SAND2023-05402C

Sandia National Laboratories

Autoencoders for Multimodal Data Fusion

George Fragkos

Renewable & Distributed Systems Integration

U.S. DEPARTMENT OF ENERGY NNSA

Sandia National Laboratories ENERGY NNSA

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

High Dimensionality & Multimodal Data

Classification

- Data classification is a fundamental task in Machine Learning (ML)
- Identify patterns and assign labels to input data
- The predictive performance of classification models is influenced by the number of input features utilized

Challenges

- Multimodal Data: Refers to datasets that combine information from multiple sources or modalities, e.g., text, images, audio, sensor data.
- High Dimensionality: Refers to datasets with a large number of features or variables
- High-dimensional data often arises in multimodal datasets due to the combination of multiple sources
- Each modality contributes its own set of features, resulting in an increased overall dimensionality

High Dimensionality & Multimodal Data

Problem

- Analyzing and extracting meaningful insights from such data require specialized techniques and handling
- As the dimensionality of the input space increases, the complexity of the classification task also increases
- There is a growing interest in reducing the dimensionality of the input space to enhance the predictive performance of classification models

Approaches for dimensionality reduction

PCA
[Principal Component Analysis]

LDA
[Linear Discriminant Analysis]

Assumption: The data lies on a linear subspace

Not always true in real-world scenarios!

Linear vs nonlinear dimensionality reduction

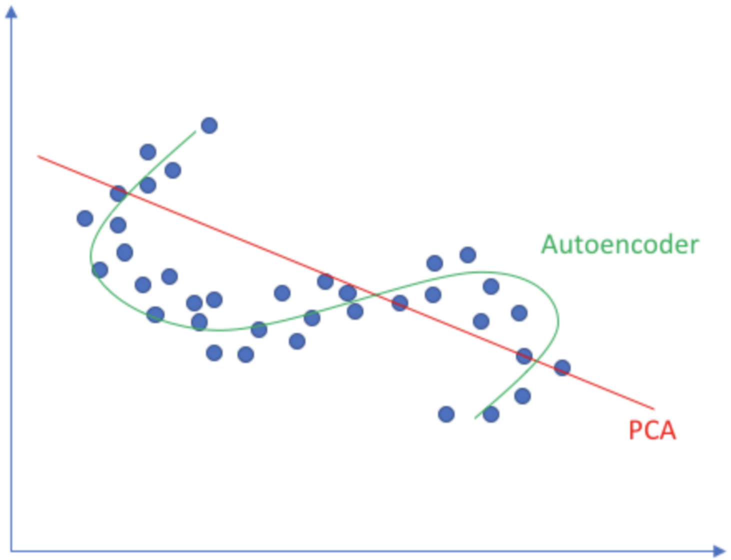

Manifold Learning & Data Fusion

- It is a dimensionality reduction technique used to understand the underlying structure of complex high-dimensional data
- It aims to uncover the intrinsic low-dimensional manifold on which the data points lie
- By preserving the local and global relationships between data points, manifold learning provides a more meaningful representation for further analysis

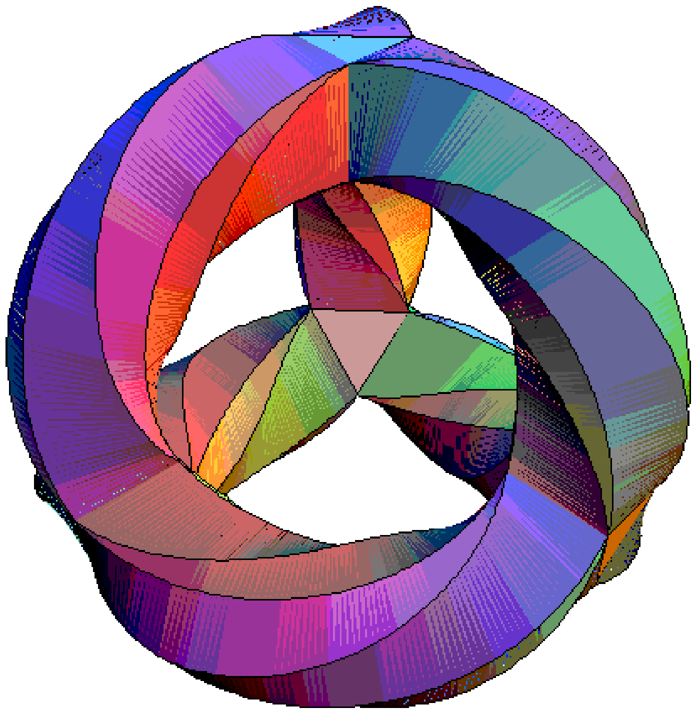

Manifold Learning & Data Fusion

Data Fusion: The process of combining and integrating multiple heterogeneous data sources or modalities to generate a unified representation, enabling enhanced insights and decision-making

↓

Manifold: A lower-dimensional, curved subspace embedded within a higher-dimensional space that captures the essential structure of the data

↓

Autoencoders: Consist one of the most fundamental Manifold Learning/Data Fusion techniques

Manifold Learning & Data Fusion

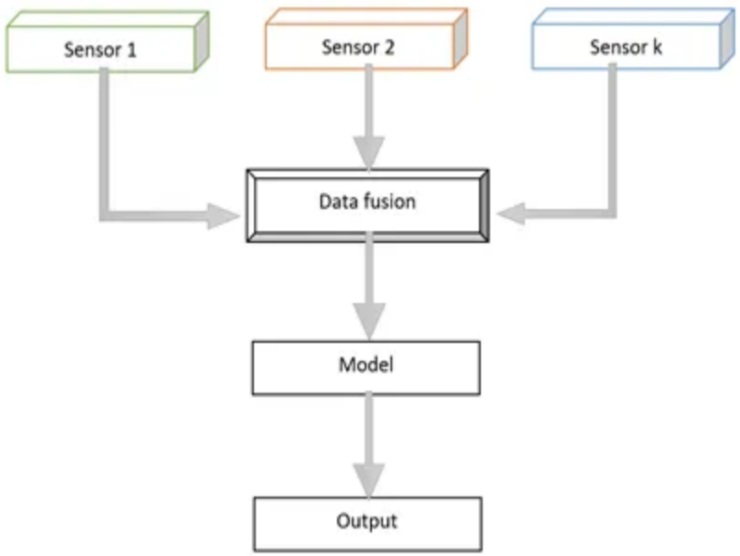

Early Fusion

- Fusing multiple data before the analysis
- Applicable on raw-data or pre-processed data obtained from sensors
- Features should be extracted from the data before fusion – synchronization problems
- Simplest Form: Concatenation

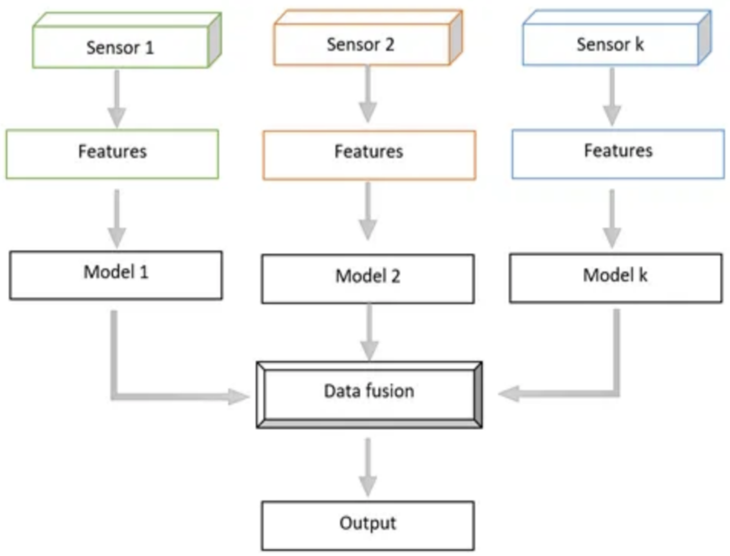

Late Fusion

- Uses data sources independently and fusion happens at a decision-making level
- Inspired by ensemble methods
- Simpler than early data fusion when data is varied (sampling rate, data dimensionality, measurement unit)
- Often better results – uncorrelated error

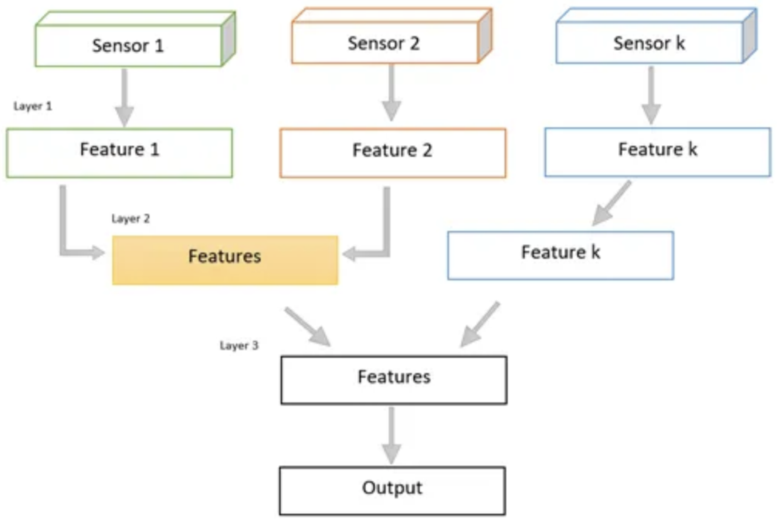

Intermediate Fusion

- Architecture built on the basis of the popular deep neural networks
- Fusion at different depths of the model
- Deep Learning fusion context: Learn a joint representation of each of the modalities
- PCA , Autoencoders

Autoencoders for Multimodal Data Fusion

Description

- Autoencoders (AE) are a type of Artificial Neural Networks (ANN) used for:
  - Unsupervised Learning
  - Dimensionality reduction/Compression
  - Data Fusion

How do they work? / Architecture

- **Encoder**: Compresses the input data into a lower-dimensional space using an encoder network – bottleneck layer
- **Decoder**: Then it reconstructs the input data back into the original space using a decoder network
- It learns an **internal representation/code** to perform useful transformations on the input data (middle layer)
- Finds a codification of the input multimodal data by learning non-linear combinations of their features

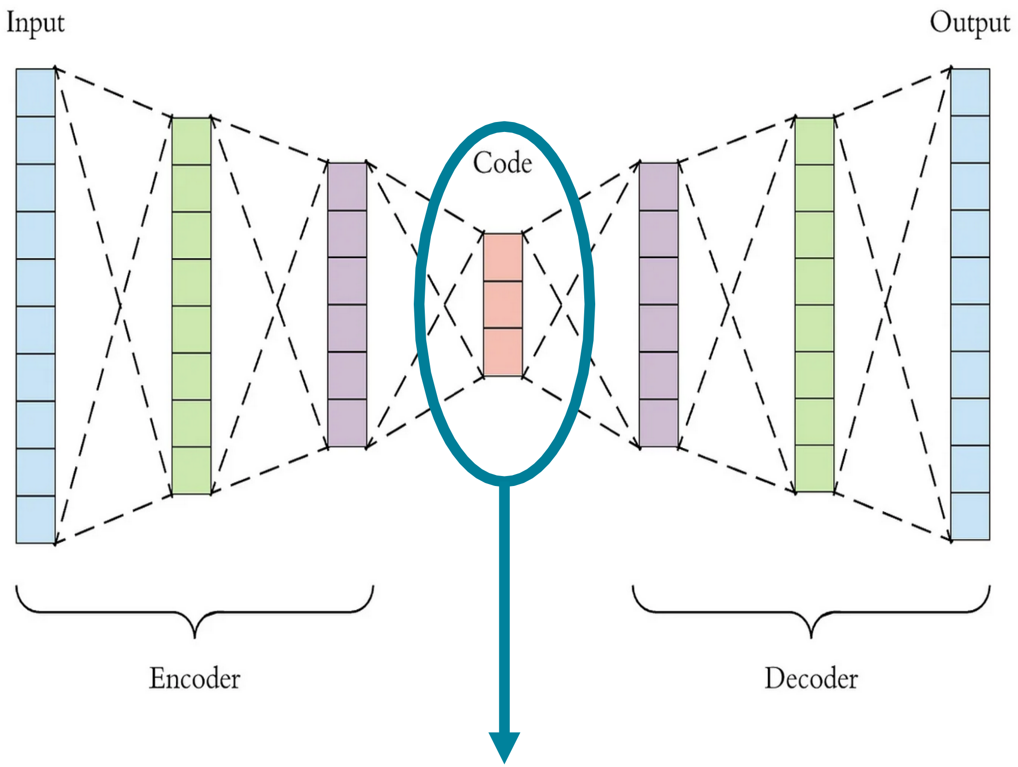

Undercomplete: New higher level variables are generated

Autoencoders for Multimodal Data Fusion

Training

- Involves minimizing the reconstruction error between the original input data and the decoded output
- Based on the Backpropagation algorithm with Mean Square Error (MSE)
- The model obtained will be able to map new examples onto the latent feature space

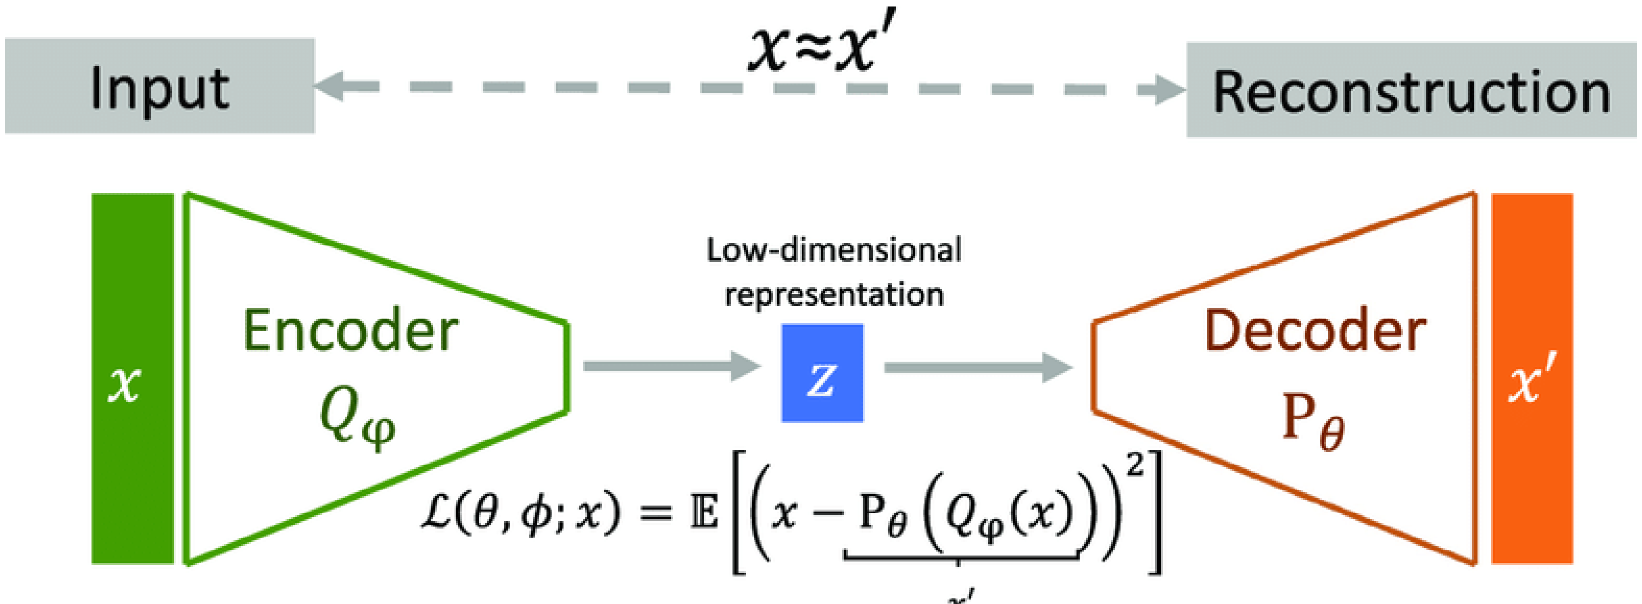

Autoencoders for Multimodal Data Fusion

Properties of AE

- **Data specific**: Only able to compress data similar to what they have been trained on
- **Lossy**: Degraded output representation of the input
- **Unsupervised**: They do not need explicit labels to train on. They can also be considered self-supervised because they generate their own labels from the training data

Types

- Fully connected
- Convolutional
- Recurrent (e.g., LSTM-based for temporal ordering of cyber-physical data)

Useful for Downstream Tasks

- Classification
- Clustering
- Anomaly detection

Note for AE:
- Output layer may be of:
  - No Importance: Extract code
  - Importance: Clean the noise of the input data

Autoencoders for Multimodal Data Fusion

Basic Autoencoder

- Feed forward ANN with symmetrical layer architecture
- The symmetry does not necessarily have to be reflected in the weights and activation functions
- Objective function corresponds to a per-instance loss function (MSE)
- Optimization of weights and biases: SGD, RMSProp, AdaGrad

Contractive Autoencoder

- Autoencoders are very sensitive to variations in the input data
- Small perturbations can generate very different encodings
- Contractive Autoencoders include a regularization term that allows them to stabilize the encodings

Autoencoders for Multimodal Data Fusion

Denoising Autoencoder

- Uses noise to build a new feature space that is more resistant to corrupt entries
- Fundamental modification: Corruption of entrance during the training phase
- Number of input variables randomly chosen – set to 0 – reconstruction error is compared to the original unmodified values
- Detects missing values

Robust Autoencoder

- Tolerate possible noise present in the training data
- Fundamental modification: Loss function used during training
- Error function based in correntropy
  - Measures the probability density that two events are similar
  - The outliers affect this measure to a lesser extent than the MSE

THANK YOU

Let's talk: gfragko@sandia.gov