

# Multi-Horizon Time Series Forecasting of Non-Parametric CDFs with Deep Lattice Networks

Niklas Erdmann<sup>1</sup>, Lars Bentsen<sup>2</sup>, Roy Stenbro<sup>3</sup>, Heine Nygard Riise<sup>3</sup>, Narada Dilp Warakagoda<sup>1,4</sup>, Paal E. Engelstad<sup>1</sup>

<sup>1</sup>Department of Technology Systems, University of Oslo, 2007 Kjeller, Norway

<sup>2</sup>BearingPoint, 0252 Oslo, Norway

<sup>3</sup>Institute for Energy Technology, 2007 Kjeller, Norway

<sup>4</sup>Norwegian Defence Research Establishment, 2007 Kjeller, Norway

{niklase, n.d.warakagoda, paal.engelstad}@its.uio.no, lars.bentsen@bearingpoint.com, {roy.stenbro, heine.riise}@ife.no

## Abstract

Probabilistic forecasting is not only a way to add more information to a prediction of the future, but it also builds on weaknesses in point prediction. Sudden changes in a time series can still be captured by a cumulative distribution function (CDF), while a point prediction is likely to miss it entirely. The modeling of CDFs within forecasts has historically been limited to parametric approaches, but due to recent advances, this no longer has to be the case. We aim to advance the fields of probabilistic forecasting and monotonic networks by connecting them and propose an approach that permits the forecasting of implicit, complete, and non-parametric CDFs. For this purpose, we propose an adaptation to deep lattice networks (DLN) for monotonically constrained simultaneous/implicit quantile regression in time series forecasting. Quantile regression usually produces quantile crossovers, which need to be prevented to achieve a legitimate CDF. By leveraging long short term memory units (LSTM) as the embedding layer, and spreading quantile inputs to all sub-lattices of a DLN with an extended output size, we can produce a multi-horizon forecast of an implicit CDF due to the monotonic constraintability of DLNs that prevent quantile crossovers. We compare and evaluate our approach's performance to relevant state of the art within the context of a highly relevant application of time series forecasting: Day-ahead, hourly forecasts of solar irradiance observations. Our experiments show that the adaptation of a DLN performs just as well or even better than an unconstrained approach. Further comparison of the adapted DLN against a scalable monotonic neural network shows that our approach performs better. With this adaptation of DLNs, we intend to create more interest and crossover investigations in techniques of monotonic neural networks and probabilistic forecasting.

## Code and Dataset —

<https://github.com/Coopez/CDF-Forecasts-with-DLNs>

## Extended version —

<https://doi.org/10.48550/arXiv.2511.13756>

## 1 Introduction

Probabilistic forecasting is a quickly growing field that aims to provide more information to a forecast than just a point

estimate of the future (e.g. (Abdar et al. 2021)). Forecasting finds application in many fields, including meteorology, economics, and energy production (Gneiting, Balabdaoui, and Raftery 2007). In probabilistic forecasting, the future is predicted as a probability distribution. This can be used to give us more details about the future, such as the uncertainty or risk associated with a prediction. Forecasts about the future will always have a degree of uncertainty, which should be modeled (Price et al. 2025). In the energy market, for example, renewable energy prices are influenced by inherently variable weather conditions. Probabilistic forecasts can capture such variabilities and provide insights that can be used for informed risk management and bidding decisions by human or machine operators (Morales et al. 2014). These are processes in which optimization leads directly to a higher return in value and efficient use of resources (such as energy).

Generally, probabilistic approaches can be categorized into parametric and non-parametric approaches. Parametric approaches assume a specific distribution of the data, in which prior knowledge about the data can be leveraged to reach better performance. Non-parametric approaches do not assume a specific distribution, but are also not bound to adhere to assumptions about an underlying distribution that may not hold true.

Quantile regression (QR) (Koenker and Bassett 1978) is a non-parametric approach that is still actively developed and applied today. Due to its non-parametric nature, QR has a natural advantage in flexibility and adaptability compared to parametric approaches. It is also very much extendable for usage in time series forecasting (Wen et al. 2017). As a result of this, QR is widely used in areas like weather and energy forecasting (Song et al. 2024; Lauret, David, and Pedro 2017). Recently, several novel advances have shown similar ways to use neural networks in combination with QR to implicitly model all quantiles between 0 and 1 continuously in a neural network's latent space (Tagasovska and Lopez-Paz 2019; Dabney et al. 2018; Gasthaus et al. 2019). Related work coined this technique as simultaneous quantile regression (SQR), or Implicit quantile networks (IQN). Here we will use SQR for brevity. SQR enables us to implicitly estimate a non-parametric cumulative distribution function (CDF). By extension, it transforms a quantile regressor function into a complete inverse CDF of all quantile values.

However, QR, and especially SQR (as shown previously in Narayan et al. (2023)), is vulnerable to so-called quantile crossover, in which output values of lower quantiles cross over higher quantile values. Such a crossover decreases interpretability and applicability to real-world solutions. A modern car navigation system may provide three time-of-arrival estimates, such as: An optimistic 5 min, a likely 10 min, and a worst-case of 8 min. This does not make sense, as the worst-case time of arrival should not be shorter than the likely case. For our purposes, it also invalidates SQR, providing an implicit CDF, as it violates the monotonicity constraint of quantiles (Tagasovska and Lopez-Paz 2019).

Quantile monotonicity must be guaranteed to ensure an accurate and legitimate approximation of a CDF. A proposed solution for quantile crossover is to develop a function approximator that can be monotonically constrained to the SQR quantile input. Gasthaus et al. (2019) proposed to use splines in a quantile function with a recurrent neural network output (SQF-RNN). Their model can transform an input into parameters for a monotonic piecewise linear function of quantiles. Independently, Brando et al. (2022) implemented an approach whereby a positive neural network is considered the partial derivative of the conditional quantile function, resulting in a non-crossing inverse CDF. Narayan et al. (2023) proposed to use a Deep Lattice Network (DLN) for this task. Lattices are monotonically constrainable lookup tables (Gupta et al. 2016; You et al. 2017). Narayan et al. also compared their work to SQF-RNNs and found lattices to show higher accuracy and faster run times. SQF-RNNs have been applied to time series forecasting tasks before (Wang et al. 2022). However, DLN parametric complexity scales with the number of input features, which led to DLNs only being applied on problems with a limited number of input features (Narayan et al. 2023). In particular, DLNs have never been used in time series forecasting.

There is ongoing development in preventing quantile crossover in QR with predefined quantile outputs. In the introduction of SQR, Tagasovska and Lopez-Paz (2019) proposed a penalty function to prevent quantile crossover, which was built upon in future work (Gupta et al. 2019; Shen et al. 2024). However, penalty functions do not provide a guarantee against crossovers. Providing such guarantees, Song et al. (2024) added a positive last layer to a QR-RNN to prevent quantile crossover in the context of a QR post-processing step for ensemble weather forecasts. Similarly, positive parameters are also used to prevent quantile crossover of QR in forecasting of wind power generation (Cui, Wan, and Song 2023). Unfortunately, positive parameter-based approaches are not transferable to SQR as they rely on a predefined set of quantiles. Similarly, Park et al. (2022) model a set of quantiles as an integral with non-negative elements to ensure monotonicity, and can use interpolation to extrapolate on non-trained quantiles. However, this does not model a complete CDF, but just approximates it.

It is also possible to simply sort the quantiles as a post-processing step (Chernozhukov, Fernández-Val, and Galichon 2010; Smyl et al. 2026). While this can reportedly increase performance, it discards the trained quantile relations

as part of the output of the function approximator.

There also exist more general monotonic network approaches, competing and featuring different trade-offs compared to DLNs. Monotonic networks can guarantee monotonicity with respect to one (or more) input feature(s), which is all that is needed to prevent quantile crossovers in SQR. Certified monotonic neural networks (CMNNs) (Liu et al. 2020) reported to offer better performance than DLNs, while constrained monotonic neural networks (Runje and Shankaranarayana 2023) were shown to beat CMNNs. Scalable monotonic neural networks (SMNNs) (Kim and Lee 2023) outperform CMNNs and DLNs. In their testing, SMNNs perform similarly to, or slightly better than, constrained monotonic neural networks. However, monotonic network comparisons do not take place in a time series forecasting context, which makes it difficult to know if they could and would work with a forecasting task.

Numerous studies in probabilistic forecasting and monotonic networks have explored different ways that can be used to prevent quantile crossover. There is also significant interest in non-parametric probabilistic forecasting from an applied forecasting perspective. Comparatively, however, there is very little work on combining these probabilistic forecasting and monotonic networks in approaches that try to forecast an inverse CDF with a network that is guaranteed to be monotonic with respect to its quantile input.

## Contributions

In this paper, we propose an adaptation of a DLN to time series forecasting with SQR. This represents a missing link between probabilistic forecasting or non-parametric CDFs and monotonic constraintable networks.

We test this novel approach on a solar irradiance forecasting task. Solar irradiance is a particularly relevant application of weather forecasting techniques to enable the day-ahead prediction of photovoltaic energy production (of which solar irradiance data is an excellent surrogate, e.g. (Antonanzas et al. 2016)). The increasing prevalence of renewable energy sources and their natural, weather-related variations necessitate making predictions about them (Yang et al. 2022). Accurate probabilistic forecasting is shown to play an increasingly important role in the general methodology. Ramps, high changes in observed irradiance due to varying cloud coverage, may only in some cases be possible to capture by point forecasters, but can be captured by a CDF of a probabilistic forecaster (Samu et al. 2021). Predictions taking uncertainty and risk into account enable more informed bidding decisions in the energy market and make photovoltaic production more economically efficient. By extension, probabilistic forecasting may facilitate even higher investment in this area (Morales et al. 2014).

With this application, we show that we can not only connect existing fields in a novel way but also show that the resulting approaches may be directly applied in relevant applications. We hope that our work can facilitate further interaction between two largely independent research streams, applied forecasting and monotonic networks, which may lead to further improvements in the future. Specifically, the list of our contributions is as follows:

- We adapt a DLN to time series forecasting with SQR, which is the first adaptation of a DLN to this task.
- We detail the implementation and special considerations needed to fit DLN to time series forecasting with SQR.
- We compare our flavor of DLN to popular solar irradiance forecasting approaches and a forecasting adaptation of SMNNs, which have been shown to outperform DLNs in other tasks. In these comparisons, our iteration of a DLN shows better results, highlighting the applicability of Lattices for use in probabilistic forecasting.

## 2 Methods

### Forecasting Task

In solar irradiance forecasting, the goal is to predict a horizon of future values given a window of past observations. This could be a useful forecasting application for owners of a photovoltaic power plant, looking to optimize their sales on the day-ahead energy market (Samu et al. 2021). Formally, we define the forecasting task as given a series of  $\chi$  past values containing  $F$  features, predict a series of  $\gamma$  future values. The past values are defined as a window  $w$  of  $\chi = (x_{t-w}, \dots, x_t)$ , with  $\chi \in \mathbb{R}^{w \times F}$ , while the future values are defined as a horizon  $h$  of  $\gamma = (y_{t+1}, \dots, y_{t+h})$  with  $\gamma \in \mathbb{R}^{h \times 1}$ . The trainable forecaster is a function  $f$  mapping the window of past values to the horizon of future values, i.e.

$$\gamma = f(\chi) \quad (1)$$

As a potential application to the day-ahead energy market, we choose a forecasting horizon  $h$  of 36 hours and 96 as the past window  $w$ , i.e., 4 days.

### Quantile Regression

QR in time series forecasting (even if monotonically constrained) is usually trained to output a list of predictions containing a predetermined set of quantiles of interest, e.g.,  $[0.1, 0.25, 0.5, 0.75, 0.9]$ . This can be advantageous against parametric approaches, as we only learn a categorical subset of the underlying distribution. However, it also contains less information than parametric approaches, as the resulting categorical distribution is not a complete CDF.

In QR, we minimize the pinball loss  $l_\tau(\gamma, \gamma')$ , with  $\gamma$  representing ground truth, and  $\gamma'$  representing the prediction of a forecaster. Pinball loss is dependent on a quantile level  $0 \leq \tau \leq 1$ . The loss function looks as follows:

$$l_\tau(\gamma, \gamma') = \begin{cases} \tau(\gamma - \gamma') & \text{if } \gamma - \gamma' \geq 0, \\ (1 - \tau)(\gamma' - \gamma) & \text{else.} \end{cases} \quad (2)$$

Usually, we would now define a set of  $\tau$ -values and evaluate model output by iterating through and summarizing the pinball loss depending on  $\tau$ .

SQR is an adaptation of default non-parametric quantile regression (Tagasovska and Lopez-Paz 2019; Dabney et al. 2018). Instead of predetermining a set of quantiles to learn, we add a random quantile to the input of a trainable function, such as a neural network, and learn to predict that. This approach enables a neural network to learn an implicit continuous quantile distribution.

With SQR, we draw  $\tau \sim U(0, 1)$  during training and use it as input into the model and loss function  $f(\chi, \tau) = \gamma'$ . Now we can update equation 1 as predicting a probability distribution when aggregating a set of  $T = (\tau_0, \dots, \tau_q)$  with  $q \in \mathbb{N}$  into a set of output values  $Y \in \mathbb{R}^\tau$  shaping the inverse CDF :

$$Y = (f(\chi, \tau_0), \dots, f(\chi, \tau_q)) \quad (3)$$

In training, we do not need to consider a set of  $Y$ , as we are drawing a different  $\tau$  for every minibatch. Applying minibatches and combining equation 3 and equation 2 we can thus optimize  $f$  as follows:

$$\arg \min_f \frac{1}{N} \sum_{i=1}^N E_{\tau_i \sim U(0,1)} l_\tau(f(\chi_i, \tau_i), \gamma_i) \quad (4)$$

$N$  here being the batch size.

Intuitively,  $f$  is now an inverse and complete CDF with respect to  $\tau$  (Narayan et al. 2023). Nevertheless, QR via neural networks as described above does not protect against quantile crossovers, which are violations of the monotonicity constraint of quantiles. Quantiles of a probability distribution are monotonically increasing if,

$$f(\chi, \tau_a) \leq f(\chi, \tau_b) \text{ for } \tau_a \leq \tau_b. \quad (5)$$

A solution to quantile crossover is to use a function approximator that is monotonically constrainable with respect to (at least) one of its inputs, such as DLNs. Only by preventing quantile crossover can we produce a legitimate inverse CDF (Narayan et al. 2023).

### DLNs

DLNs are described as monotonically constrainable look-up tables (Gupta et al. 2016). A one-dimensional lattice function maps an input to an interpolated piecewise linear function. A higher-dimensional lattice function represents the same piecewise-linear function, but with a  $D$ -dimensional input.  $D$ -lattices will have  $k^D$  number of trainable parameters, where  $k$  represents the knots, or keypoints, in the piecewise linear function. In other words, a lattice is a hypercube, spanned by its input features, such that  $[0, 1]^D$ . This property enables it to map all of its inputs to an output value, just as a lookup table would. It achieves a smooth and nonlinear function by interpolating inputs between its vertices. Lattices can be trained using backpropagation and are shown to be able to approximate any bounded, continuous function if they are supplied with enough knots in their piece-wise function. In mathematical terms, a lattice maps  $f(\chi) : \mathbb{R}^D \rightarrow \mathbb{R}$  in which

$$f_l(\chi) = \theta^T \psi(\chi) \quad (6)$$

$\theta$  represents the values of the lattice and  $\psi(x)$  the interpolation of the input over the interpolation weights (Narayan et al. 2023). In Gupta et al. (2016), the interpolation step of a lattice is tested with a multilinear or a simplex approach. Here we only use multilinear. With  $v_k \in [0, 1]^D$  being the  $k$ th vertex of the unit hypercube, the interpolation weight on  $v_k$  is shown as:

$$\psi_k(\chi) = \prod_{d=0}^{D-1} \chi[d]^{v_k[d]} (1 - \chi[d])^{1-v_k[d]} \quad (7)$$

As  $v_k$  is either 0 or 1, it determines that  $\chi$  is multiplied either as  $\chi[d]$  or  $1 - \chi[d]$ .

Intuitively, lattices can be applied similarly to regular neural networks, although featuring some prominent differences: A single lattice only ever outputs a scalar, even if the input is a vector (i.e., in practice, the output is a vector of the same size as the mini-batch size). As mentioned above, Lattices’ complexity scaling is exponential with respect to the input dimensions  $D$  ( $k^D$ ). A large  $D$  quickly becomes unfeasible to fit onto modern computers, which is why Lattices are usually applied to problems with a limited number of features. One solution to this is to use an ensemble in which every lattice only takes a small subsample of the overall input features (You et al. 2017). In this ensemble, low-dimensional lattices are initialized to a random subset of features and then trained.

The originally proposed DLNs, which leverage such ensembles, can be built up of several distinct layers as described in detail by You et al. (2017).

- **Constrainable linear embedding layer:** Featuring two matrices and one bias, which linearly embeds a monotonic input. One matrix is used for monotonic inputs, in which a monotonic constraint is applied by keeping its weights restricted to be non-negative. The other matrix is used for non-monotonic inputs. If splitting between monotonic weights  $\theta^q \geq 0$  and non-monotonic weights  $\theta^m$ ,

$$\gamma = \begin{bmatrix} \theta^q \chi^q \\ \theta^m \chi^m \end{bmatrix} + b \quad (8)$$

This has been used in conjunction with calibrators (described below) as a form of input layer.

- **Calibration layer:** A layer of parallel one-dimensional lattice functions, embedding each feature separately. Any calibrator is a one-dimensional lattice with  $k$  key-value pairs mapping input to output ( $a \in \mathbb{R}^k, b \in \mathbb{R}^k$ ). Two values of  $b$  that correspond to  $a$  values close to the input are interpolated to get the output. This layer is used as a preprocessing step or to give the DLN more expressiveness with respect to every feature.
- **Lattice ensemble:** A set of lattices  $S$  in which each lattice  $f_l : \mathbb{R}^A \rightarrow \mathbb{R}, A \subset D$ , in which all  $A$  are unique with respect to each other  $A \neq A'$ . Outputs of this layer can be sent into another lattice ensemble until the output becomes scalar, be summarized, e.g., by averaging, or be sent through a final calibration step.

Even though these layers are suggested to be used in a certain way in DLNs, they are all monotonically constrainable and may be stacked freely in any order.

**Monotonicity of DLNs** In contrast to neural networks, lattice monotonicity constraints are relatively easy to enforce. As long as every pair of adjacent linearly interpolated lookup table parameters  $\theta_s$  and  $\theta_r$  obey  $\theta_s > \theta_r$  in the same direction, then the lattice is monotonic (Gupta et al. 2016) (also see Gupta et al. (2016) for the formal proof of monotonicity of lattices which applies the calibration layer and the lattice ensemble). All three different DLN layers, the

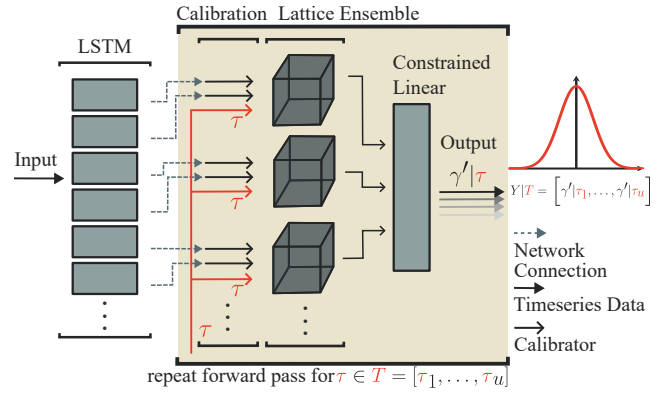


Figure 1: Architecture of the LSTM-DLN for SQR. If the DLN is looped over a set of quantiles, a correspondingly detailed approximation of a CDF can be accumulated.

constrainable linear, the calibration layer, and the lattice ensemble, preserve monotonicity across layers by treating the monotonic output from previous layers as monotonic input in subsequent layers (You et al. 2017). In practice, lattice monotonicity is enforced by solving the constrained optimization with Dykstra’s projection algorithm (Narayan et al. 2023).

### DLN adaptation to Time Series Forecasting

Even with the changes to lattices introduced by DLNs, they cannot, out of the box, be applied to higher-dimensional forecasting problems. First, ensembles are not a perfect tool to shrink the complexity scaling of  $k^D$ . Lattices with a smaller  $k$  or  $D$  might be too small as they will be less expressive, or not be able to model many feature interactions. Too big lattices will run into a complexity bottleneck. Second, DLNs’ scalar output size limits architectural freedom when needing to meet a predefined horizon of forecasts. One lattice function only maps to one scalar output.

Fortunately, we do not need a DLN to interpret time series data; we only need it to produce an implicit CDF of the forecast. As such, we can embed higher-dimensional data into lower dimensions, with e.g. a neural network, and use that as input to a DLN. This idea was first discussed in Narayan et al. (2023), but never actually tested. Long-short-term memory (LSTM) models are very popular and perform relatively well in the solar irradiance forecasting (e.g. (Qing and Niu 2018; Srivastava and Lessmann 2018; Huang et al. 2021)) For this work, we select an LSTM as a trainable embedding function for the time series data. This is an interchangeable decision, so more sophisticated models like Deep-AR (Salinas et al. 2020), or Transformers (Vaswani et al. 2017) may result in even better forecasts. However, LSTMs show great performance in forecasting tasks and are a good experimental baseline for our intended comparison of showing that DLNs can be used to accurately forecast implicit CDFs. In addition to an LSTM, an ensemble of lattices is used to minimize the input-dependent size of the DLN.

Conversely, fitting the output of a DLN to a predefined horizon can be achieved by employing a set of parallel lat-

---

**Algorithm 1: Training of the proposed SQR Architecture.**

---

**Input:** Collection of minibatches of past and target,  $(X, \Gamma)$ **Output:** Trained LSTM  $f_1$  and DLN  $f_2$ 

```
1: Initialize  $f_1$  and  $f_2$ 
2: for  $e \leftarrow 1$  to  $E$  do
3:   Sample  $\chi_j, \gamma_j$  from  $(X, \Gamma)$ 
4:   Sample  $\tau_j$  from  $\mathcal{U}(0, 1)$ 
5:   Obtain embedded values  $B$  from  $f_1(\chi_j)$ 
6:   Obtain  $\gamma_j^{\tau}$  from  $f_2(B, \tau_j)$ 
7:   Compute loss  $L(\gamma_j, \gamma_j^{\tau}, \tau_j)$  with eq. 2
8:   Apply backpropagation to  $f_1$  and  $f_2$  with  $L$ 
9: end for
10: return  $f_1$  and  $f_2$ 
```

---

---

**Algorithm 2: Exploitation step of the proposed SQR Architecture.**

---

**Input:** Time series of past values  $\chi$  and a set of quantiles  $T = (\tau_0, \dots, \tau_q)$ **Output:** Set of forecasts  $Y = (f(\chi, \tau_0), \dots, f(\chi, \tau_q))$ 

```
1: Initialize empty set  $Y$ 
2:  $B \leftarrow f_1(\chi)$ 
3: for  $\tau$  in  $T$  do
4:    $Y_\tau \leftarrow f_2(B, \tau)$ 
5: end for
6: return  $Y$ 
```

---

tices that take the same input and produce a different output timestep, or using the constrained linear embedding proposed by You et al. (2017) as a constrained linear output layer. The latter is much preferred, as a linear output layer also permits us to reduce the usage of high-dimensional lattices to a minimum and only use a single layer of ensembles.

Finally, we change the quantile input as proposed by Narayan et al. (2023). Instead of feeding quantiles into only one lattice in the ensemble and letting constraints be propagated throughout the ensemble, we reserve one input feature in every lattice in the ensemble for quantiles. This way, every feature receives a directly trained relation to the implicit, complete CDF. Additionally, calibrators are added before and after the DLN layer. The resulting architecture is shown in figure 1. We refer to the LSTM as the input model  $f_1$  and the DLN as the output model  $f_2$ .  $f_2$  needs one forward pass to model one quantile of the complete CDF. During training,  $f_1$  and  $f_2$  are called in sequence for each randomly drawn quantile, as can be seen in algorithm 1. When desired (for example, during an exploitation step),  $f_2$  can also be run for every quantile of interest, e.g., to form a highly detailed approximation of a CDF. This is shown in figure 1 and algorithm 2.

## Dataset

The solar irradiance dataset consists of five years (2016-2020) of hourly surface downwelling shortwave radiation

measurements from 20 stations around southern Norway<sup>1</sup>. As additional features, satellite-based irradiance estimates (from CAMS-RAD<sup>2</sup>) and estimated weather features (from NORA3<sup>3</sup>) from the coordinates of all 20 sites are added. The weather data taken from NORA3 consists of integrated surface net downward short-wave flux, relative humidity, wind speed and direction, snowfall, precipitation, temperature, cloud area fraction, and surface air pressure. A clear sky index, indicating irradiance on a cloudless day, is added from CAMS McClear<sup>4</sup>. Lastly, time features for hour of day, day of week, and week of year are embedded in a circular sine and cosine component and added to the dataset. Overall, this results in 12 features per 20 stations with an additional 6 embedded time features, adding up to 246. One station (station 11, near the west coast of Norway) is chosen as the target for the forecaster, as it has the least amount of missing data. Other missing data were mostly night values and were replaced with zeroes. Norway experiences much longer days with much higher irradiance amplitudes in summer than in winter, which models will have to adapt to. Night values are often removed from the dataset as irradiance is zero. However, as the window of historical values used for a forecast is fixed and longer than a day, we need these values to keep the relative days of available data consistent.

**Data Preprocessing** The respective data sources provide cleaned data sets, so only a min-max scaling is applied to all features. Additionally, the dataset is split into a training, validation, and testing set. Testing and validation received one year each (2016 and 2017, respectively) to be able to infer performance on a whole year of data. 2017 and 2016 were the least correlated to the rest (2018-2020, which was used as a training set) to increase forecast difficulty.

## Evaluation

We evaluate our adapted DLN on a set of 11 equidistant quantiles within the interval:  $[0.025, 0.975]$ . Training and Hyperparameter optimization are done on the training and validation data, while final performance testing is done on the test set.

**Comparisons** We designed the adapted DLN to be able to provide more information about a forecast by forecasting an implicit complete CDF. We consider this successful if we can manage to achieve similar performance to unconstrained versions of the same general architecture. This is the case, as unconstrained architectures can be expected to be more flexible than constrained ones, thus should return better results if the architecture can return well-calibrated CDFs. So, we are interested in testing only different versions of  $f_2$  while  $f_1$  is retrained, but with hyperparameters kept constant across different  $f_2$ .

---

<sup>1</sup><https://thredds.met.no/thredds/projects/sunpoint.html><sup>2</sup><https://www.soda-pro.com/web-services/radiation/cams-radiation-service><sup>3</sup><https://thredds.met.no/thredds/projects/nora3.html><sup>4</sup><https://www.soda-pro.com/web-services/radiation/cams-mcclear>

We hypothesize that SQR requires a nonlinear model in the output. Thus, the first comparison is made with a linear and a constrained linear embedding layer as a singular output layer. Next, we need to test our constrained approach against an unconstrained neural network to see if our choice in architecture mitigates the decline in accuracy due to constraints. Further, we identified SMNNs as one of the most accurate state of the art monotonic networks (Kim and Lee 2023), reportedly showing better results than regular DLNs in several regression tasks, making the current comparison very relevant.

We also wanted to compare performance to non-SQR settings. For this purpose, we use two more LSTM-Linear models, one point predictor trained with mean absolute loss and a QR with a predefined set of 11 quantiles. These are intended to show potential trade-offs made by using a probabilistic model with modeling capacity reserved to represent the implicit CDF instead of a set of discrete quantiles.

Lastly, in solar irradiance forecasting, smart persistence is seen as the most consistent baseline forecaster, and is used as a performance validation for more sophisticated models (van der Meer, Widén, and Munkhammar 2018). Smart persistence assumes that over a short time, irradiance is not changing significantly, and combines it with future information about irradiance under clear sky conditions. If a model cannot at least match the accuracy of smart persistence, it is considered ill-adjusted.

$$f(t) = \frac{x(t)}{o(t)} * o(t+h) \quad (9)$$

With  $x(t)$  being observed irradiance, and  $o(t)$  being clear sky irradiance, we estimate future irradiance (in  $h$  steps) based on the difference in observed and clear irradiance at  $t$ . In our condition of forecasting 36 hours into the future, we use the last day (i.e., 24 hours) as our source of information for the next one-and-a-half days. First, the previous 24 hours predict the next 24 hours. Then, the remaining 12 hours of the forecast are predicted by the first half (12 hours) of the 24 hours of the previous day. This way, each timestep is predicted by a previous day’s timestep at the same time of day.

**Metrics** All models are evaluated following many of the recommendations for probabilistic solar irradiance forecasting evaluations (Lauret, David, and Pinson 2019). Further, it should be noted that all metrics are calculated as the average across the length of the complete forecast horizon. This averaging is omitted from the equations for brevity.

First, we evaluate the point prediction metrics mean absolute error (MAE), root mean squared error (RMSE), and skill score (SS), which are evaluated on the 0.5 quantile values. SS is the fraction between  $MSE_f$  of the LSTM-DLN and  $MSE_p$  from a smart persistence model, such that:

$$SS = 1 - \frac{MSE_f}{MSE_p} \quad (10)$$

Further, we examine continuous rank probability scores (CRPS) as a probabilistic representation of distributional ac-

curacy. CRPS is usually computed in continuous form as follows,

$$CRPS = \frac{1}{N} \sum_{i=1}^N \int_{-\infty}^{+\infty} (\hat{F}_i(x) - 1_{x\gamma_i})^2 dx \quad (11)$$

In our case, we employ an approximated version of CRPS computed by the sum of pinball loss over all evaluated quantiles (Narayan et al. 2023):

$$CRPS = \frac{1}{N} \sum_{i=0}^N l_{\tau}(\gamma, \gamma') \quad (12)$$

Examining the coverage properties of the CDF, we plot prediction interval coverage percentage (PICP) and reliability diagrams in addition to computing the average coverage error (ACE) score. PICP is a percentage of interval points correctly captured in each predication interval  $c$  between the points  $[\hat{I}_{0.5-\frac{c}{2}}, \hat{I}_{0.5+\frac{c}{2}}]$ , which are the pair of opposing quantiles. PICP is better the closer it is to the target percentage coverage of its respective interval.

$$PICP_c = \frac{1}{N} \sum_{i=1}^N s_i, s_i = \begin{cases} 1, & \text{if } \gamma_i \in [\hat{I}_{0.5-\frac{c}{2}}, \hat{I}_{0.5+\frac{c}{2}}] \\ 0, & \text{otherwise} \end{cases} \quad (13)$$

ACE is the average difference of all PICP scores (and thus all intervals  $C$ ) and the target percentage coverage of their respective interval.

$$ACE = \frac{1}{|C|} \sum_{c \in C} |c - PICP_c| \quad (14)$$

Lastly, reliability diagrams are calculated just as PICP, but instead of counting values between intervals, we count values correctly captured within a quantile level.

### Implementation

Model training, optimization, and testing are done on a system with an AMD EPYC 7702 64-Core CPU and an Nvidia RTX3090 GPU. All code is written in Python with PyTorch (Paszke et al. 2019) and the package pytorch-lattice (Bakst 2024).

All models were optimized for hyperparameters in a sequential (coarse-to-finer) grid search with subsequent manual adjustment and visual inspection of probability distribution fit. General performance was assessed with CRPS and ACE, but additional metrics were also taken into account. A detailed description of hyperparameter optimization is given in Appendix A, while optimized model parameters are detailed in Appendix B. Some hyperparameters were kept constant for brevity and based on experience in performance development. Constant hyperparameters are a batch size of 64 and the Adam optimizer. All probabilistic models were trained using pinball loss from equation 2.

## 3 Results

Metric performance observations based on the test dataset are shown in table 1. Scores of the models are based on the

Models	CRPS	MAE	RMSE	ACE	SS
SP	-	52.903	96.472	-	-
LSTM-PP	-	43.152(0.336)	74.933(0.602)	-	0.397(0.019)
LSTM-QR	30.506(0.083)	43.875(0.226)	74.949(0.343)	0.052(0.002)	0.396(0.011)
SQR: LSTM-Lin	24.427(0.247)	45.556(0.360)	77.165(0.414)	0.173(0.003)	0.360(0.014)
SQR: LSTM-NN	27.690(0.099)	42.612(0.525)	72.930(1.009)	0.075(0.003)	0.429(0.032)
SQR: LSTM-CLin	25.920(0.184)	44.619(0.294)	75.472(0.397)	0.190(0.003)	0.388(0.013)
SQR: LSTM-SMNN	29.333(0.657)	47.548(0.974)	81.731(1.691)	0.079(0.657)	0.282(0.059)
SQR: LSTM-DLN	29.175(0.216)	43.266(0.534)	74.520(0.850)	0.061(0.003)	0.403(0.027)

Table 1: Mean model performance comparison over 5 training runs with a different random seed. Standard deviation is displayed in brackets. SP is smart persistence, PP is the point predictor, and QR is regular quantile regression. SQR models use simultaneous quantile regression with an LSTM as  $f_1$  and different  $f_2$ . L is a linear layer, NN is a non-linear neural network, CL is a constrained linear layer, SMNN is a scalable constrained neural network, and DLN is a deep lattice network.




Figure 2: Two 36h forecasts of the test data from LSTM-DLN (row 1), LSTM-SMNN (row 2), and LSTM-NN (row 3). Shown are 11 quantiles between 0.025 and 0.975. Axis labels are consistent between subplots.

average metric value over a horizon of 36 hours ahead. All metrics displayed are better when smaller, except SS, which is better when larger. Results are the average performance over five optimization runs with different random seeds. In parentheses, the standard deviation of the metric is shown. Generally, all models display positive SS and thus are better forecasters than the standard baseline. In MAE and RMSE, the unconstrained LSTM-NN and our LSTM-DLN are not significantly different. The point predictor (LSTM-PP) is also not significantly different from LSTM-NN and LSTM-DLN in MAE, but shows a significant performance decline in RMSE compared to LSTM-NN. SQR approaches showing better performance may be due to regularization caused by representing all quantiles. In the ACE score, LSTM-QR

performs best, with the second best being the LSTM-DLN. CRPS results are more complex. Linear and constrained linear appear to be the best performers. However, this is due to the way CRPS is approximated here (equation 12). In pin-ball loss (equation 2), any value of  $\gamma'$  close to  $\gamma$  will result in lower loss, no matter the quantile value. As such, this loss can favour a distribution in which all quantiles have the same value close to the median of the target values. In both linear conditions, the CDF is observed to be small in width. This is why their respective ACEs are much worse than all others and appear to be ill-calibrated. Their performance in MAE and RMSE is similar to the non-linear competitors (LSTM-NN, SMNN, DLN), likely due to the networks being trained to focus on the median quantile. Due to this ill-calibration, they are generally performing worse.

LSTM-SMNN displays slightly worse performance than LSTM-DLN in MAE and RMSE, while not being significantly different in CRPS and ACE. In ACE, this is the case, as its ability to form a wide and well-calibrated CDF appears to be inconsistent across random seeds. Even for well-adjusted examples, LSTM-DLN displays better results.

Figure 2 shows two different full 36h forecasts from NN, SMNN, and DLN. All plots in the same column are from the same timestamp. We can observe that the fit is not perfect for either of them, as the forecasting of changes in irradiance from one day to the next is not trivial. DLN shows the largest CDF width, which seems to help capture observations within its CDF. DLNs also seem to display a more skewed and non-Gaussian distribution compared to NN. SMNN shows more erratic behavior and poorer fit, reflected by its worse performance metrics.

In figure 3, we show CDF coverage performance for all models across intervals and quantiles. Ideal performance is marked by the black bar. Apart from constrained linear (CLin) and linear (Lin), all models seem to be close to the ideal. This echoes the previous observation that both linear approaches are ill-calibrated even though their median fit is good. In terms of interval coverage, the smaller intervals closer to the median show slight overestimation, while the lateral intervals do not always include all extreme values. In addition to mistakes in the extreme values, the reliability diagram shows quantiles between 0.2-0.5 to be more underestimated. This likely means that even the better-performing




Figure 3: PICP (left) and reliability diagram (right). PICP starts at approx. 0.2 as that is the smallest interval around the median. The values are means across all 5 models, with standard deviations displayed as colored areas around the curve.

models make mistakes in the median of the CDF, leading to less correct values within the lower and higher quantiles around the median, which is not captured by the interval.

LSTM-NN, as the most well-performing competitor to our LSTM-DLN, is unconstrained with respect to quantiles. The overall average percentage of quantile crossovers displayed by the LSTM-NN over the whole test set and all quantiles is 3.075 % with a standard deviation of 0.0744. On the other hand, DLNs are still much more complex in terms of parameter numbers than other networks. The DLN implementation has more than twice the number of parameters as SMNN and takes more than 2s for every forward pass when run on a local laptop. Exact times and trainable parameter numbers for all models can be seen in Appendix C.

## 4 Discussion

This work aimed to show the effectiveness of using a DLN as a CDF multi-horizon forecaster. We do this by comparing against unconstrained approaches, as they are more flexible, thus should perform better, and a state of the art monotonic network approach, which showed better results than DLNs in a different application context.

The results show that our adaptation has similar performance to the unconstrained alternative, but shows better coverage in its quantiles. In conclusion, the expected limited expressiveness in the DLN induced by monotonic constraints is successfully mitigated by our adaptations. Our DLN adaptation also shows higher performance than an SMNN-based implementation, highlighting the competitiveness of this approach for time series forecasting within the state of the art of monotonic networks. Further, we show the necessity of a non-linear forecaster by comparing it against a linear representation of the implicit CDF, as LSTMs with both types of linear output layers display notably worse calibration compared to all other models. Well-performing SQR approaches do not show many tradeoffs to point predictor or predetermined QR. QR shows an edge in coverage performance; however, the other top-performing models also have very good coverage.

Overall, our approach represents a successful way to apply DLNs on multi-horizon time series forecasting of CDFs.

It should be highlighted that this implementation is highly modular, and we would encourage further investigation with other embedding models, more comparisons across quantile modelling techniques, or the usage of the DLN output layer in different contexts. In its current implementation, the SQR-based DLN can be considered useful for highly detailed modeling of a variable-sized (partial) inverse CDF. Contrary to approaches outside of monotonic networks, it can also be used to constrain more features (than just quantiles) monotonically. Lastly, an implementation that does not need further embedding before applying a Lattice may benefit from its increased explainability. In contrast to neural networks, lattices are interpolated look-up tables that model the whole range of possible outputs. As such, relations between input and output features are not hidden in a black-box and may even be plotted.

Nevertheless, DLNs still suffer from a very high complexity in the number of trainable parameters and execution time, the decrease of which remains an ongoing research question.

## 5 Conclusion

In this work, we show that monotonous network research can be successfully connected with probabilistic forecasting applications. DLNs can be applied to time series forecasting, yielding a complete, implicit CDF with performance comparable to that of the unconstrained alternative and surpassing the state of the art monotonic network competitor.

## Acknowledgements

We would like to thank our friends and colleagues at the Department of Technology Systems and the Institute for Energy Technology for the (ongoing) helpful discussions and support. We would also like to thank the anonymous reviewers for their insightful feedback and enriching comments on the draft of this work.

## References

- Abdar, M.; Pourpanah, F.; Hussain, S.; Rezazadegan, D.; Liu, L.; Ghavamzadeh, M.; Fieguth, P.; Cao, X.; Khosravi, A.; Acharya, U. R.; Makaremkov, V.; and Nahavandi, S. 2021. A Review of Uncertainty Quantification in Deep Learning: Techniques, Applications and Challenges. *Information Fusion*, 76: 243–297.
- Antonanzas, J.; Osorio, N.; Escobar, R.; Urraca, R.; Martinez-de-Pison, F.; and Antonanzas-Torres, F. 2016. Review of Photovoltaic Power Forecasting. *Solar Energy*, 136: 78–111.
- Bakst, W. 2024. `pytorch-lattice`. <https://github.com/willbakst/pytorch-lattice>. Accessed: 12/05/2025.
- Brando, A.; Center, B. S.; Gimeno, a. J.; Rodriguez-Serrano, J.; and Vitria, J. 2022. Deep Non-crossing Quantiles through the Partial Derivative. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, 7902–7914. PMLR.
- Chernozhukov, V.; Fernández-Val, I.; and Galichon, A. 2010. Quantile and Probability Curves Without Crossing. *Econometrica*, 78(3): 1093–1125.

- Cui, W.; Wan, C.; and Song, Y. 2023. Ensemble Deep Learning-Based Non-Crossing Quantile Regression for Nonparametric Probabilistic Forecasting of Wind Power Generation. *IEEE Transactions on Power Systems*, 38(4): 3163–3178.
- Dabney, W.; Ostrovski, G.; Silver, D.; and Munos, R. 2018. Implicit Quantile Networks for Distributional Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning*, 1096–1105. PMLR.
- Gasthaus, J.; Benidis, K.; Wang, Y.; Rangapuram, S. S.; Salinas, D.; Flunkert, V.; and Januschowski, T. 2019. Probabilistic Forecasting with Spline Quantile Function RNNs. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, 1901–1910. PMLR.
- Gneiting, T.; Balabdaoui, F.; and Raftery, A. E. 2007. Probabilistic Forecasts, Calibration and Sharpness. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 69(2): 243–268.
- Gupta, A.; Shukla, N.; Marla, L.; Kolbeinsson, A.; and Yellepeddi, K. 2019. How to Incorporate Monotonicity in Deep Networks While Preserving Flexibility? arXiv:1909.10662.
- Gupta, M.; Cotter, A.; Pfeifer, J.; Voevodski, K.; Canini, K.; Mangylov, A.; Moczydlowski, W.; and van Esbroeck, A. 2016. Monotonic Calibrated Interpolated Look-Up Tables. *Journal of Machine Learning Research*, 17(109): 1–47.
- Huang, X.; Li, Q.; Tai, Y.; Chen, Z.; Zhang, J.; Shi, J.; Gao, B.; and Liu, W. 2021. Hybrid Deep Neural Model for Hourly Solar Irradiance Forecasting. *Renewable Energy*, 171: 1041–1060.
- Kim, H.; and Lee, J.-S. 2023. Scalable Monotonic Neural Networks. In *The Twelfth International Conference on Learning Representations*.
- Koenker, R.; and Bassett, G. 1978. Regression Quantiles. *Econometrica*, 46(1): 33–50.
- Lauret, P.; David, M.; and Pedro, H. 2017. Probabilistic Solar Forecasting Using Quantile Regression Models. *Energies*, 10(10): 1591.
- Lauret, P.; David, M.; and Pinson, P. 2019. Verification of Solar Irradiance Probabilistic Forecasts. *Solar Energy*, 194: 254–271.
- Liu, X.; Han, X.; Zhang, N.; and Liu, Q. 2020. Certified Monotonic Neural Networks. In *Advances in Neural Information Processing Systems*, volume 33, 15427–15438. Curran Associates, Inc.
- Morales, J. M.; Zugno, M.; Pineda, S.; and Pinson, P. 2014. Electricity Market Clearing with Improved Scheduling of Stochastic Production. *European Journal of Operational Research*, 235(3): 765–774.
- Narayan, T.; Wang, S. L.; Canini, K. R.; and Gupta, M. 2023. Expected Pinball Loss For Quantile Regression And Inverse CDF Estimation. *Transactions on Machine Learning Research*.
- Park, Y.; Maddix, D.; Aubet, F.-X.; Kan, K.; Gasthaus, J.; and Wang, Y. 2022. Learning Quantile Functions without Quantile Crossing for Distribution-free Time Series Forecasting. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, 8127–8150. PMLR.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Köpf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: an imperative style, high-performance deep learning library. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Red Hook, NY, USA: Curran Associates Inc.
- Price, I.; Sanchez-Gonzalez, A.; Alet, F.; Andersson, T. R.; El-Kadi, A.; Masters, D.; Ewalds, T.; Stott, J.; Mohamed, S.; Battaglia, P.; Lam, R.; and Willson, M. 2025. Probabilistic Weather Forecasting with Machine Learning. *Nature*, 637(8044): 84–90.
- Qing, X.; and Niu, Y. 2018. Hourly Day-Ahead Solar Irradiance Prediction Using Weather Forecasts by LSTM. *Energy*, 148: 461–468.
- Runje, D.; and Shankaranarayana, S. M. 2023. Constrained Monotonic Neural Networks. In *Proceedings of the 40th International Conference on Machine Learning*, 29338–29353. PMLR.
- Salinas, D.; Flunkert, V.; Gasthaus, J.; and Januschowski, T. 2020. DeepAR: Probabilistic Forecasting with Autoregressive Recurrent Networks. *International Journal of Forecasting*, 36(3): 1181–1191.
- Samu, R.; Calais, M.; Shafiullah, G. M.; Moghbel, M.; Shoeb, M. A.; Nouri, B.; and Blum, N. 2021. Applications for Solar Irradiance Nowcasting in the Control of Microgrids: A Review. *Renewable and Sustainable Energy Reviews*, 147: 111187.
- Shen, G.; Jiao, Y.; Lin, Y.; Horowitz, J. L.; and Huang, J. 2024. Nonparametric Estimation of Non-Crossing Quantile Regression Process with Deep ReQU Neural Networks. *Journal of Machine Learning Research*, 25(88): 1–75.
- Smyl, S.; Oreshkin, B. N.; Peřka, P.; and Dudek, G. 2026. Any-Quantile Probabilistic Forecasting of Short-Term Electricity Demand: Fusing Uncertainties from Diverse Sources. *Information Fusion*, 127: 103637.
- Song, M.; Yang, D.; Lerch, S.; Xia, X.; Yagli, G. M.; Bright, J. M.; Shen, Y.; Liu, B.; Liu, X.; and Mayer, M. J. 2024. Non-Crossing Quantile Regression Neural Network as a Calibration Tool for Ensemble Weather Forecasts. *Advances in Atmospheric Sciences*, 41(7): 1417–1437.
- Srivastava, S.; and Lessmann, S. 2018. A Comparative Study of LSTM Neural Networks in Forecasting Day-Ahead Global Horizontal Irradiance with Satellite Data. *Solar Energy*, 162: 232–247.
- Tagasovska, N.; and Lopez-Paz, D. 2019. Single-Model Uncertainties for Deep Learning. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 576, 6417–6428. Red Hook, NY, USA: Curran Associates Inc.

- van der Meer, D. W.; Widén, J.; and Munkhammar, J. 2018. Review on Probabilistic Forecasting of Photovoltaic Power Production and Electricity Consumption. *Renewable and Sustainable Energy Reviews*, 81: 1484–1512.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention Is All You Need. arXiv:1706.03762.
- Wang, K.; Zhang, Y.; Lin, F.; Wang, J.; and Zhu, M. 2022. Nonparametric Probabilistic Forecasting for Wind Power Generation Using Quadratic Spline Quantile Function and Autoregressive Recurrent Neural Network. *IEEE Transactions on Sustainable Energy*, 13(4): 1930–1943.
- Wen, R.; Torkkola, K.; Narayanaswamy, B.; and Madeka, D. 2017. A Multi-Horizon Quantile Recurrent Forecaster. <https://arxiv.org/abs/1711.11053v2>.
- Yang, D.; Wang, W.; Gueymard, C. A.; Hong, T.; Kleissl, J.; Huang, J.; Perez, M. J.; Perez, R.; Bright, J. M.; Xia, X.; van der Meer, D.; and Peters, I. M. 2022. A Review of Solar Forecasting, Its Dependence on Atmospheric Sciences and Implications for Grid Integration: Towards Carbon Neutrality. *Renewable and Sustainable Energy Reviews*, 161: 112348.
- You, S.; Ding, D.; Canini, K.; Pfeifer, J.; and Gupta, M. 2017. Deep Lattice Networks and Partial Monotonic Functions. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.