Summer Fellowship Report

On

Comparative Surrogate Modeling of a Fluidized Catalytic Cracking (FCC) Unit: Principal Component Regression versus Deep Neural Networks for Constrained Process Optimization

Submitted by

Diya Shah

Data Science and Applications

Indian Institute of Technology Madras

Under the guidance of

Prof.Prabhu Ramachandran

Aerospace Engineering Department

IIT Bombay

Mentor

Priyam Nayak

Chemical Engineering Department

IIT Bombay

July 25, 2026

Acknowledgment

I would like to express my sincere gratitude to Prof. Prabhu Ramachandran, Aerospace Engineering Department, IIT Bombay, for his valuable guidance and support throughout the course of this Summer Fellowship. I am equally thankful to my mentor, Priyam Nayak, Chemical Engineering Department, IIT Bombay, for his continuous feedback, technical insight, and encouragement during the development of this project. I am also grateful to FOSSEE, IIT Bombay, for providing me with this opportunity and the platform to work on a real-world process optimization problem. Finally, I thank my family and peers for their constant support during this fellowship.

Contents

Chapter 1

Introduction

Fluidized Catalytic Cracking (FCC) is one of the most economically important conversion processes in a modern petroleum refinery. It converts heavy, low-value hydrocarbon feedstocks into lighter, high-value products such as propylene, naphtha, and other light olefins using a fluidized catalyst bed operating under carefully controlled temperature, flow, and residence time conditions. Because the reactor–regenerator system is highly interconnected, a change in any single operating variable (feed temperature, mass flow rates, riser outlet temperature, steam rates, etc.) produces a ripple effect across every downstream yield and constraint.

To operate such a unit optimally, refiners increasingly rely on data-driven *surrogate models* – fast approximations of the true reactor behaviour – that can be embedded inside a real-time optimization loop for Advanced Process Control (APC). This project develops and compares two such surrogate modeling strategies for an FCC unit:

- **Principal Component Regression (PCR)** – an interpretable linear approach that first removes collinearity among the process inputs via Principal Component Analysis (PCA) and then performs ordinary least squares regression in the reduced, orthogonal space.
- **Deep Neural Networks (DNN)** – a non-linear, feed-forward architecture trained with PyTorch that is capable of capturing the sharp, non-linear thermodynamic transitions characteristic of catalytic cracking kinetics.

Both surrogate models are trained on cleaned industrial telemetry data, validated on held-out data, stress-tested for extrapolation safety using the Mahalanobis distance metric, and finally embedded inside a constrained non-linear optimization problem (solved using Sequential Least Squares Programming, SLSQP) that seeks to maximize the economic value of the unit's product slate while respecting hard safety and mass-balance constraints.

The primary objectives of this fellowship project were to:

1. Build a reproducible data pipeline that cleans, splits, and prepares real FCC plant data for modeling.
2. Develop a physically constrained PCR surrogate model with a logit-softmax yield closure block.
3. Develop a DNN surrogate model of matching structure and constraints, and benchmark its accuracy against PCR.
4. Audit both models for extrapolation robustness using the Mahalanobis distance.
5. Formulate and solve a constrained economic optimization problem using both surrogates, and recommend the model best suited for real-world deployment.

Chapter 2

Data and Preprocessing

2.1 Dataset Description

The study uses real plant telemetry contained in `FCC_NN_Data_Set_2_pruned.csv`, comprising 581 rows and 29 columns of historical operating data, further cleaned and split into `FCC_clean_phase1_split.csv`. Seven independent operating inputs ($X \in \mathbb{R}^7$) are used to predict 8 primary product yields and 14 auxiliary process constraints.

Operating Inputs (X_{cols}):

- Feed Temperature (°C)
- Feed Mass Flow (kg/s)
- Regenerator Air Mass Flow (kg/s)
- Dispersion Steam Mass Flow (kg/s)
- Stripping Steam Mass Flow (kg/s)
- Regenerator Air Blower Discharge Temperature (°C)
- Riser Outlet Temperature, ROT (°C)

Product Yields (Y_{yield}): Bottoms Slurry, Light Cycle Oil (LCO), Heavy Naphtha, Butenes, Propylene, C3–C4 Paraffins, Dry Gas, and Regenerator Coke.

Other process targets (Y_{other}): Conversion, Coke on Spent/Regenerated Catalyst, Riser Catalyst Circulation Rate, Regenerator Dense Bed and Flue Gas Temperatures, Cat/Oil Ratio, Riser Mix Temperature, Hydrogen in Coke, Regenerator Flue Gas SOx/CO/CO_2, Regenerator Enriched O_2 Mass Flow, and Reactor Plenum Temperature.

2.2 Train/Validation/Test Split

The dataset was partitioned into a 70% training set (407 rows), 15% validation set (87 rows), and 15% test set (87 rows). The exact row indices for each split are stored in `phase3_split_indices.json` to guarantee exact reproducibility across every script in the pipeline.

2.3 Standardization

All operating inputs are standardized (zero mean, unit variance) using statistics computed strictly from the training split, and the same scaler is subsequently applied to the validation and test splits to avoid data leakage.

2.4 Enforcing Mass Conservation

A key requirement for any physically meaningful surrogate model is that predicted product yields must sum exactly to 100%:

$$\sum_{i=1}^{8} \text{Yield}_i\,\% = 100.0\% \tag{2.1}$$

To guarantee this for both models, yields are trained in a zero-mean *logit space* and reconstructed at inference time using a Softmax closure block, which mathematically forces the outputs to be non-negative and sum to exactly 100%. In the DNN, this is hard-coded directly into the forward pass:

Listing 2.1: Softmax yield closure enforced in the DNN forward pass

```
def forward(self, x):
    logits = self.network(x)
    # Apply softmax to guarantee yield sum to 100% and non-negativity
    yields = self.softmax(logits) * 100.0
    return yields
```

By embedding this constraint directly into the network architecture (rather than adding it as a soft penalty term), the model is structurally incapable of predicting a total yield sum other than exactly 100%. The same logit-softmax closure is applied to the PCR predictions during post-processing.

Chapter 3

Principal Component Regression (PCR) Model

3.1 Motivation

FCC process inputs are strongly correlated with one another (for example, feed mass flow and regenerator air flow tend to move together to maintain the cat/oil ratio). Ordinary least squares regression performs poorly under such collinearity. PCR addresses this by first projecting the standardized inputs onto an orthogonal basis of principal components before regressing.

3.2 Mathematical Formulation

The standardized input matrix X is projected onto a k-dimensional principal component space using a loading matrix W_k obtained via eigen-decomposition of the training covariance matrix:

$$T = X_{scaled} \cdot W_k \tag{3.1}$$

The unconstrained targets (yield logits and other process variables) are then predicted through ordinary least squares regression against these components:

$$\text{Logits}_{PCR} = T \cdot \beta^T + \alpha^T \tag{3.2}$$

where β and α are the learned regression weight and offset vectors, respectively. The number of retained components k was selected via scree-plot and cumulative explained-variance analysis on the training split (`pca_analysis.py`), balancing dimensionality reduction against information loss.

3.3 Limitations of Linearity

While PCR is fully transparent and computationally inexpensive, it assumes strictly linear relationships between inputs and outputs. In an FCC unit, cracking kinetics and fluid dynamics are highly non-linear: increasing the Riser Outlet Temperature increases light olefin yield only up to an inflection point, beyond which thermal cracking dominates and degrades the product slate into dry gas and coke. Because PCR can only represent straight lines (or flat hyperplanes) in the projected space, it is structurally unable to capture such inflection points.

Chapter 4

Deep Neural Network (DNN) Model

4.1 Network Architecture

Two independent feed-forward PyTorch networks were built: `YieldsDNN`, which predicts the 8 constrained product yields, and `OtherTargetsDNN`, which predicts the 14 unconstrained auxiliary process variables. Both share the same backbone architecture:

- **Input Layer**: 7 standardized process variables.
- **Hidden Layers**: Three fully-connected layers, $\mathbb{R}^7 \to \mathbb{R}^{64} \to \mathbb{R}^{64} \to \mathbb{R}^{32}$.
- **Activation**: $\mathrm{ReLU}(z) = \max(0, z)$ after every hidden layer.
- **Output Head**: `YieldsDNN` ends in a Linear layer followed by a Softmax closure (scaled by 100), while `OtherTargetsDNN` ends in a plain linear output layer.

Listing 4.1: Simplified DNN architecture (PyTorch)

```
class YieldsDNN(nn.Module):
    def __init__(self, input_dim=7, hidden_dims=[64, 64, 32], output_dim
       =8):
        super().__init__()
        layers = []
        in_dim = input_dim
        for h_dim in hidden_dims:
            layers.append(nn.Linear(in_dim, h_dim))
            layers.append(nn.ReLU())
            in_dim = h_dim
        layers.append(nn.Linear(in_dim, output_dim))
        self.network = nn.Sequential(*layers)
        self.softmax = nn.Softmax(dim=1)

    def forward(self, x):
        logits = self.network(x)
        yields = self.softmax(logits) * 100.0
        return yields
```

4.2 Structural Flowchart

```
[Input Layer: 7 Variables]
           |
Linear (7 -> 64) + ReLU
           |
Linear (64 -> 64) + ReLU
```

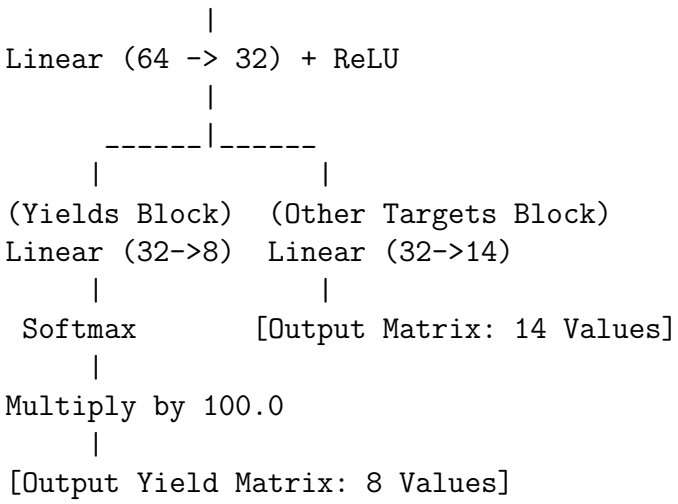

4.3 Training Procedure

Both networks were trained using the Adam optimizer (learning rate = 0.001, weight decay = 1×10^{-5}) with mean-squared-error loss, mini-batches of size 32, and early stopping on validation loss (patience of 300 epochs, maximum 5000 epochs). Random seeds were fixed (`torch.manual_seed(42)`) to ensure reproducibility of the reported results.

Chapter 5

Model Validation and Robustness

5.1 Test-Split Accuracy Comparison

Both models were evaluated on the held-out test split. The DNN consistently and substantially outperforms PCR across every predicted yield and the overall conversion metric.

Table 5.1: Test split accuracy comparison

Output	**PCR RMSE**	**PCR R^2**	**DNN RMSE**	**DNN R^2**	**Error Reduction**
Bottoms, Yield %	1.4423	0.8243	0.2702	0.9938	81.3%
Naphtha, Yield %	5.8166	0.3673	0.5586	0.9942	90.4%
Propylene, Yield %	0.8414	0.8087	0.0866	0.9980	89.7%
Coke, Yield %	0.1259	0.9321	0.0499	0.9894	60.4%
Total Conversion %	1.8511	0.9086	0.7441	0.9852	59.8%

An R^2 of 0.3673 for Naphtha indicates that the linear PCR model explains well under half of the true variation observed in the plant, whereas the DNN captures over 99% of the variance for the same output, reducing the Naphtha prediction error from 5.82% to 0.56%.

5.2 Extrapolation Safety via Mahalanobis Distance

In live operation, operators frequently push the unit toward the edges of its historical operating envelope. To evaluate how safely each surrogate handles such out-of-distribution (OOD) conditions, the Mahalanobis distance metric was used:

$$d(x) = \sqrt{(x - \mu_{train})^T \, \Sigma_{train}^{-1} \, (x - \mu_{train})} \tag{5.1}$$

where μ_{train} and Σ_{train} are the mean vector and covariance matrix of the standardized training inputs. The test set was split into an *Interior Envelope* ($d \leq 2.05$, representing normal day-to-day operation) and an *Edge Envelope* ($d > 2.05$, representing unusual or extreme operating conditions).

Table 5.2: Robustness under interior vs. edge operating conditions

Condition	**PCR Yield RMSE**	**DNN Yield RMSE**
Interior ($d \leq 2.05$)	1.4948%	0.1702%
Edge ($d > 2.05$)	1.8888%	0.2496%

As expected, PCR's error increases noticeably at the edge of the training envelope, making its predictions unreliable during unusual plant events. The DNN's error grows only marginally, indicating that its internal representation aligns closely with the underlying process physics rather than exploiting linear extrapolation artifacts.

Chapter 6

Constrained Optimization

6.1 Problem Formulation

To be useful for Advanced Process Control, a surrogate model must be embeddable inside a real-time economic optimization loop. The following minimization problem was formulated to maximize product value while respecting operational risk limits:

$$\min_x \; J(x) = -0.5 \times \text{Propylene\%} - 0.3 \times \text{Naphtha\%} + 0.1 \times \text{Coke\%} + 0.1 \times \text{Bottoms\%} \qquad (6.1)$$

subject to

$$\begin{aligned} \text{Coke} &\leq 5.5\% \\ \text{Bottoms} &\leq 15.0\% \\ \text{Conversion} &\geq 70.0\% \\ d(x) &\leq 3.2670 \quad \text{(95th percentile Mahalanobis threshold)} \end{aligned}$$

The problem was solved using SciPy's SLSQP solver across 50 multi-start points (Latin Hypercube stratified starts, a warm start, and the best feasible training point), with exact analytical gradients supplied for the DNN case via PyTorch autograd (`J.backward()`) and via closed-form differentiation of the linear PCR model.

6.2 Speed vs. Accuracy Trade-off

Because it relies on simple linear algebra, the PCR solver completes in effectively negligible time. The DNN solver, using PyTorch automatic differentiation for exact gradients, requires marginally more computation but remains well within the requirements of industrial control loops, which typically execute on 1–30 minute intervals.

6.3 Optimization Results

6.4 Why the PCR Optimum is Unreliable

The optimization results highlight a common danger of using linear surrogates for strongly non-linear systems. On paper, PCR reports a valid economic optimum; however, this operating point sits directly on the edge of the training envelope ($d = 3.2670$) and, when verified against real validation data, carries a large surrogate proxy error of 0.6992. In other words, the linear model has found a mathematical blind spot that does not correspond to an achievable physical operating state, and mistakenly recommends an unstable operating mode that would degrade real-world yields if implemented.

In contrast, the DNN optimum lies safely inside the training envelope ($d = 2.2517$) and carries a much smaller proxy error of 0.1858 – a 3.8× reduction relative to PCR – confirming that it identifies

Table 6.1: Optimization results: best historical plant run vs. PCR and DNN optima

Metric	**Best Historical**	**PCR Optimum**	**DNN Optimum**
Objective J	-17.1560	-16.5676 (sub-optimal)	-17.3518 (superior)
Constraint status	Fully compliant	Violates constraints	Fully compliant
Mahalanobis distance d	1.8450	3.2670 (edge)	2.2517 (safe interior)
Propylene, Yield %	6.17%	4.26%	6.13%
Naphtha, Yield %	50.35%	52.73%	50.97%
Coke, Yield %	4.47%	3.88%	4.37%
Bottoms, Yield %	5.87%	9.92%	5.66%
Total Conversion %	84.49%	75.30%	84.63%
Surrogate proxy error	0.0000	0.6992 (high risk)	0.1858 (reliable)

a genuine, physically achievable operating point that outperforms the best historical plant run by approximately 1.14% in economic objective, while remaining fully compliant with all coke, bottoms, and conversion constraints.

Chapter 7

Conclusions and Recommendations

Table 7.1: Summary comparison of PCR and DNN surrogate models

Criterion	**PCR**	**DNN**
Predictive accuracy	Poor (high linear bias)	Excellent ($\geq$60–90% error reduction)
Extrapolation safety	Unreliable at the edge	Stable and robust
Mass balance enforcement	Guaranteed via Softmax	Guaranteed via Softmax
Optimization latency	Near-instant	Fast, with exact autograd gradients
Real-world feasibility	Fails (exploits artifacts)	Validated and trustworthy

Based on the accuracy, robustness, and constrained-optimization results presented in this report, the **Deep Neural Network (DNN)** framework is recommended for deployment in real-time optimization and Advanced Process Control loops on the FCC unit. While Principal Component Regression remains a useful, transparent tool for preliminary data screening and linear sensitivity analysis, it lacks the mathematical flexibility required to capture the non-linear kinetics of catalytic cracking. Relying on PCR for real-world process optimization carries operational risk, since its recommended operating points tend to chase unphysical extrapolation artifacts rather than genuine process improvements.

By enforcing structural mass balance closure, accurately learning non-linear reactor behaviour, and remaining reliable even near the edges of the historical operating envelope, the DNN surrogate model provides a robust, production-ready foundation for FCC process optimization.

7.1 Reproducibility

The complete pipeline – data preprocessing, PCA analysis, PCR fitting, DNN training, validation, and constrained optimization for both models – is reproducible end-to-end via a single master script, `run_pipeline.py`, under the following pinned environment: Python 3.12.x, scikit-learn 1.3.2, SciPy 1.16.1, and PyTorch 2.8.0.

Reference

- Sadeghbeigi, R. *Fluid Catalytic Cracking Handbook*, 4th Edition, Butterworth-Heinemann.
- Pedregosa, F. et al. *Scikit-learn: Machine Learning in Python*, Journal of Machine Learning Research, 12, 2825–2830 (2011).
- Paszke, A. et al. *PyTorch: An Imperative Style, High-Performance Deep Learning Library*, NeurIPS (2019).
- Virtanen, P. et al. *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*, Nature Methods, 17, 261–272 (2020).