

Fusion Forecasting for Residential PV: Reference Implementation Documentation

Erik Schmidt
`erik@f97.io`

August 28, 2025

Abstract

This document provides technical documentation for the reference implementation of Fusion Forecasting, an adaptive regression framework that blends AI-based and physics-based (GTI) solar forecasts with online calibration. The implementation follows the methodology described in the original paper, providing a complete, operational system for residential PV forecasting, including a command-line interface and detailed data specifications.

Contents

1	Introduction	2
2	Mathematical Formulation	2
2.1	Linear Blending Layer	2
2.2	Regression Objective	2
2.3	Online Calibration	2
2.4	Uncertainty Quantification	3
3	Implementation Architecture	3
3.1	Class Structure	3
3.2	Core Methods	3
3.2.1	Model Training	3
3.2.2	Online Calibration	3
3.2.3	Forecast Generation	3
3.2.4	Uncertainty Estimation	3
4	Command-Line Interface	3
4.1	Usage Examples	5
4.2	Command-Line Options	5
5	Data Requirements and CSV Format	5
5.1	CSV File Specification	5
5.2	Example CSV File	5
5.3	Output Format	6
6	Implementation Details	6
6.1	Dependencies	6
6.2	Key Parameters	6
7	Usage Example	6
7.1	Basic Usage	6
7.2	Daily Operation	7
8	Performance Considerations	7
8.1	Computational Complexity	7
8.2	Memory Requirements	8

9 Extensions and Customization	8
9.1 Custom Loss Functions	8
9.2 Additional Features	8
9.3 Advanced Calibration	8
10 Conclusion	8

1 Introduction

Fusion Forecasting addresses the challenge of photovoltaic (PV) output forecasting by combining the strengths of AI-based and physics-based approaches while mitigating their individual weaknesses. This reference implementation provides a complete, operational system based on the methodology described in the original paper.

The system consists of four core components:

1. Linear blending of AI and GTI forecasts
2. Robust regression with time-series cross-validation
3. Online calibration using median-ratio scaling
4. Uncertainty quantification with robust statistics

This enhanced implementation adds a comprehensive command-line interface and detailed specifications for data input formats.

2 Mathematical Formulation

2.1 Linear Blending Layer

The fusion model predicts daily energy production through an affine combination:

$$\hat{y}_{\text{lin}}(t) = w_{\text{AI}} \cdot \hat{y}_{\text{AI}}(t) + w_{\text{GTI}} \cdot \hat{y}_{\text{GTI}}(t) + b$$

where:

- $\hat{y}_{\text{AI}}(t)$ is the AI-based forecast
- $\hat{y}_{\text{GTI}}(t)$ is the GTI-based forecast
- $w_{\text{AI}}, w_{\text{GTI}}$ are learned weights
- b is the intercept term

2.2 Regression Objective

Parameters are learned by minimizing a robust, regularized objective function:

$$\min_{\theta} \frac{1}{N} \sum_{t \in \mathcal{T}_{\text{train}}} \rho_{\delta}(y(t) - \hat{y}_{\text{lin}}(t)) + \lambda (w_{\text{AI}}^2 + w_{\text{GTI}}^2)$$

where $\rho_{\delta}(\cdot)$ is the Huber loss function and λ is the ridge regularization parameter.

2.3 Online Calibration

The calibration scale factor is computed as:

$$\text{scale} = \text{median} \left(\text{winsorize}_{\alpha} \left[\frac{y(d)}{\hat{y}_{\text{lin}}(d)} \right]_{d \in \mathcal{D}_K} \right)$$

where \mathcal{D}_K contains the most recent K days with valid data, and winsorization reduces the influence of outliers.

2.4 Uncertainty Quantification

Prediction uncertainty is estimated using Median Absolute Deviation (MAD):

$$\hat{\sigma} = 1.4826 \cdot \text{MAD}(\{\varepsilon(d)\}_{d \in \mathcal{D}_K})$$

where $\varepsilon(d) = y(d) - \hat{y}_{\text{fusion}}(d)$.

3 Implementation Architecture

3.1 Class Structure

The implementation is organized around the `FusionForecaster` class, which encapsulates the complete forecasting pipeline.

```
1 class FusionForecaster:
2     def __init__(self,
3                  train_window_size=90,
4                  calibration_window_size=14,
5                  alpha=0.05,
6                  p_max=10.0,
7                  min_train_samples=30,
8                  fallback_weights=(0.7, 0.3)):
9         # Initialization code
```

Listing 1: `FusionForecaster` class initialization

3.2 Core Methods

3.2.1 Model Training

The `train_fusion` method implements the training procedure:

Algorithm 1 Model Training Procedure

```
1: procedure TRAIN_FUSION( $\mathcal{D}_{\text{train}}$ )
2:   if  $|\mathcal{D}_{\text{train}}| < N_{\min}$  then
3:     return Fallback weights (0.7, 0.3)
4:   end if
5:   Prepare feature matrix  $X$  and target vector  $y$ 
6:   Perform time-series cross-validation to select  $\lambda$ 
7:   Fit Huber regression model with best  $\lambda$ 
8:   return learned parameters  $w_{\text{AI}}, w_{\text{GTI}}, b$ 
9: end procedure
```

3.2.2 Online Calibration

The `calibrate_scale` method implements online calibration:

3.2.3 Forecast Generation

The `forecast_day` method generates predictions:

3.2.4 Uncertainty Estimation

The `estimate_uncertainty` method quantifies prediction uncertainty:

4 Command-Line Interface

The enhanced implementation includes a comprehensive command-line interface for operational use.

Algorithm 2 Online Calibration Procedure

```
1: procedure CALIBRATE_SCALE( $\mathcal{D}_{\text{recent}}$ )
2:   Initialize empty ratio list  $R$ 
3:   for each day  $d \in \mathcal{D}_{\text{recent}}$  do
4:     if  $y(d) > 0$  and  $\hat{y}_{\text{lin}}(d) > 0$  then
5:        $r(d) \leftarrow y(d)/\hat{y}_{\text{lin}}(d)$ 
6:       Append  $r(d)$  to  $R$ 
7:     end if
8:   end for
9:   if  $|R| < \max(3, 0.3 \cdot K)$  then
10:    return previous scale
11:   end if
12:    $R_{\text{winsorized}} \leftarrow \text{winsorize}(R, \alpha)$ 
13:   return median( $R_{\text{winsorized}}$ )
14: end procedure
```

Algorithm 3 Forecast Generation Procedure

```
1: procedure FORECAST_DAY( $\hat{y}_{\text{AI}}, \hat{y}_{\text{GTI}}$ , daylight_hours)
2:    $\hat{y}_{\text{lin}} \leftarrow w_{\text{AI}} \cdot \hat{y}_{\text{AI}} + w_{\text{GTI}} \cdot \hat{y}_{\text{GTI}} + b$ 
3:    $\hat{y}_{\text{fusion}} \leftarrow \text{scale} \cdot \hat{y}_{\text{lin}}$ 
4:    $\hat{y}_{\text{fusion}} \leftarrow \max(0, \min(\hat{y}_{\text{fusion}}, P_{\text{max}} \cdot \text{daylight\_hours}))$ 
5:   return  $\hat{y}_{\text{fusion}}$ 
6: end procedure
```

Algorithm 4 Uncertainty Estimation Procedure

```
1: procedure UNCERTAINTY( $\mathcal{D}_{\text{recent}}$ )
2:   Initialize empty error list  $E$ 
3:   for each day  $d \in \mathcal{D}_{\text{recent}}$  do
4:     if valid data exists for day  $d$  then
5:        $\varepsilon(d) \leftarrow y(d) - \hat{y}_{\text{fusion}}(d)$ 
6:       Append  $\varepsilon(d)$  to  $E$ 
7:     end if
8:   end for
9:   if  $|E| < 5$  then
10:    return None
11:   end if
12:    $\text{MAD} \leftarrow \text{median}(|E - \text{median}(E)|)$ 
13:    $\hat{\sigma} \leftarrow 1.4826 \cdot \text{MAD}$ 
14:   return  $\hat{\sigma}$ 
15: end procedure
```

4.1 Usage Examples

```

1 # Basic usage with a CSV file
2 python fusion_forecast.py --train-data historical_data.csv --forecast-dates
   2023-08-01,2023-08-02
3
4 # With custom parameters
5 python fusion_forecast.py --train-data data.csv --train-window 60 --calibration-window
   14 --p-max 8.5 --forecast-dates 2023-08-01
6
7 # Save results to a JSON file
8 python fusion_forecast.py --train-data data.csv --forecast-dates 2023-08-01 --output
   results.json
9
10 # Verbose output
11 python fusion_forecast.py --train-data data.csv --forecast-dates 2023-08-01 --verbose

```

Listing 2: Command-line usage examples

4.2 Command-Line Options

Option	Default	Description
-train-data	(required)	CSV file with historical training data
-forecast-dates	(required)	Dates to forecast (YYYY-MM-DD, comma-separated)
-output	None	Output file for results (JSON format)
-train-window	90	Days used for training
-calibration-window	14	Days used for calibration
-p-max	10.0	System capacity (kW)
-min-train-samples	30	Minimum samples for training
-alpha	0.05	Winsorization parameter
-fallback-ai-weight	0.7	Fallback weight for AI forecast
-fallback-gti-weight	0.3	Fallback weight for GTI forecast
-retrain-frequency	30	Retraining frequency (days)
-confidence	0.95	Confidence level for prediction intervals
-verbose	False	Enable verbose output

Table 1: Command-line options

5 Data Requirements and CSV Format

5.1 CSV File Specification

The system requires a CSV file with historical data for training. The file must contain the following columns:

Column	Type	Description
date	Date (YYYY-MM-DD)	Date of observation
y_ai	Numeric (kWh)	AI-based forecast
y_gti	Numeric (kWh)	GTI-based forecast
y_actual	Numeric (kWh)	Actual energy production
daylight_hours	Numeric (hours)	Number of daylight hours

Table 2: CSV file column specifications

5.2 Example CSV File

```

1 date,y_ai,y_gti,y_actual,daylight_hours
2 2023-01-01,8.5,7.2,8.1,9.5

```

```

3 2023-01-02,9.2,8.1,9.5,9.6
4 2023-01-03,7.8,6.9,7.2,9.6
5 2023-01-04,6.5,5.8,6.8,9.7
6 2023-01-05,10.2,9.5,10.5,9.7

```

Listing 3: Example CSV data format

5.3 Output Format

The program outputs results in JSON format with the following structure:

```

1 [
2   {
3     "date": "2023-08-01",
4     "forecast": 12.45,
5     "uncertainty": 1.23,
6     "w_ai": 0.65,
7     "w_gti": 0.42,
8     "b": -0.32,
9     "scale": 1.05,
10    "prediction_interval": {
11      "lower": 10.12,
12      "upper": 14.78,
13      "confidence": 0.95
14    }
15  }
16 ]

```

Listing 4: Output JSON format

6 Implementation Details

6.1 Dependencies

The implementation requires the following Python packages:

- numpy $\geq 1.20.0$
- pandas $\geq 1.3.0$
- scikit-learn $\geq 1.0.0$
- scipy $\geq 1.7.0$

6.2 Key Parameters

Parameter	Default	Description
train_window_size	90	Days used for training
calibration_window_size	14	Days used for calibration
alpha	0.05	Winsorization parameter
p_max	10.0	System capacity (kW)
min_train_samples	30	Minimum samples for training
fallback_weights	(0.7, 0.3)	Weights when falling back

Table 3: Key configuration parameters

7 Usage Example

7.1 Basic Usage

```

1 # Initialize the forecaster
2 forecaster = FusionForecaster(
3     train_window_size=60,
4     calibration_window_size=14,
5     p_max=10.0 # System capacity: 10 kW
6 )
7
8 # Add historical observations
9 for i in range(len(historical_data)):
10     forecaster.add_observation(
11         date=dates[i],
12         y_ai=ai_forecasts[i],
13         y_gti=gti_forecasts[i],
14         y_actual=actual_production[i],
15         daylight_hours=daylight_hours[i]
16     )
17
18 # Generate a forecast
19 forecast = forecaster.forecast_day(
20     y_ai=next_ai_forecast,
21     y_gti=next_gti_forecast,
22     daylight_hours=next_daylight_hours
23 )
24
25 # Get prediction interval
26 point_forecast, lower, upper = forecaster.prediction_interval(
27     y_ai=next_ai_forecast,
28     y_gti=next_gti_forecast,
29     daylight_hours=next_daylight_hours,
30     confidence=0.95
31 )

```

Listing 5: Basic usage example

7.2 Daily Operation

For operational use, the `daily_operation` method provides a complete workflow:

```

1 # Complete daily operation
2 result = forecaster.daily_operation(
3     date=current_date,
4     y_ai=ai_forecast,
5     y_gti=gti_forecast,
6     daylight_hours=daylight_hours,
7     y_actual=actual_production, # If available
8     retrain_frequency=30 # Retrain every 30 days
9 )
10
11 # Result contains forecast, uncertainty, and model parameters
12 print(f"Forecast: {result['forecast']:.2f} kWh")
13 print(f"Uncertainty: {result['uncertainty']:.2f} kWh")
14 print(f"95% Prediction Interval: [{result['lower_95']:.2f}, {result['upper_95']:.2f}]")

```

Listing 6: Daily operational workflow

8 Performance Considerations

8.1 Computational Complexity

The implementation is designed for efficiency on commodity hardware:

- Training: $O(N \cdot M)$ where N is training window size and M is number of λ values tested
- Calibration: $O(K)$ where K is calibration window size
- Forecasting: $O(1)$ per forecast

8.2 Memory Requirements

Memory usage is primarily determined by:

- History storage: $O(T)$ where T is the maximum of training and calibration windows
- Model parameters: Constant size (3 weights + scale factor)

9 Extensions and Customization

The implementation can be extended in several ways:

9.1 Custom Loss Functions

The Huber loss can be replaced with other robust loss functions:

```
1 from sklearn.linear_model import RANSACRegressor
2
3 # Use RANSAC for more robust regression
4 model = RANSACRegressor(
5     base_estimator=Ridge(alpha=best_lambda),
6     min_samples=0.8
7 )
```

Listing 7: Custom loss function example

9.2 Additional Features

The model can be extended to include additional features:

```
1 # Extended feature vector at time t
2 # The variables are assumed to be defined elsewhere in the code.
3 x_t = [
4     y_hat_ai(t),      # AI model's prediction
5     y_hat_gti(t),    # GTI model's prediction
6     temperature(t),  # Temperature reading
7     humidity(t)      # Humidity reading
8 ]
9 # x_t is now a list that can be used as a feature vector.
```

Listing 8: Creating an extended feature vector

9.3 Advanced Calibration

The calibration procedure can be enhanced with adaptive window sizing or different robust estimators.

10 Conclusion

This reference implementation provides a complete, operational version of the Fusion Forecasting system described in the original paper. The system combines AI-based and physics-based forecasts using adaptive regression blending with online calibration, resulting in improved accuracy and robustness for residential PV forecasting.

The enhanced implementation includes a comprehensive command-line interface and detailed data specifications, making it suitable for operational deployment. The implementation is designed for practical use, with careful attention to operational safeguards, missing data handling, and computational efficiency. The modular architecture allows for customization and extension to meet specific requirements of different PV forecasting applications.