporates psychoacoustic models to capture speech characteristics more accurately.
- **How it works:**
 - Similar to MFCCs, PLP applies a series of transformations to the speech signal, but it emphasizes perceptual properties such as the critical-band spectrum.
 - PLP often provides better performance in noisy environments.
- **Applications:** Frequently used in speaker identification and recognition systems due to its sensitivity to speaker-specific traits.



Source: <https://www.mdpi.com/2079-9292/9/7/1157>



Erasmus+

Acoustic Features for Speech Recognition

2. Mel-Frequency Cepstral Coefficients (MFCC):

- **Overview:** MFCCs are the most widely used features in speech recognition, designed to mimic the human auditory system's response to sound.
- **How it works:**
 - The speech signal is divided into short frames (20-40ms).
 - Each frame undergoes a Fourier transform to capture frequency information.
 - Frequencies are then mapped onto the Mel scale (a perceptually motivated scale), and the logarithm of the power spectrum is taken.
 - Finally, the Discrete Cosine Transform (DCT) is applied to obtain the MFCCs.
- **Applications:** used in both traditional and modern speech recognition systems due to their efficiency and robustness.

3. Log Filterbank Energies (FBANK):

- **Overview:** unlike MFCC, FBANK features skip the DCT step, keeping the log-filtered energies. This makes FBANK more suitable for deep learning models, which can process raw filterbank features directly.



Erasmus+

Traditional Models in Speech Recognition

Before the advent of deep learning, speech recognition relied heavily on probabilistic models that could handle the sequential nature of speech.


1. Hidden Markov Models (HMMs):

- **How it works:**

- Speech is modeled as a sequence of hidden states (e.g., phonemes), with transitions governed by a Markov process.
- Each state emits an observable output (acoustic feature vectors), modeled by a probability distribution.
- The Viterbi algorithm is used to find the most likely sequence of states corresponding to the observed speech features.

- **Strengths:** effective for modeling time-varying sequences and widely used for speech recognition tasks for decades.

- **Limitations:** HMMs assume conditional independence between observations, which limits their ability to capture long-term dependencies in speech.



Source: <https://wisdomml.in/hidden-markov-model-hmm-in-nlp-python/>




Erasmus+

Traditional Models in Speech Recognition

2. Gaussian Mixture Models (GMMs):


- **How it works:**
 - GMMs model the probability distribution of acoustic feature vectors in each HMM state.
 - Each state in an HMM is represented by a mixture of Gaussian distributions, allowing the system to capture variations in speech signals within each state.
- **Strengths:** GMM-HMM combinations were the dominant approach in early speech recognition systems due to their robustness and flexibility in modeling complex distributions.
- **Limitations:** GMMs require large amounts of data and struggle with high-dimensional features compared to modern deep learning approaches.




Source: <https://www.mdpi.com/1424-8220/24/22/7219>

Limitations of Traditional Models


Despite their success, HMM-GMM models have several limitations:



Limited Capacity for Long-Term Dependencies: HMMs can only model short-term dependencies due to their Markovian nature.



Over-reliance on Feature Engineering: models require careful feature extraction and design (e.g., MFCCs), whereas modern systems allow features to be learned from data.



Inability to Capture Complex Speech Patterns: GMMs are linear models and may struggle to capture non-linear relationships in speech signals, particularly in noisy or real-world conditions.

With the advent of deep learning, modern models like Recurrent Neural Networks (RNNs) and Transformers have largely replaced these traditional approaches.

Deep Learning Models in Speech Recognition

Deep learning models have revolutionized speech recognition by automatically learning hierarchical representations from raw audio data.

Recurrent Neural Networks (RNNs):

Overview: RNNs are neural networks designed to model sequential data. They maintain an internal state that can process inputs over time, making them well-suited for speech signals.

Strengths: RNNs can model temporal dependencies, overcoming the short-term limitation of HMMs.

Limitations: standard RNNs suffer from vanishing gradient problems, making them ineffective at learning long-range dependencies.

Long Short-Term Memory (LSTM) Networks:

Overview: LSTMs are a variant of RNNs designed to capture long-term dependencies by using memory cells that control when to store or forget information.

Applications: widely used in modern speech recognition systems for tasks requiring the understanding of temporal context, such as decoding phoneme sequences.

Bidirectional RNNs (Bi-RNNs):

Overview: Bi-RNNs process speech data in both forward and backward directions, capturing context from both past and future frames.

Strengths: provides better performance in speech recognition by capturing more complete contextual information.



Erasmus+

Transformer Models in Speech Recognition

Recently, Transformers have become state-of-the-art for speech recognition, especially in large-scale systems like Google's WaveNet or Whisper from OpenAI. How Transformers work:

- Self-Attention Mechanism: unlike RNNs and LSTMs, Transformers use self-attention to model dependencies between all elements in a sequence, allowing them to capture both short-term and long-term dependencies more effectively.
- Parallelization: transformers can process entire sequences simultaneously, leading to faster training and inference compared to RNN-based models.

Applications:

- End-to-End Speech Recognition (ASR): modern systems like DeepSpeech and Whisper use Transformer-based architectures to transcribe audio directly to text, often bypassing the need for handcrafted features like MFCCs.

Advantages of Transformers:

- Handles Long-Term Dependencies: captures both local and global patterns effectively.
- End-to-End Learning: learns feature representations and mapping to text directly from data.
- Scalability: transformers are highly scalable, especially for large datasets and real-time speech recognition.



Erasmus+

Hybrid Systems and End-to-End Models

Hybrid HMM-DNN Systems:

Overview: a combination of traditional HMMs and deep neural networks (DNNs) where DNNs are used to estimate HMM state probabilities.

Advantage: improved acoustic modeling compared to GMM-HMM systems, while still leveraging the temporal structure of HMMs.

Applications: many large-scale commercial speech recognition systems (like Google's ASR) initially used hybrid architectures.

End-to-End Speech Recognition Systems:

Overview: End-to-end models, such as Connectionist Temporal Classification (CTC) and Attention-based models, directly map speech inputs to text outputs without needing intermediate phoneme or state-level alignments.


Advantages:

- Simpler architecture.
- Requires less domain-specific knowledge.
- Highly scalable and adaptable to various languages and dialects.




Erasmus+


Summary of Speech Recognition Features and Models




Acoustic Features like MFCC and PLP capture the essential characteristics of speech signals but require careful design and tuning.



Traditional Models (HMMs, GMMs) were effective in early speech recognition but have limitations in modeling complex dependencies.



Deep Learning (RNNs, LSTMs, Transformers) has significantly improved speech recognition by learning hierarchical, data-driven representations and handling long-term dependencies.



End-to-End Models represent the future of speech recognition, simplifying the process by directly mapping raw audio to text, achieving state-of-the-art performance in various applications.

Text Recognition: From Characters to Sentences



Politechnika
Śląska




UCZELNIA
BADAWCZA
INICJATYWA INNOWACYJNA


VIBE
PROJECT

Text Recognition: From Characters to Sentences


Text recognition involves converting images of printed, handwritten, or typed text into machine-encoded text. The challenge spans from recognizing individual characters to understanding entire sentences, with applications in document digitization, automatic transcription, and real-time translation. Text recognition systems face various challenges:




Font Variability: different fonts, sizes, and styles can make character identification complex.



Handwriting Recognition: variations in handwriting pose a major challenge, with ambiguities in letter shapes and spacing.



Noise and Distortions: low-quality images, distortions, or background noise can degrade the quality of recognition.



Contextual Understanding: recognizing characters alone is insufficient; understanding words and sentences in context is key to improving accuracy.

Character-Level Recognition

Traditional Feature-Based Approaches:

Edge Detection and Contour Analysis: early OCR systems relied on simple geometric features such as edges, corners, and contours to detect characters. For instance, characters were recognized based on their shapes, such as lines, curves, and intersections.

Template Matching: involves comparing detected shapes against a database of known character templates.

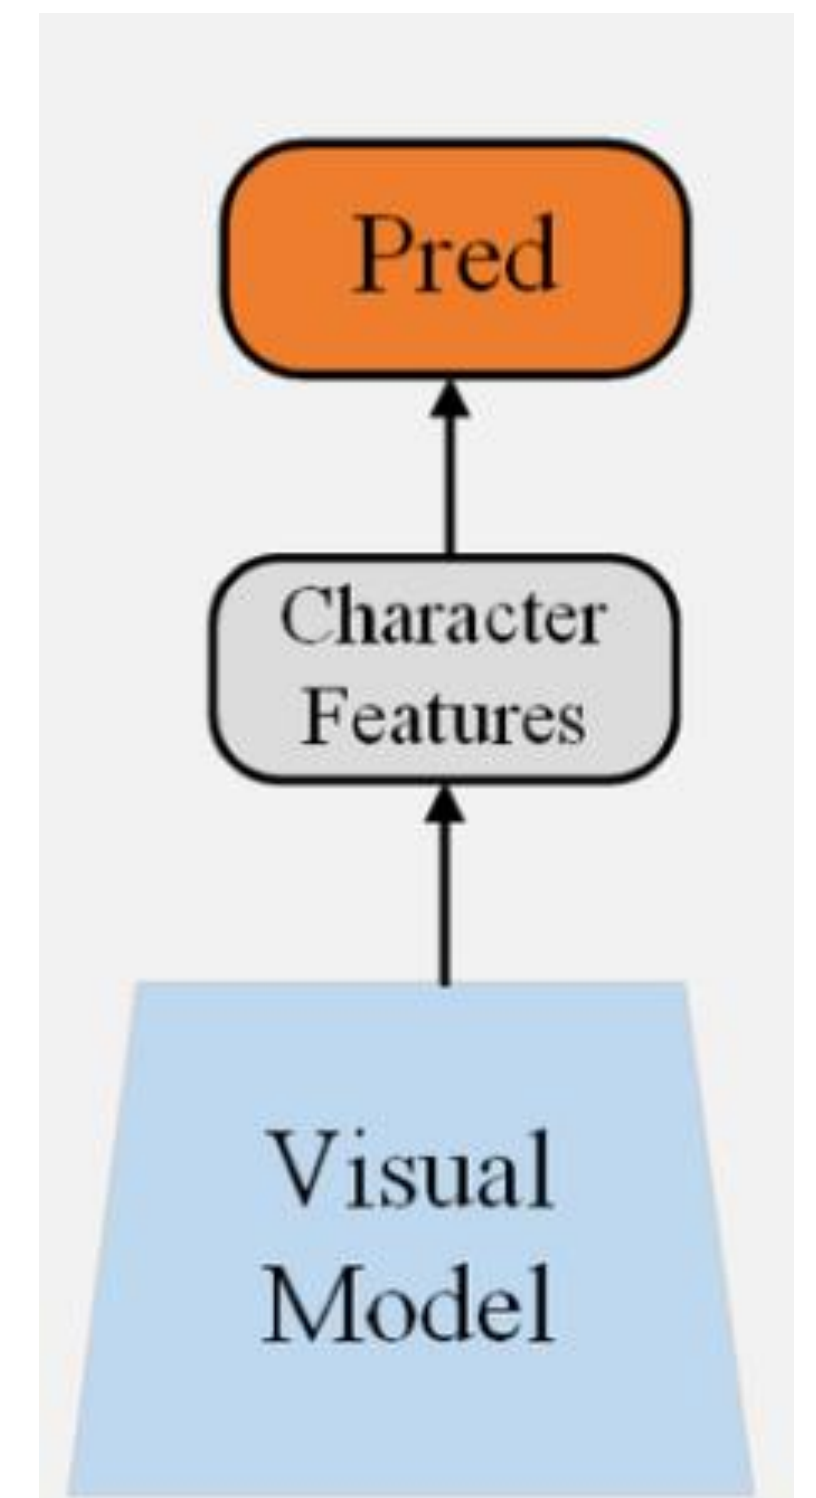
Limitations: methods were highly sensitive to variations in font, style, and noise. They also lacked the ability to adapt to handwriting or complex layouts.

Machine Learning Models:

Support Vector Machines (SVMs) and k-Nearest Neighbors (k-NN) were used in early OCR systems to classify character shapes based on extracted features.

A CNN-based method that explores each character and affinity between characters to effectively detect text areas, even for arbitrarily-shaped text

Handcrafted Features like geometric descriptors, zoning, and projections were used as inputs to these classifiers.




Source:


<https://www.researchgate.net/publication/371785428>
Chinese text recognition enhanced by glyph and character semantic information

Word-Level Recognition


Word-level recognition brings additional complexities, as systems must account for:



Inter-character Spacing: variations in spacing between characters can cause misinterpretation of words, especially in handwritten text.



Ligatures: in cursive handwriting or certain fonts, characters can blend together, requiring sophisticated segmentation algorithms.




Dictionary Matching: some systems use lexicons to improve recognition accuracy by ensuring recognized words match entries in a dictionary


N-gram models and Hidden Markov Models (HMMs) have been used historically to model character sequences for word recognition, ensuring that character combinations form meaningful words.

Sentence-Level Recognition and Context

Recognizing sentences involves more than just identifying individual characters or words. It requires understanding:



Contextual Dependencies: the recognition of one word can influence the interpretation of surrounding words. For example, sentence structure and grammar help resolve ambiguities between similarly shaped characters or words (e.g., “their” vs. “there”).



Language Models: modern text recognition systems use probabilistic language models to improve sentence-level accuracy. These models predict the likelihood of a word or character sequence based on context.

N-gram Language Models: predict the likelihood of a word given the preceding n -words, providing context for recognition.

Markov Models: extend the basic n-gram approach by modeling sentence structures probabilistically.

Grammar-based Models: incorporate grammatical rules to correct errors in sentence recognition.



Erasmus+

Deep Learning in Text Recognition

Deep learning has transformed text recognition by eliminating the need for handcrafted features and allowing models to learn directly from data.

Convolutional Neural Networks (CNNs):

How it works: CNNs are used to automatically extract visual features from images of text, learning the shapes and patterns of characters through layers of convolutional filters.

Applications: CNNs are used for recognizing printed text, handwritten text, and even distorted or noisy text.

Recurrent Neural Networks (RNNs):

How it works: RNNs, especially Long Short-Term Memory (LSTM) networks, are well-suited for modeling sequential data, such as character sequences in words or sentences.

Role in OCR: RNNs help to maintain context across longer sequences, ensuring that text is interpreted in context rather than as isolated characters.

End-to-End Models:

Overview: modern OCR systems like Google's Tesseract or the Transformer-based OCR use deep learning models that process entire text regions in images without relying on prior segmentation of individual characters.

Applications: End-to-end models handle complex layouts, handwritten notes, and text in real-world scenes (e.g., street signs, documents).



Politechnika
Śląska



UCZELNIA
BADAWCZA

INICJATYWA INNOWACYJNA

VIBE
PROJECT



Erasmus+

Attention Mechanisms in Text Recognition

Attention mechanisms have brought significant improvements to text recognition, especially in sentence-level understanding.

Self-Attention (Transformers):

How it works: self-attention mechanisms, like those used in Transformers, allow the model to focus on relevant parts of the input when recognizing a word or sentence. Instead of processing text in a strictly sequential manner, the model attends to all parts of the sequence simultaneously.

Impact: Transformers can capture long-range dependencies more effectively than RNNs, making them highly accurate for recognizing long or complex text sequences.

Sequence-to-Sequence Models (Seq2Seq):

How it works: models process input sequences (e.g., images of text) and generate corresponding output sequences (e.g., transcribed sentences), often using an encoder-decoder architecture with attention. This allows for context-aware sentence recognition.

Applications: Seq2Seq models are used in real-time translation, automatic subtitle generation, and document digitization.

Applications of Text Recognition


Text recognition is a fundamental technology across various fields:

- Document Digitization:
 - used in libraries, archives, and businesses to digitize paper documents and make them searchable.
- Handwriting Recognition:
 - applied in fields like postal address reading, signature verification, and digital note-taking apps (e.g., Microsoft OneNote, Google Keep).
- Real-Time Translation and Accessibility:
 - systems like Google Translate use OCR for instant text translation from images.
 - screen readers and assistive technology use OCR to convert written text into speech for visually impaired users.
- License Plate and Identity Recognition:
 - OCR is used in surveillance, security systems, and smart cities to automatically detect and transcribe license plates or identity documents.




Erasmus+


Summary: From Characters to Sentences




Text recognition involves converting images of text into machine-readable characters, words, and sentences.




Character-level recognition: traditional methods relied on handcrafted features and geometric analysis, while modern systems use CNNs for automatic feature extraction.



Word and sentence recognition: RNNs, LSTMs, and language models improve recognition accuracy by incorporating context and sequential dependencies.



Deep learning and attention mechanisms: these have revolutionized the field, enabling end-to-end models capable of handling complex, real-world text recognition tasks.



Applications: from document digitization to real-time translation and accessibility, text recognition plays a crucial role in various industries and daily applications.

Pattern Recognition in Multimodal Systems



Politechnika
Śląska



UCZELNIA
BADAWCZA
INICJATYWA INNOWACYJNA

VIBE
PROJECT



Erasmus+

Pattern Recognition in Multimodal Systems

Multimodal systems integrate and process information from multiple sensory or data modalities (e.g., images, speech, text, and video). These systems recognize patterns across different types of data, providing a more holistic understanding and improving performance in tasks like recognition, classification, and interaction. A multimodal system processes data from more than one modality to recognize patterns or make decisions. Common modalities include:


- visual (images, video),
- auditory (speech, sound),
- text (language, Natural Language Processing),
- tactile (touch sensors, haptics),
- other (biosignals, sensor data).

Examples of Multimodal Systems:


- Autonomous Vehicles: combine vision (cameras), lidar (depth sensing), and GPS data to navigate.
- Virtual Assistants: use speech recognition, natural language processing (NLP), and visual inputs for human-computer interaction.
- Healthcare Systems: analyze medical images, patient history (text), and sensor data (e.g., heart rate, EEG) to diagnose conditions.

Multimodal Recognition advantages


Pattern recognition in multimodal systems offers several advantages:



Complementary Information: each modality provides unique insights; for example, text might capture semantic content while images provide spatial context.




Improved Robustness: systems can handle missing or noisy data from one modality by relying on others, enhancing overall performance.




Contextual Understanding: multimodal fusion allows for a deeper understanding of the situation or task by combining different perspectives (e.g., recognizing emotions from facial expressions and speech tone).

Challenges in Multimodal Recognition


Despite the benefits, multimodal pattern recognition comes with challenges:




Data Heterogeneity: different modalities have varying data formats, structures, and characteristics. Combining image pixels, sound waves, and text sequences into a unified model can be complex.



Synchronization and Alignment: synchronizing inputs (e.g., aligning video frames with audio) is crucial but non-trivial.



Fusion Techniques: deciding how to combine modalities, whether early fusion (merging raw data) or late fusion (merging decisions from unimodal systems), impacts the system's effectiveness.



Data Imbalance: some modalities may provide richer or more frequent information, while others might be sparse or unreliable.

Fusion Techniques in Multimodal Systems

Fusion strategies are crucial to how multimodal systems integrate information:

Early Fusion (Feature-Level Fusion):

How it works: data from multiple modalities are combined at the feature extraction level, before any pattern recognition takes place.

Advantages: captures interactions between modalities early on, which can lead to more comprehensive feature representations.

Challenges: requires data from different modalities to be well-aligned and represented in a common format (e.g., vectors).

Late Fusion (Decision-Level Fusion):

How it works: each modality is processed independently, and the results are merged at the decision stage (e.g., voting or averaging predictions).

Advantages: more flexible and resilient to modality-specific noise or failure.

Challenges: fails to capture cross-modal interactions during feature learning, which can limit performance.

Hybrid Fusion:

How it works: combines both early and late fusion, leveraging the strengths of both approaches. Feature extraction happens for each modality, followed by cross-modal interactions before final decision-making.

Applications: Used in complex systems like autonomous vehicles or multimodal biometric systems.

Cross-Modal Learning

Cross-modal learning involves using data from one modality to improve learning in another modality. This can be particularly powerful when one modality has limited data, or when the modalities have complementary information.

Cross-Modal Representation Learning:

Objective: learn representations that capture the commonalities between different modalities, enabling joint understanding.

Techniques:

- Canonical Correlation Analysis (CCA): a traditional method for finding relationships between two modalities by maximizing their correlation in a shared latent space.
- Deep Cross-Modal Networks: neural networks that align multimodal data through shared layers, learning joint feature representations.

Zero-Shot or Few-Shot Learning:

Objective: Recognize patterns in one modality based on information learned from another, even without direct examples (e.g., learning to recognize objects in images based on their textual descriptions).

Example: Visual Question Answering (VQA), where the system answers questions about images using both text and visual data.

Multimodal Learning Architectures

Modern multimodal systems use advanced architectures to capture and fuse information:

Recurrent Neural Networks (RNNs) and LSTMs:

Use case: for sequential data like video (frames) and speech, LSTMs model temporal dependencies, helping align and fuse data from different modalities.

Applications: in video analysis with synchronized audio, LSTMs fuse temporal data for better event detection.

Convolutional Neural Networks (CNNs):

Use case: CNNs are typically used for spatial data like images but can be adapted for multimodal fusion. In multimodal systems, CNNs extract spatial features from images and pair them with text or audio features.

Applications: used in systems that analyze both image and text, such as visual search engines or image captioning systems.

Transformer-Based Models:

Use case: Transformers excel at processing sequences (text, audio, video) and have become popular for multimodal fusion due to their ability to attend to data across different modalities simultaneously.

Applications: used in models like OpenAI's CLIP or Google's Vision-Language models, where both text and images are processed together to make sense of the context.



Multimodal Pattern Recognition Applications


Multimodal systems have wide-ranging applications, each benefiting from combining different forms of data:

Autonomous Vehicles: Overview: vehicles use cameras (vision), radar (depth), lidar (range sensing), and GPS (location) data to recognize obstacles, navigate, and interact with the environment. Example: Tesla’s self-driving system and Google’s Waymo combine multimodal inputs for driving decisions.	Healthcare Systems: Overview: multimodal learning helps in diagnostics by combining medical imaging (CT scans, X-rays), patient history (text), and biosignals (heart rate, EEG). Example: cancer detection systems combining imaging data with genetic analysis and patient records for accurate diagnosis.	Virtual Assistants: Overview: assistants like Alexa, Siri, or Google Assistant use multimodal inputs—speech for commands, NLP for understanding context, and sometimes visual data (like images or gestures) for interaction. Example: Amazon Echo devices with screens (Echo Show) integrate voice and video for richer user interaction.	Multimodal Emotion Recognition: Overview: recognizing human emotions by combining facial expressions (visual), tone of voice (audio), and text (words spoken). Example: used in customer service systems to detect frustration or satisfaction.	Augmented Reality (AR) and Virtual Reality (VR): Overview: AR/VR systems combine visual, auditory, and sometimes haptic feedback to create immersive experiences. Example: medical training simulators use multimodal inputs to provide realistic surgical training with visual, auditory, and tactile feedback.
--	---	---	--	---




Erasmus+


Summary of Multimodal Systems




Multimodal systems integrate data from multiple sources (e.g., vision, speech, text) to enhance pattern recognition, creating richer, more reliable models.



Fusion strategies like early, late, and hybrid fusion help in combining multimodal data.



Cross-modal learning enables systems to leverage knowledge from one modality to improve another.



Applications: range from autonomous vehicles to healthcare, virtual assistants, and emotion recognition, where multimodal inputs significantly improve accuracy and performance.

Evaluation and Error Analysis in Pattern Recognition





Erasmus+

Evaluation and Error Analysis in Pattern Recognition

Evaluation and error analysis are critical steps in developing robust pattern recognition systems. They provide insight into the system's performance, help identify weaknesses, and guide improvements. Evaluation ensures that a pattern recognition model:

- measures accuracy: verifies the model's ability to correctly recognize patterns (e.g., classification accuracy, recognition rate),
- identifies weaknesses: highlights the areas where the model underperforms, such as specific classes or scenarios,
- guides improvements: points to potential optimizations in feature extraction, model architecture, or data preprocessing,
- validates generalization: confirms that the model works well not only on training data but also on unseen test data (generalization capability).

Systematic evaluation provides benchmarks for future enhancements and ensures reliability in real-world deployments.

Key Evaluation Metrics

Several metrics are used to evaluate pattern recognition systems, depending on the problem domain:

Accuracy:

- measures the proportion of correct predictions among all predictions,
- formula:

$$\text{Accuracy} = \frac{\text{True Positives (TP)} + \text{True Negatives (TN)}}{\text{Total Predictions}}$$

Precision:

- proportion of true positive results among all positive predictions,
- important when false positives are costly,
- formula:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{False Positives (FP)}}$$

Recall (Sensitivity):

- proportion of actual positives correctly identified,
- critical when false negatives are costly (e.g., in medical diagnostics),
- formula:

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{False Negatives (FN)}}$$

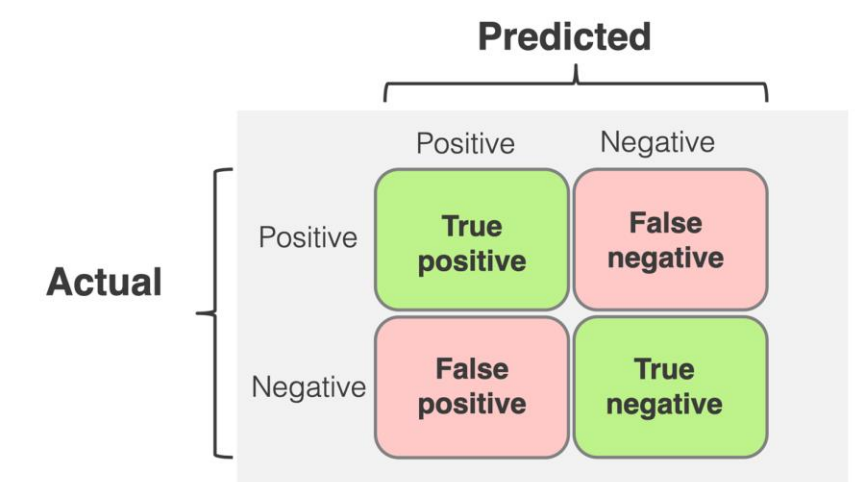
F1-Score:

- harmonic mean of precision and recall, balancing both,
- used when you need a balance between precision and recall,
- formula:

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

Confusion Matrix:

- a matrix showing true positives, true negatives, false positives, and false negatives for each class, allowing a detailed view of performance across categories.



Source:

<https://link.springer.com/article/10.1007/s11469-024-01343-8>




Erasmus+

Receiver Operating Characteristic (ROC) Curve

The ROC curve is a graphical representation of a model's performance across different classification thresholds, balancing between true positive and false positive rates:

- True Positive Rate (TPR) or Recall: the fraction of positive instances correctly classified as positive,
- False Positive Rate (FPR): the fraction of negative instances incorrectly classified as positive.

Area Under the Curve (AUC) measures the overall ability of the model to distinguish between classes. A higher AUC indicates better model performance across various thresholds.



Source: <https://link.springer.com/article/10.1245/s10434-008-0133-6>

Source: <https://www.fastercapital.com/keyword/curve-auc.html/2>



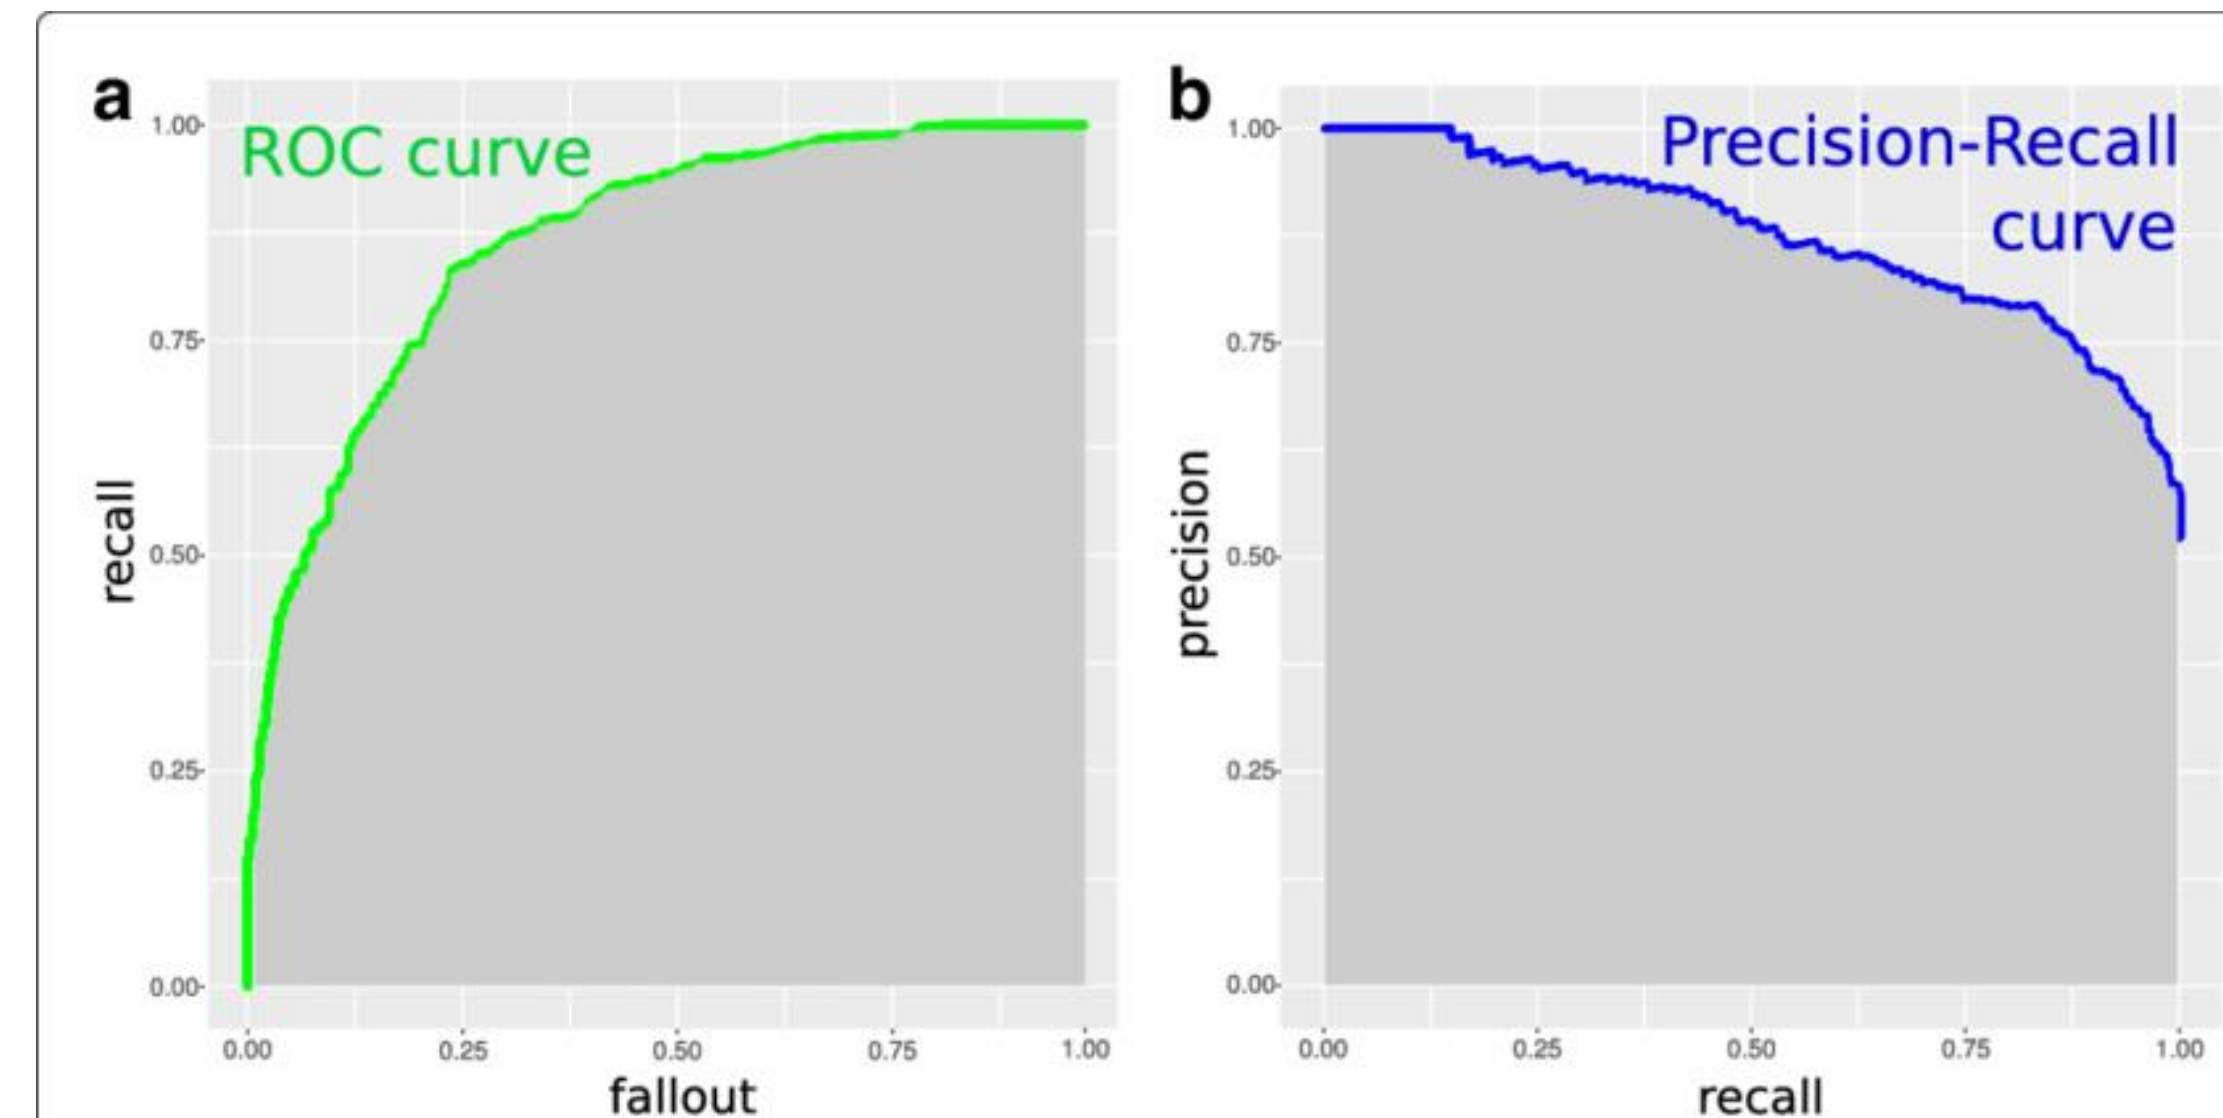
Erasmus+

Precision-Recall Curve

While ROC curves are useful for binary classification, the Precision-Recall (PR) curve is more informative when dealing with imbalanced datasets, where one class significantly outnumbers the other:

- Precision: measures how many of the predicted positive instances are actually positive,
- Recall (Sensitivity): measures how many actual positive instances are correctly identified.

The PR curve focuses on the trade-off between precision and recall, making it ideal for evaluating models in scenarios like rare event detection or medical diagnosis, where true positives are far fewer than negatives.



Source:

https://www.researchgate.net/publication/321572019_Ten_quick_tips_for_machine_learning_in_computational_biology

CV



Erasmus+

Error Analysis in Pattern Recognition

Error analysis helps identify and understand the reasons behind the model's incorrect predictions. Key steps include:

- Confusion Matrix Analysis:
 - identify which classes are most frequently confused,
 - spot patterns in false positives (incorrectly classified as positive) and false negatives (missed detections).
- Bias and Variance:
 - High Bias: Underfitting—model is too simple and fails to capture the complexity of the data,
 - High Variance: Overfitting—model fits the training data too closely and fails to generalize to new data,
 - solutions include regularization, increasing data diversity, or simplifying the model.
- Class Imbalance:
 - Imbalance in data distribution can cause a model to favor the majority class. Address by using techniques like oversampling, undersampling, or class-weighted loss functions to give higher priority to the minority class.
- Error Attribution:
 - categorize errors by source: Is the error due to poor feature extraction, incorrect labeling, noise in the data, or model limitations?



Erasmus+

Common Sources of Errors

Several common sources of errors exist in pattern recognition systems:

- Noise in Data:
 - real-world data is often noisy, whether it's blurry images, overlapping speech, or poorly scanned text,
 - preprocessing techniques like denoising, filtering, or data augmentation help mitigate this.
- Misleading or Poor Features:
 - features that don't capture relevant patterns can lead to poor recognition. Feature selection or engineering techniques should be applied to ensure meaningful patterns are captured.
- Model Complexity:
 - too simple models (underfitting) fail to capture complex patterns, while too complex models (overfitting) can memorize the training data rather than generalizing.
- Domain Shift:
 - when training data and test data come from different distributions (e.g., training on daytime images but testing on nighttime images), performance can degrade significantly. Domain adaptation or transfer learning techniques help address this.



Erasmus+

Strategies to Improve Performance

Once errors have been identified, various strategies can be applied to improve the system's performance:

- Data Augmentation:
 - expand the training dataset by applying transformations (e.g., rotations, scaling, adding noise) to improve generalization.
- Regularization:
 - techniques like L1/L2 regularization, dropout, or early stopping can help reduce overfitting and improve the model's generalization capability.
- Cross-Validation:
 - use k-fold cross-validation to get a more reliable estimate of the model's performance, ensuring it generalizes well across different subsets of the data.
- Error Corrective Feedback:
 - use active learning or human-in-the-loop systems to correct and retrain the model based on errors in real-time, improving the model iteratively.



Erasmus+

Advanced Error Analysis Techniques

Beyond simple error tracking, advanced techniques help further diagnose and fix issues in pattern recognition models:

Class Activation Mapping (CAM):

- helps identify which parts of an image a CNN model focuses on when making decisions, aiding in diagnosing misclassifications due to irrelevant or noisy features.

Saliency Maps:

- visualize what parts of the input (image, text, or speech) contributed most to the model's decision, giving insights into why errors occur.

Sensitivity Analysis:

- measures the impact of changing certain features or input values on the model's predictions, helping identify weak or unreliable features.


Outlier Detection:

- identifying and addressing outliers can reduce errors. Techniques like isolation forests or autoencoders are useful for detecting anomalies that degrade model performance.




Erasmus+


Summary: Evaluation and Error Analysis




Evaluation metrics (accuracy, precision, recall, F1-score, ROC, PR curve) are essential for assessing a model's performance in different contexts.



Error analysis allows you to pinpoint the sources of errors, such as misclassification, bias-variance trade-off, or data quality issues.



Addressing errors through feature refinement, model regularization, or advanced analysis techniques can significantly enhance the performance of the pattern recognition system.



Continuous improvement: evaluating and iteratively analyzing errors leads to more robust and generalizable models, capable of performing well in diverse real-world environments.

Examples of Pattern Recognition Systems



Politechnika
Śląska



UCZELNIA
BADAWCZA
INICJATYWA INNOWACYJNA

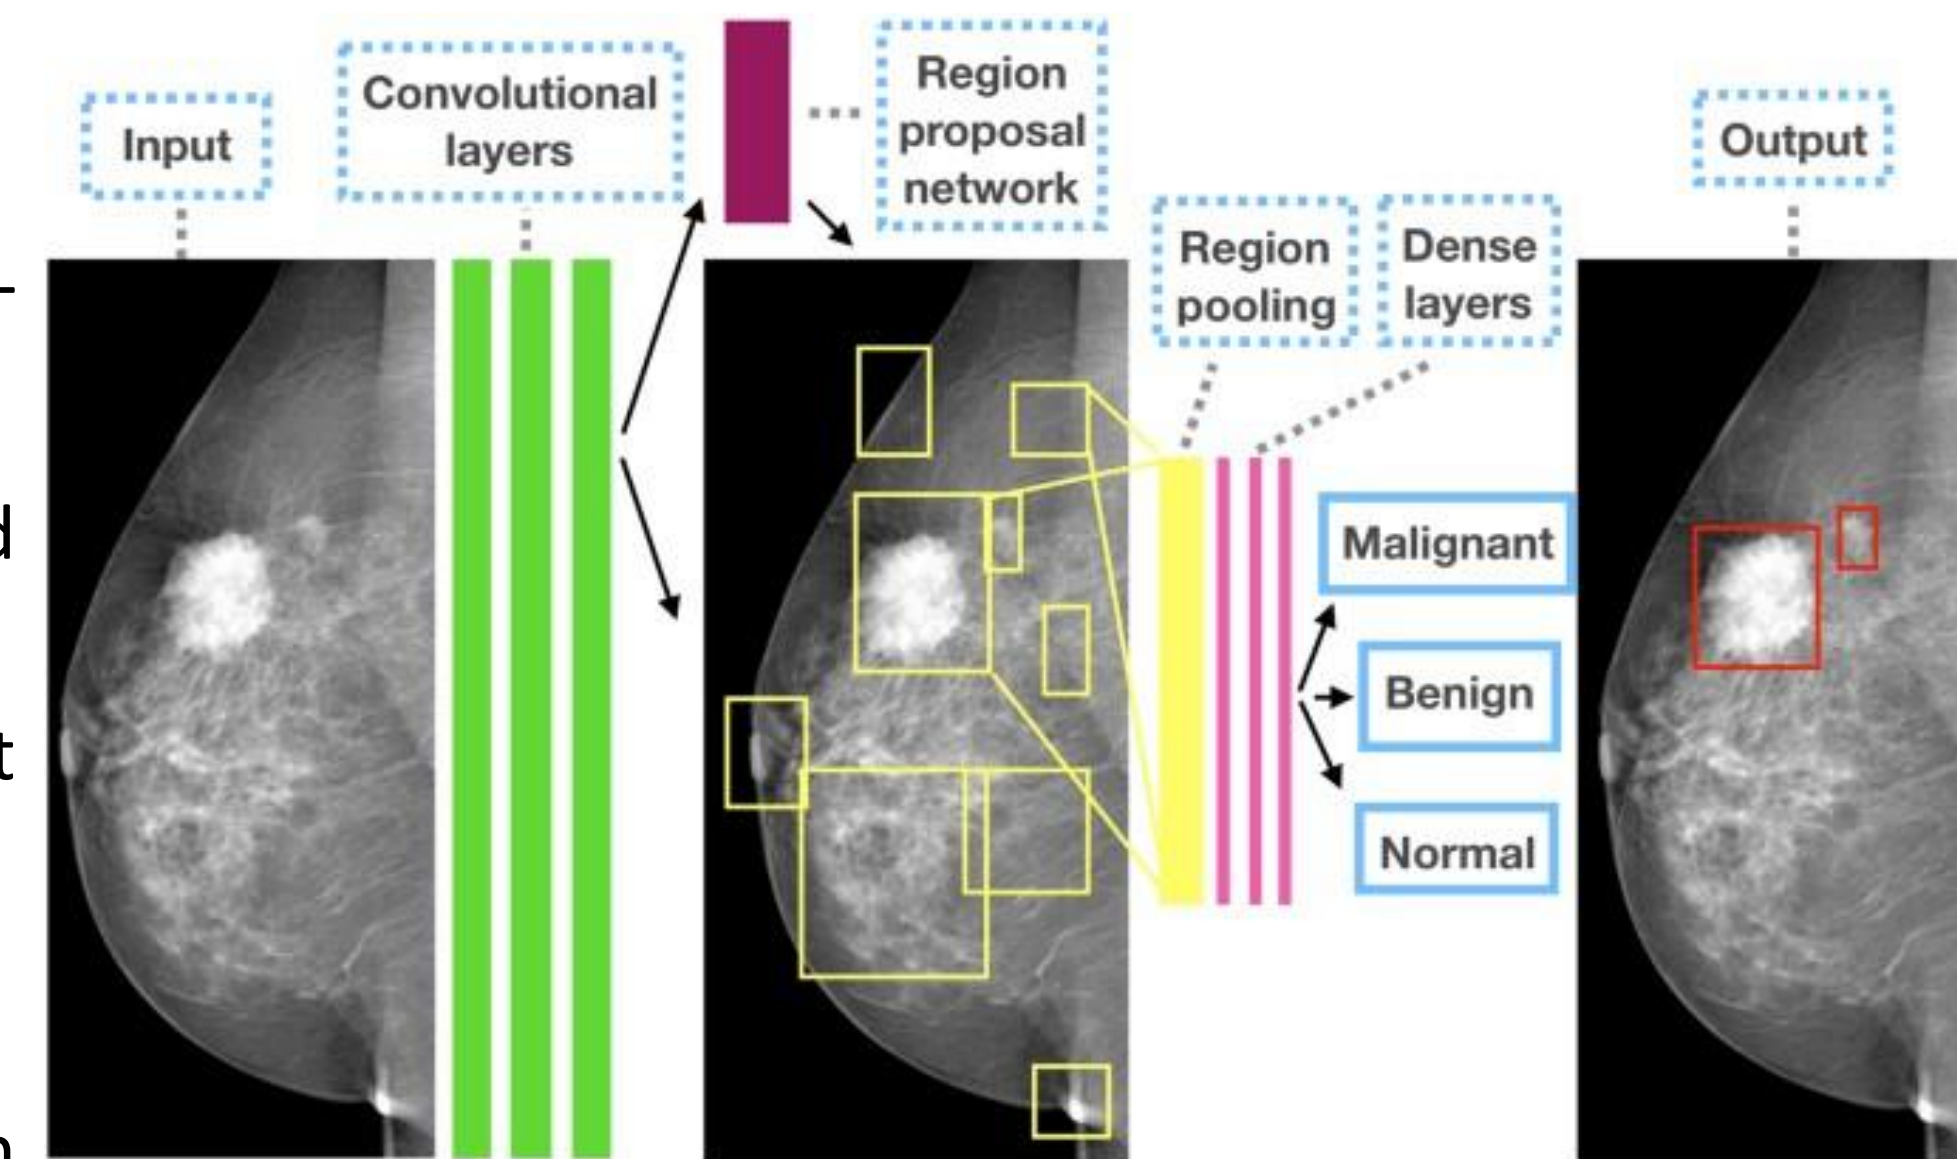
VIBE
PROJECT

Example 1: Image Recognition in Medical Diagnostics

Application: Medical Imaging (e.g., cancer detection, organ segmentation):

- **Model:** Convolutional Neural Networks (CNNs),
- **Task:** identifying tumors or lesions in medical images (e.g., X-rays, MRI, or CT scans),
- **Key Features:** CNNs automatically extract spatial patterns, such as edges and textures, and recognize abnormalities,
- **Impact:** early and accurate diagnosis of diseases, aiding doctors in treatment decisions and reducing diagnostic errors.

Example: deep learning models trained on large datasets of mammograms can detect breast cancer with a level of accuracy comparable to expert radiologists.



Source:

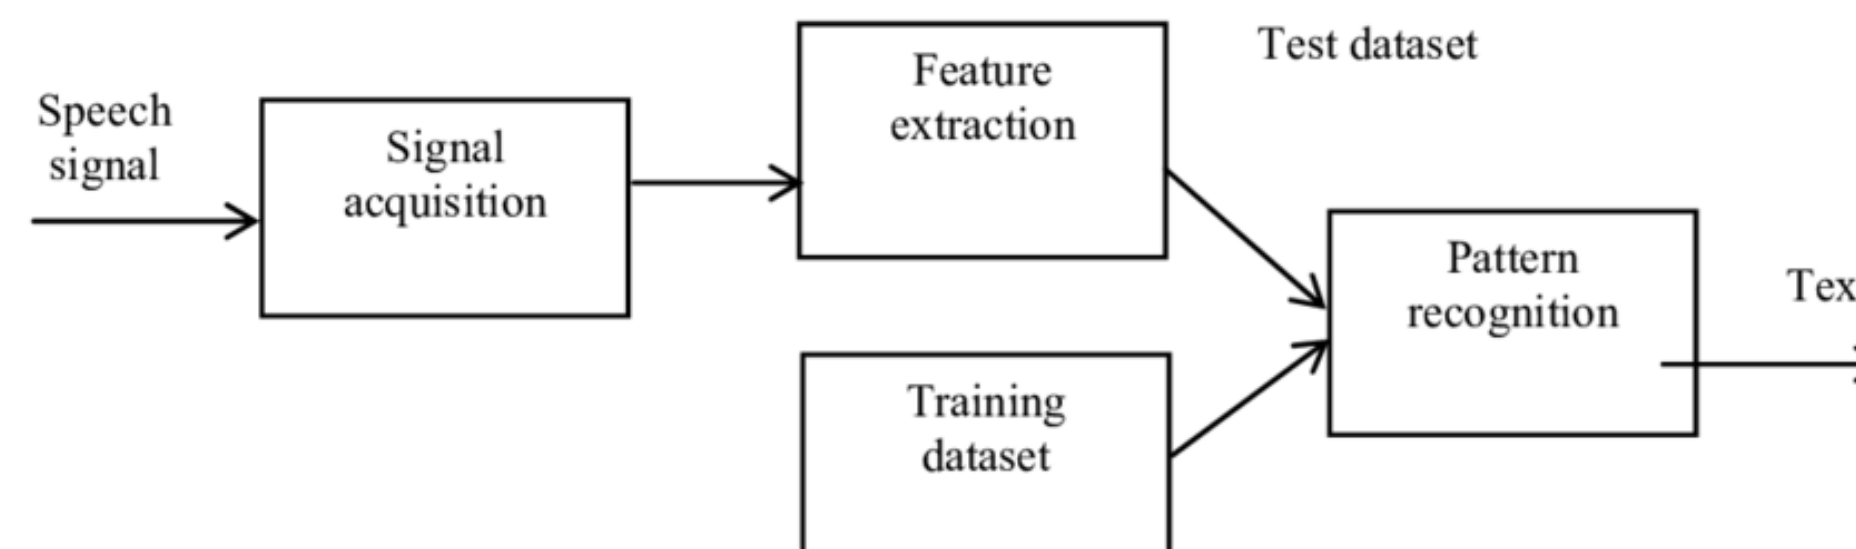
<https://www.researchgate.net/publication/352932542> Cancer detection using deep learning techniques

Example 2: Speech Recognition for Virtual Assistants

Application: Voice-Activated Virtual Assistants (e.g., Siri, Google Assistant, Alexa)

- **Model:** Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) networks or Transformer-based architectures,
- **Task:** convert spoken language into text (Automatic Speech Recognition - ASR) and understand user commands,
- **Key Features:** models capture the temporal dependencies in speech data and are robust to variations in accents and noise,
- **Impact:** improved accessibility, hands-free device control, and personalized user experiences.

Example: Google's speech recognition system boasts over 95% accuracy and supports numerous languages and dialects, making it a global tool for communication.



Source: https://www.researchgate.net/publication/338761069_Recurrent_neural_network-based_speech_recognition_using_MATLAB



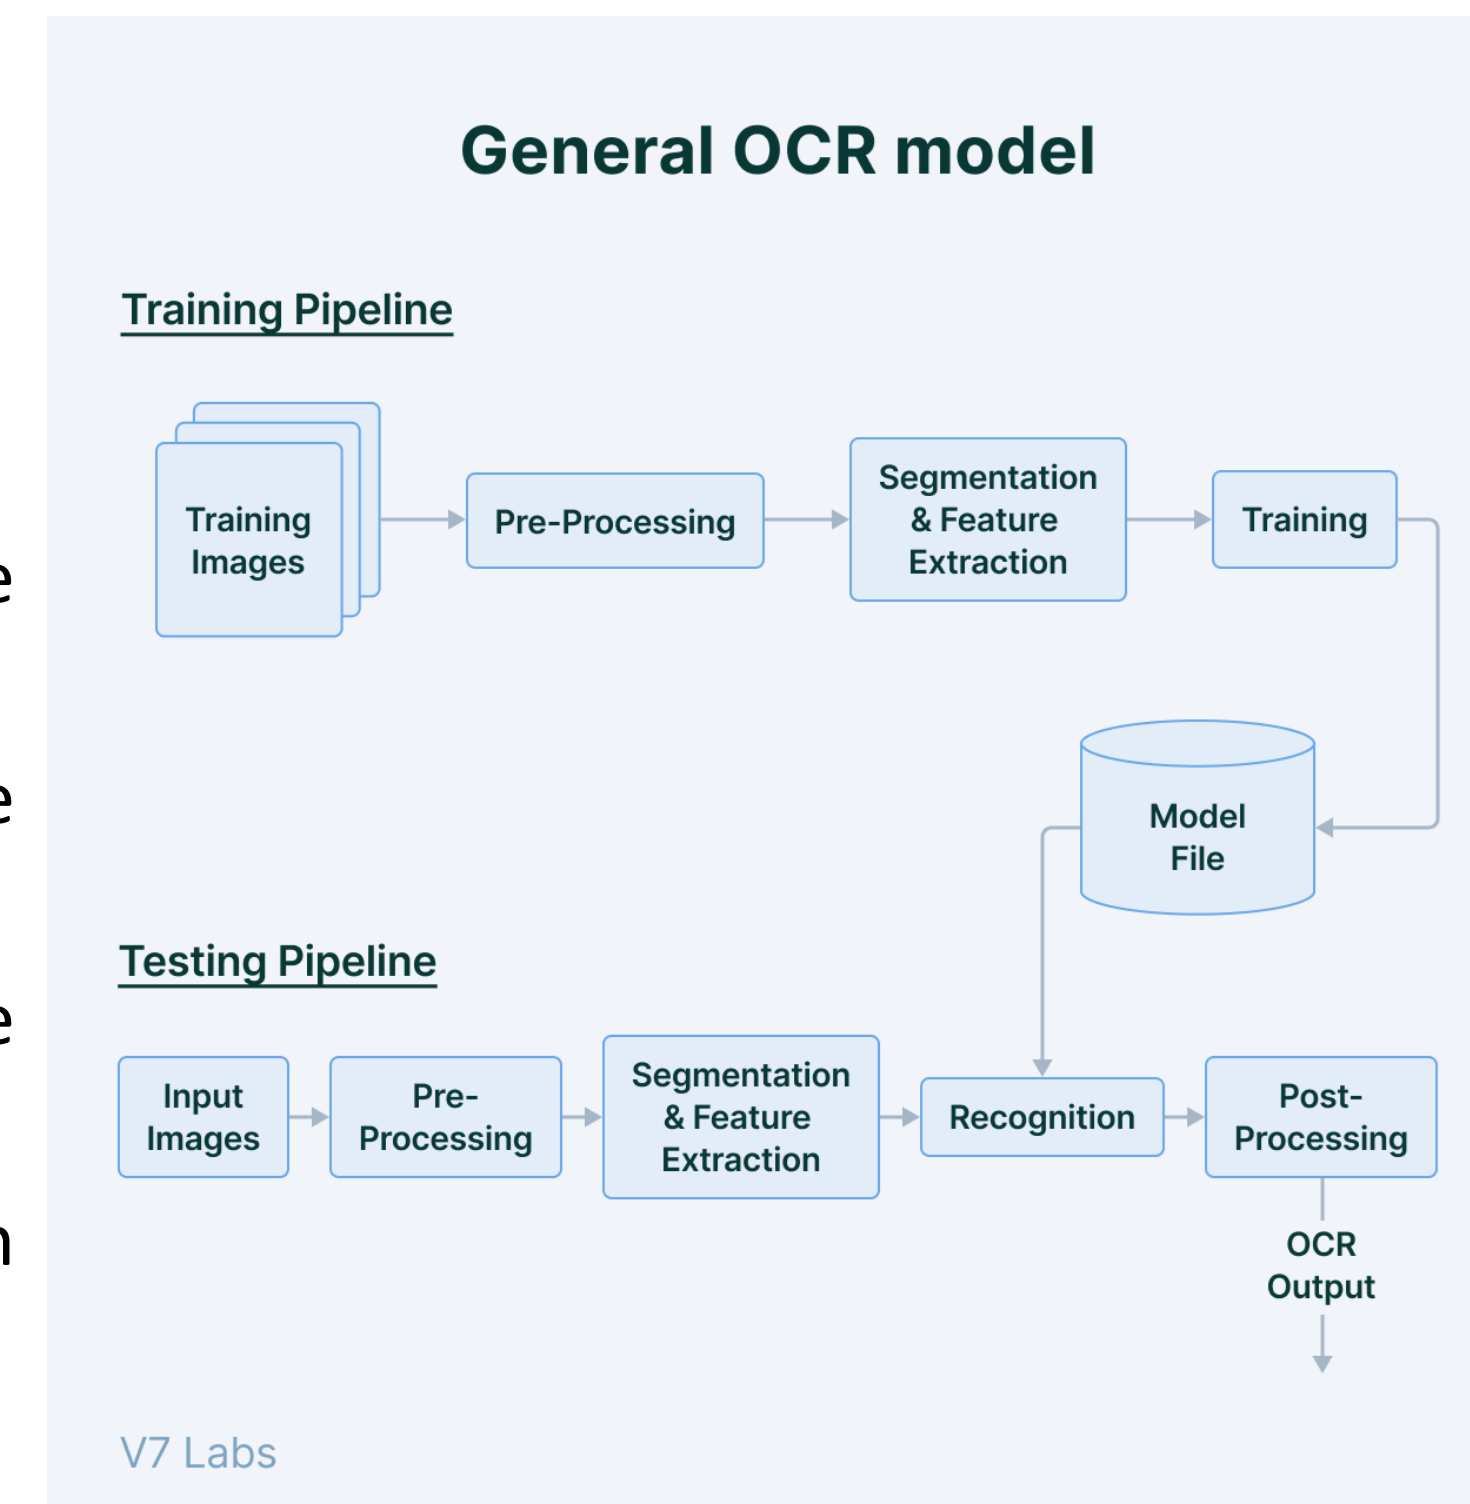
Erasmus+

Example 3: Optical Character Recognition (OCR) for Document Processing

Application: Document Digitization and Automation (e.g., financial documents, legal contracts)

- **Model:** Optical Character Recognition (OCR) systems using CNNs combined with sequence models (LSTM, Transformer),
- **Task:** extract text from scanned documents or images and convert it into machine-readable text,
- **Key Features:** models handle varying fonts, layouts, and image quality, and recognize characters and words even in noisy or distorted images,
- **Impact:** automates the processing of large volumes of documents, increasing efficiency in industries like banking, insurance, and law.

Example: OCR systems used by banks automatically extract data from checks or invoices, significantly reducing manual entry and processing time.



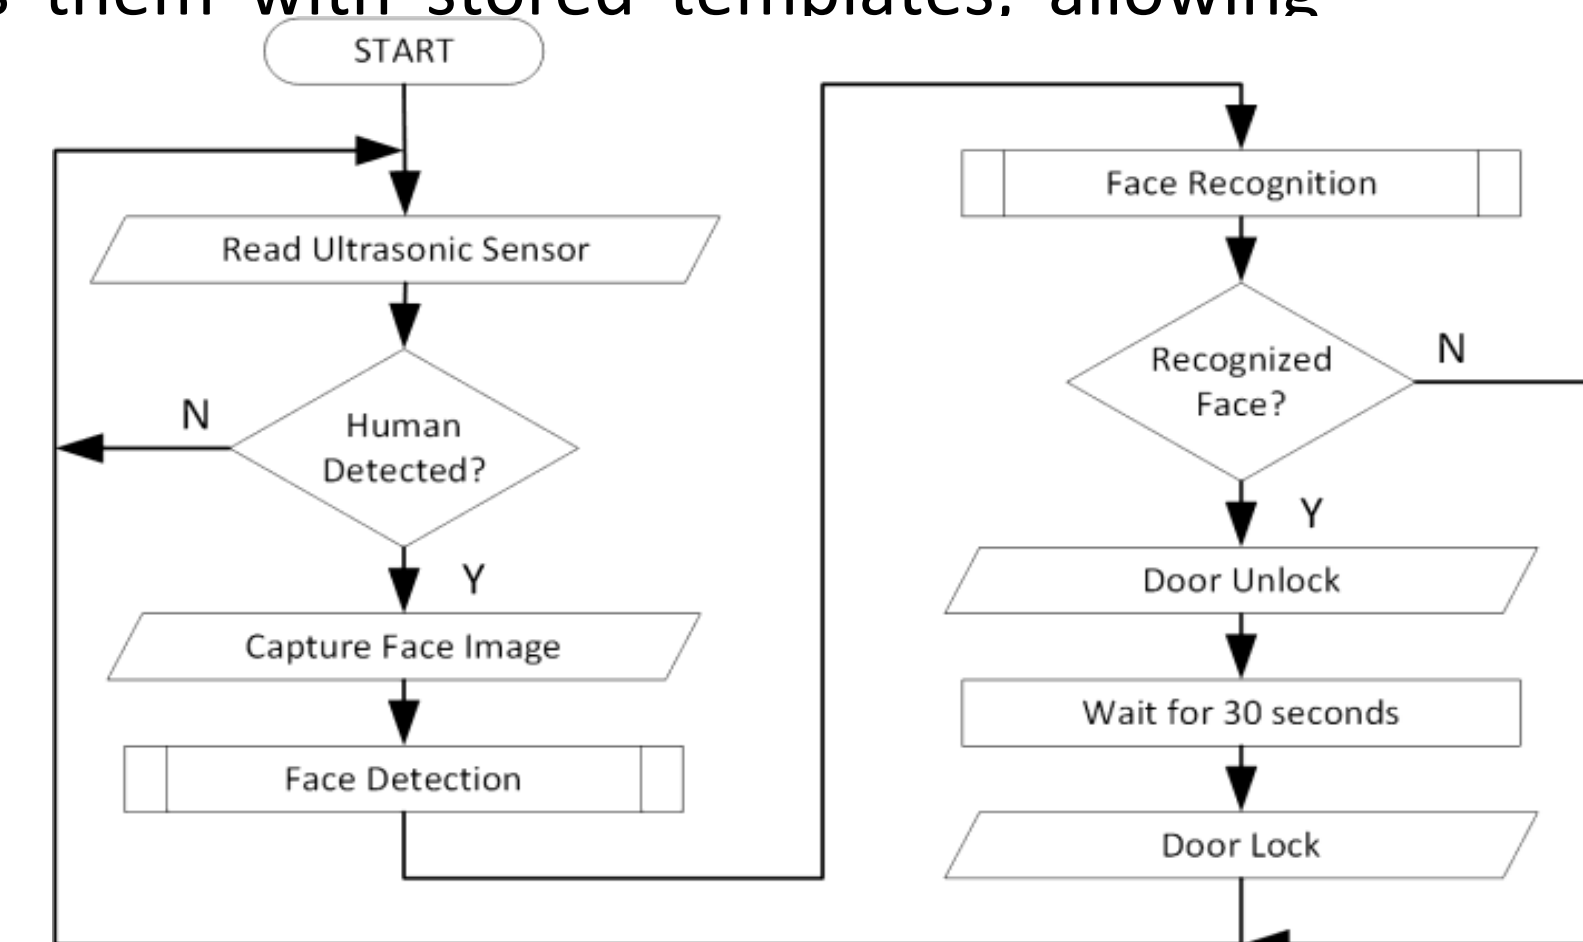
Source: <https://www.ijraset.com/research-paper/revolutionizing-data-entry-an-in-depth-study-of-optical-character-recognition-technology>

Example 4: Facial Recognition for Security Systems

Application: Security and Surveillance (e.g., facial recognition at airports, secure access control)

- **Model:** CNNs, often enhanced with deep metric learning or face embedding techniques (e.g., FaceNet),
- **Task:** identify or verify individuals from images or video feeds by comparing facial features,
- **Key Features:** the system extracts high-dimensional facial embeddings and compares them with stored templates, allowing recognition across varied conditions (lighting, angles, occlusions),
- **Impact:** increased security in sensitive areas, faster identification, and verification processes.

Example: airports and government agencies use facial recognition systems to streamline border control and identify persons of interest in real-time surveillance.



Source:

https://www.researchgate.net/publication/322009772_Development_of_Face_Recognition_on_Raspberry_Pi_for_Security_Enhancement_of_Smart_Home_System



Erasmus+

Example 5: Multimodal Recognition for Autonomous Vehicles

Application: Autonomous Driving (e.g., self-driving cars)

- **Model:** multimodal systems that combine CNNs for image recognition (vision) and LSTMs or transformers for sensor fusion (e.g., lidar, radar, GPS),
- **Task:** recognize objects (pedestrians, vehicles, traffic signals), understand the driving environment, and make real-time decisions,
- **Key Features:** integration of multiple data streams (visual, spatial, and temporal data) enhances situational awareness, enabling safe navigation,
- **Impact:** autonomous vehicles rely on real-time recognition for safe and efficient operation in dynamic environments.

Example: Tesla's self-driving cars use multiple cameras and sensors to detect road signs, lanes, and obstacles, allowing autonomous navigation in complex urban settings.

Case study —



Politechnika
Śląska



UCZELNIA
BADAWCZA
INICJATYWA INNOWACYJNA

VIBE
PROJECT



Erasmus+

Case Study 1: Image Recognition for Early Disease Detection in Healthcare

Introduction: in healthcare, accurate and early diagnosis is critical for effective treatment. Image recognition technologies have enabled medical professionals to analyze complex imaging data, such as MRI or CT scans, for early disease detection, particularly in oncology and radiology.

Problem Statement: Delayed Diagnosis in Oncology

Objective: develop an image recognition system capable of accurately identifying early-stage tumors in medical images, thereby reducing diagnostic delays and improving patient outcomes.



Erasmus+

Case Study 1: Image Recognition for Early Disease Detection in Healthcare

Proposed Solution: Deep Learning-Based Image Analysis

A deep learning-based image recognition model using Convolutional Neural Networks (CNNs) was developed to automatically detect tumors in MRI and CT scans. The model was trained on a large dataset of labeled medical images, learning to distinguish between normal and abnormal tissues.

Key Features:

- **Automated Feature Extraction:** CNNs automatically learn spatial features, such as edges and textures, without the need for manual intervention,
- **Accuracy:** the model was fine-tuned to minimize false positives and false negatives, ensuring high accuracy in tumor detection,
- **Scalability:** the system can process large volumes of imaging data quickly, enabling faster diagnosis.



Erasmus+

Case Study 1: Image Recognition for Early Disease Detection in Healthcare

The deployment of the CNN-based system in a clinical setting resulted in:

- **Increased diagnostic accuracy:** the model achieved an accuracy rate of 94% in identifying early-stage tumors, outperforming traditional methods,
- **Reduced diagnostic time:** the automated system reduced the time required for radiologists to analyze images by 50%, enabling faster treatment decisions,
- **Fewer false positives:** through careful training and validation, the system minimized the number of false alarms, increasing confidence in diagnoses.

Impact: hospitals using the system saw a 20% improvement in early cancer detection rates, leading to better patient outcomes and more effective treatment plans.



Erasmus+

Case Study 2: Speech Recognition in Customer Service Automation

Introduction: speech recognition technology has revolutionized customer service by automating routine inquiries and enabling businesses to handle large volumes of customer requests efficiently. This case study examines the implementation of a speech recognition system for a major telecommunications company.

Problem Statement: High Call Volume in Customer Support

Objective: implement a speech recognition system to automate frequently asked questions (FAQs) and direct customers to appropriate resources, reducing the burden on human agents.



Erasmus+

Case Study 2: Speech Recognition in Customer Service Automation

Proposed Solution: Automatic Speech Recognition (ASR) System

An Automatic Speech Recognition (ASR) system was deployed, using a combination of deep learning models like Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks for speech-to-text conversion and Natural Language Processing (NLP) models for intent recognition.

Key Features:

- Real-time transcription: the ASR system accurately converted customer queries into text in real time, identifying key phrases and intents,
- Intent recognition: NLP models classified customer requests, routing them to the appropriate automated response or human agent,
- 24/7 Availability: the system provided round-the-clock support, significantly improving customer experience.



Erasmus+

Case Study 2: Speech Recognition in Customer Service Automation

After the implementation of the ASR system:

- **Call resolution times** decreased by 40%, as FAQs were automatically addressed without human intervention.
- **Customer satisfaction** increased by 25% due to faster query resolution.
- **Employee productivity** improved: human agents were able to focus on complex and high-priority issues, reducing burnout.

Impact: the system handled 60% of all customer inquiries automatically, greatly reducing operational costs and enabling more efficient use of human resources.



Erasmus+

Case Study 3: Text Recognition in Legal Document Automation

Introduction: in the legal industry, handling large volumes of documents can be time-consuming and prone to error. Text recognition systems powered by OCR (Optical Character Recognition) have significantly streamlined document processing by automating the extraction and analysis of legal texts.

Problem Statement: Manual Document Processing in Law Firms

Objective: develop an OCR-based text recognition system to automate the extraction of key information from legal documents, reducing processing time and improving accuracy.



Erasmus+

Case Study 3: Text Recognition in Legal Document Automation

Proposed Solution: OCR for Legal Document Automation

The OCR system, built using a combination of Convolutional Neural Networks (CNNs) for image recognition and sequence models (LSTMs) for text extraction, was designed to handle complex legal language and varied document formats.

Key Features:

- Text extraction: OCR automatically converts scanned documents into editable and searchable text,
- Key phrase identification: NLP models were used to extract important legal terms, clauses, and dates from contracts and filings,
- Multi-format support: the system could handle various document formats, including PDFs, images, and scanned documents.



Erasmus+

Case Study 3: Text Recognition in Legal Document Automation

The OCR system led to significant improvements in document handling:

- **Document processing time** reduced by 70%, enabling law firms to handle more cases in less time,
- **Error rate decreased:** automated extraction reduced the number of errors associated with manual entry and review,
- **Improved document search:** lawyers could quickly search for relevant terms and clauses, enhancing case preparation efficiency.

Impact: law firms using the system saw operational costs drop by 30%, as automation freed up valuable time for legal professionals to focus on more strategic tasks.

Conclusion —



Politechnika
Śląska



UCZELNIA
BADAWCZA
INICJATYWA INŻYNIERSKA

VIBE
PROJECT



Erasmus+

Conclusion

In this presentation, we've explored the advanced applications of pattern recognition across different domains, including image, speech, and text recognition, as well as their integration into multimodal systems. Key takeaways include:

- Key Mathematical Foundations: understanding the core algorithms, including Maximum Likelihood Estimation (MLE), Maximum A Posteriori (MAP), and Hidden Markov Models (HMMs), is essential for robust pattern recognition systems,
- Feature Engineering: whether in images, speech, or text, selecting the right features is critical for model performance and accuracy,
- Multimodal Systems: combining multiple data types enhances recognition and decision-making capabilities, as seen in fields like autonomous vehicles,
- Error Analysis & Evaluation: thorough evaluation methods ensure that models perform well not just in theory but in real-world applications.
- Recent Trends: emerging technologies like transformers, generative models, and self-supervised learning are reshaping the future of pattern recognition.