MODELLING AND FORECASTING PV PRODUCTION IN THE ABSENCE OF BEHIND-THE-METER MEASUREMENTS

Tomas Landelius[1], Sandra Andersson[1] and Roger Abrahamsson[2]
Swedish Meteorological and Hydrological Institute[1], Tekniska Verken Linköping Nät AB[2]
601 76 Norrköping, Sweden[1], Brogatan 1, 582 78 Linköping[2]

ABSTRACT: This paper deals with the prediction of the net load from photovoltaic installations. The forecasts were only based on information from a numerical weather prediction model and measured net load. The study thus tackles the problem of estimating the contribution of PV power to the grid without knowing the actual production and consumption “behind-the-meter”. The main approach using a physical model was compared with persistence (forecast with yesterday’s values) and black box models represented by linear regression and an artificial neural network. The artificial neural network performed best with a normalized (with installed power) RMSE of about 10 %. Linear regression was only slightly inferior while the physical model performed worse (about 15 %). The physical model only predicts the gross production and needs information about the on site consumption in order to predict the net load. Here the consumption was predicted using persistence. This assumption did not hold up for the physical model to match the performance of the black box models. Thus, a better description of the gross consumption is needed in order to make the physical model more competitive.
Keywords: Modelling, Photovoltaic, Solar Radiation

1 INTRODUCTION

Meeting a growing worldwide demand for energy, while addressing climate change, is a challenge. One way of tackling it is to increase the generation of renewable energy, such as solar PV power. Renewable energy depends on weather conditions and high quality forecasts are required by transmission system operators, independent services operator (ISOs), grid owners and energy suppliers for managing the energy mix, maintaining grid balance and trading on the energy market. As an example, reducing the forecast error by just 0.1 percentage points is said to save the California ISO and California ratepayers more than $2 million per year [9].

In order to make a good PV power forecast for an installation, information about its orientation, installed peak power and geographical location is needed along with a reliable prediction of the relevant weather parameters. However, for most behind-the-meter (BTM) systems producing electricity primarily for on-site use, before delivering excess energy to the grid, only data about installed nominal power and the address is recorded.

Moreover, measurements of BTM PV generation are scarce. What is available is the measured net load which does not equal actual electricity consumption/production since some portion of it may come from BTM solar PV generation. This causes a disconnection between the measured load and the electricity demand, when the former is used in models to predict the electricity demand. This problem will be worsened with a growing portion of solar power and adoption of BTM storage, vehicle charging and time dependent electricity rates for instance, introducing more factors that influence the forecast accuracy. Further, the solar power is an unstable and difficult to predict source of energy which increases forecast errors and uncertainties in predictions used for planning and day ahead trading.

Here we describe a method to estimate parameters in a physical PV model (PVLIB Python, [5]). The estimation is based on hourly numerical weather prediction (NWP) data, net load measurements and an average daily cycle of the gross electricity consumption. Information about the latter is assumed to be obtained from net load measurements prior to the PV installation. The model is then used to forecast net electricity production (here defined as the gross production minus the gross consumption) under the assumption that the gross consumption can be predicted using persistence (only weekdays considered here). The result using this model was then compared with persistence and black box models represented by linear regression and an artificial neural network (ANN) using NWP data and measured net load from the previous day as input.

The aim of the study presented in the paper is to develop a model for predicting regional net PV power production and estimating its parameters when BTM PV measurements are unavailable. Ultimately the goal is to predict PV power production at different scales in a regional electricity network, e.g. in terms of total production, production per entry point or per secondary substation.

Previous work (e.g. [4, 12]) has tackled the problem with the general lack of BTM measurements of PV production by up-scaling information from a few well monitored representative sites. However, finding such representative sites could be difficult since there are no requirements on measuring the gross production.

In this study we do use BTM measurements of gross consumption. However, we assume that the daily cycle of the mean hourly gross consumption can be estimated based on data prior to the PV installation when BTM measurements are not available. Such consumption patterns are well-studied and known to be predictable. However, once a PV installation is in place the consumption pattern may very well change [9]. Our approach could be extended to take this into account by also including a model for the daily cycle of the average hourly consumption pattern in the optimization.

2 DATA

In this section we describe the data used for the study. It consists of two parts; measured energy production and consumption, along with NWP forecasts for the corresponding dates. Both data sets consist of hourly data covering the time period March - October 2016.

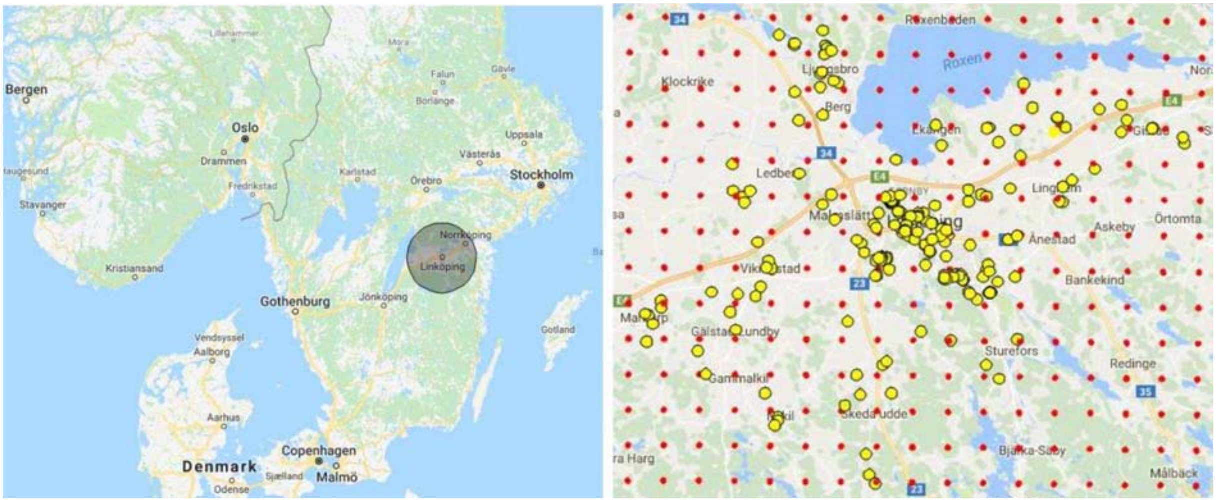

Figure 1: Left: approximate location of the Linköping network. Right: NWP grid points (red) and PV installations (yellow).

The data was divided into one training set, consisting of 80 % of the data (3360 hourly values per site), and one evaluation set with the remaining 20 % (840 hourly values per site). The division was made with a random sampling using a uniform distribution. This was done in order to end up with similar probability distributions for the training and evaluation data sets.

2.1 Measured production and consumption

The electricity measurements were provided by Tekniska verken that is responsible for the electrical grid in the Municipality of Linköping, parts of the Municipality of Mjölby and large parts of the Municipality of Katrineholm in Sweden. Hourly measurements of BTM PV production are available from 11 out of the 220 sites with PV installations in the Linköping net, see Fig. 1. Since we only had data from sites with PV installations and only from the year 2016 we had no information about the gross consumption prior to the installations. Instead we were restricted to use data from the 11 sites where this information was available. Two of these 11 sites were left out due to changes in their installations during the time period. The remaining 9 installations are located on the roofs of five households, four apartment complexes and one office building. For each installation information about the geographical location and the installed effect is also known.

2.2 NWP model data

NWP data was obtained from MetCoOp (Meteorological Co-operation on Operational NWP) where the meteorological services of Sweden, Norway and Finland run a common ensemble prediction system (MetCoOp EPS - MEPS). MEPS is developed in the framework of Aire Limitée Adaptation Dynamique Developpement InterNational (ALADIN) - High-Resolution Limited Area Model (HIRLAM) NWP system. The system can run with different configurations. Here we used the current cycle (40h1.1) of the HIRLAM-ALADIN Regional Meso-scale Operational NWP In Europe-Application of Research to Operations at Mesoscale (HARMONIE-AROME) [2]. The radiation model is based on the radiation scheme by Morcrette [10] and uses the Rapid Radiative Transfer Model of [8]. The model domain contains 900x960 points with 2.5 km grid spacing (see example in Fig. 1) and 65 levels covering a Nordic region. For this study we only used the deterministic MEPS control forecasts started at 00 UTC with a length of 24 hours. The parameters of our interest were accumulated hourly values of GHI and DNI together with instantaneous values of two meter temperature and 10 meter wind speed.

3 PV MODELS

To predict the net load we employed a physical model and included three statistical ones for comparison. By a physical model we here mean a model that simulates the physics of the PV power production based on information about the module characteristics, its orientation and the state of the atmosphere. The statistical models are here represented by persistence, linear regression and an ANN. In order to estimate the parameters of the physical model we need some information about the BTM production. This is not the case for the statistical models which can be fitted directly to the measured net load.

3.1 Physical

As a main approach, PVLIB Python was used to model the BTM PV power production. This is a open source community supported tool that provides a set of functions and classes for simulating the performance of photovoltaic energy systems. It was originally based on a toolbox developed at Sandia National Laboratories. For our purposes we chose a module and inverter that should match common installations in Sweden were selected from the Sandia library in PVLIB; SunPower SPR 220 and ABB MICRO 0 25 respectively.

The installations at each of the 9 sites with BTM measurements were modelled with a scaling factor times the output from one module in a single orientation. This approach could be seen as semi physical since the installation may consist of modules in different orientations and may be affected by shading. Moreover, it needs to be combined with a forecast for the gross error consumption in order to predict the net load. The tilt, azimuth and scaling factor (parameter vector w) was estimated by optimizing the summed squared error for the daily cycle of the mean hourly net production using a Nelder-Mead simplex algorithm (Python function

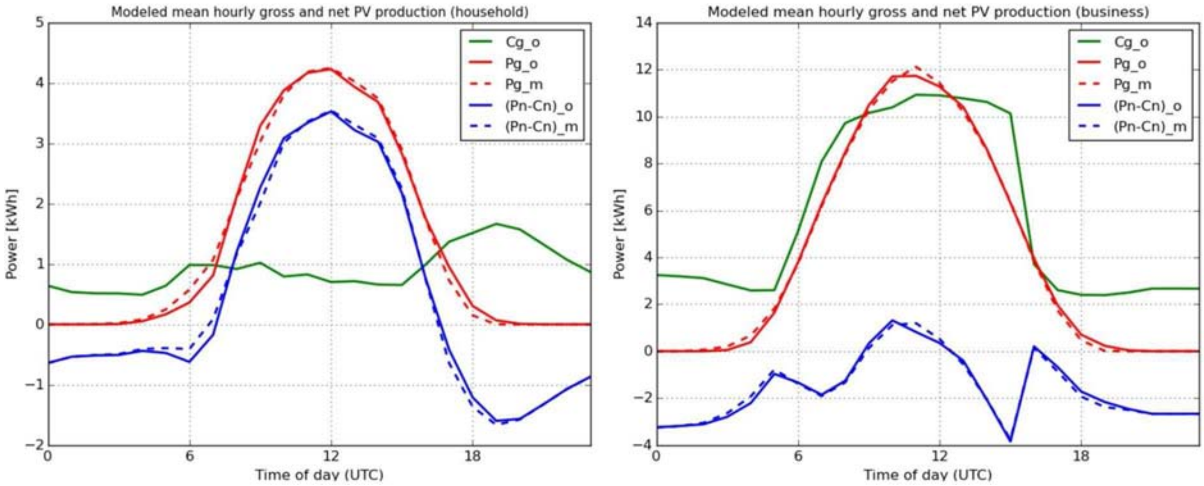

Figure 2: Left: Example of optimization with PVLIB targeting the daily cycle of the net production (solid blue line) and the resulting model values (dashed blue line) at a household installation. Predicted (dashed red line) and observed (solid red line) gross production and consumption (green line). Right: As in left panel but for a site at an office building.

scipy.optimize.fmin):

$$w^* = \operatorname{argmin} \sum \left(\bar{p}_g^m(w,t) - \bar{c}_g^o(t) - (\bar{p}_n^o(t) - \bar{c}_n^o(t))\right)^2. \quad (1)$$

Here the observed (superscript o) net (subscript n) and gross (subscript g) production and consumption for a given day and hour is related as:

$$p_g^o(d,t) - c_g^o(d,t) = p_n^o(d,t) - c_n^o(d,t). \quad (2)$$

Now, including the net consumption in the modelled (superscript m) net production (i.e. modelling the negative net load), the daily cycles of the mean hourly values are calculated as:

$$\bar{c}_g^o(t) = \frac{1}{n_d} \sum c_g^o(d,t) \quad (3)$$
$$\bar{p}_n^o(t) = \frac{1}{n_d} \sum p_n^o(d,t) \quad (4)$$
$$\bar{p}_n^m(w,t) = \frac{1}{n_d} \sum p_g^m(w,d,t) - c_g^o(d,t) \quad (5)$$

for the hourly intervals t = 0−1,…,23−24 UTC during the time period March – October 2016.

Here, the mean daily pattern of the gross consumption was calculated with data from the training data set. In practice, this information could be derived from measurements of the load (N.B. not BTM) during a time period before the PV installation was made.

Once the optimal PVLIB model parameters are estimated we have a physical model for the hourly gross (BTM) production. This model can then be used to predict the net PV production by assuming persistence for the gross consumption, i.e. that $c_g(d+1,t) \approx c_g(d,t)$. Hence the gross consumption can be eliminated from the equation and the net production predicted as:

$$p_n(d+1,t) = p_g^m(w,d+1,t) - c_g(d+1,t) \quad (6)$$
$$\approx p_g^m(w,d+1,t) - c_g(d,t) \quad (7)$$
$$= p_n(d,t) + p_g^m(w,d+1,t) - p_g^m(w,d,t). \quad (8)$$

The idea is that the daily pattern in the gross consumption is more stable than either the daily net or the daily gross production. This leads to a model for the net load that is given by persistence modified with the difference in gross production between the day ahead and today.

3.2 Statistical

Another alternative is to turn to statistical or black box models and model the net load directly. Here the idea is to fit a parametric model to the data. The assumption is that the residual is given by some model noise source when the optimal parameters have been estimated.

3.2.1 Persistence

The persistence model does not involve any parameter estimation and the input is the same as the output - the observed value for the corresponding hour from the previous day:

$$p_n^m(d+1,t) = p_n^m(d,t), \; t = 1, \ldots, 24. \quad (9)$$

3.2.2 Linear least squares

The linear and the ANN model have different parameters but share the same input. As input parameters (x) for these two models we used NWP forecast data (for the next 24 hours as well as for the past 24 hours) together with measured net load from the previous day. The idea behind this is that model should be able to make a connection between yesterday's state of the atmosphere and the net production. We also included the cosine of the solar zenith angle to provide some information about the time of the day and the season. It should also help to account for the effect of PV on tilted surfaces. For the linear model we also added a constant to the input vector in order for the model to be able to add a bias.

The linear model is given by

$$p_n^m(w,d+1,t) = w^T x. \quad (10)$$

We assume that the residual error is described by a normal distribution and employ a linear least squares method to estimate the model parameters (Python function numpy.linalg.lstsq):

Table I: Model performance in terms of mean r^2 values and RMSEn for gross PV production using PVLIB and persistence at the different sites (mean value, H:household, A: apartment block, O: office).

		Mean	H1	H2	H3	H4	H5	A1	A2	A3	O1
PVLIB	r^2	0.80	0.80	0.80	0.78	0.82	0.78	0.81	0.79	0.80	0.82
Persistence	r^2	0.45	0.46	0.45	0.48	0.42	0.44	0.45	0.43	0.43	0.46
PVLIB	RMSEn [%]	11	11	12	10	11	11	11	11	12	9.2
Persistence	RMSEn [%]	21	18	22	17	29	18	18	29	21	17

$$w^* = \text{argmin} \sum \sum (w^T x - (p_n^o(d,t) - c_n^o(d,t)))^2 \quad (11)$$

3.2.3 Artificial neural network

Using machine learning to train non-linear models has a long history within the area of energy forecasting [11, 6]. Here we used an off the shelf ANN from TensorFlow [1] to see if it a non-linear black box model could offer some improvements.

Using the standard feedforward ANN (DNNRegressor) from TensorFlow we set up a network with a three layer feedforward topology with one input, one hidden, and one output layer. In our set up it had 11 inputs, 32, 64 or 96 nodes in the hidden layer and 1 node in the output layer.

Determining the number of hidden neurons in the hidden layer(s) is a trade off between the networks ability to generalize from the training data (not too many neurons) and its representative power (not too few). Here we were guided by the empirical relation for the number of hidden layer neurons proposed by [6]: $n_{hl} = 0.5(n_{in}+n_{out}) + (n_{train})^{1/2}$, where n_{in} , n_{out} and n_{train} denote the number of input, output and the size of the training data set respectively. In our case the number of inputs equals 11 and we have one single output (the net load).

The number of cases in the training data set was 3360. Hence the suggested number of nodes in the hidden layer becomes 64. To check the robustness of this choice we also tried with 32 and 96 neurons in the hidden layer. The network with 32 neurons actually performed slightly better than the others on the evaluation data set so the results presented here are based on the outputs from that network.

In order to harmonize the amplitudes of the input variables we remove the mean (m) and normalize the input vector by multiplying it with the inverse of the square root sample covariance matrix (C), calculated from the training data set. The TensorFlow minimization algorithm then finds the solution to

$$w^* = \text{argmin} \sum \sum (w_o^T f(w_h, C^{-\frac{1}{2}}(x - m)) - (p_n^o(d,t) - c_n^o(d,t)))^2 \quad (12)$$

using an iterative procedure. We ran the minimization for 10,000 iterations (saving the result at each 100:th iteration) at which point the error for the evaluation data sets had started to increase for all sites. The prediction network was then given by the parameters from the iteration for which the evaluation error had a minimum.

4 RESULTS

In order to evaluate the performance of the different models we compared the models by looking at the root mean squared error (RMSE) normalized by the nominal installed power (RMSEn). We also calculated the square of the Pearson's correlation coefficient (r^2) between the modelled and observed net load. Only hours when the sun was over the horizon were included in the calculations.

First we look at the results for the main approach using PVLIB. Fig. 2 shows two examples of the fit of the optimized PVLIB model (blue dashed line) to the measured net production $(P_n-C_n)_o$ (blue solid line) which is the error criteria in equation 1. Using the model to predict gross PV production (dashed red line) results in a good fit to observed values (red solid line). The left and right panels show the results for a household and an office building respectively. Here the daily pattern for gross consumption (green lines) of the household shows a typical structure with peaks during the morning and afternoon while the other installation shows a consumption patterns related to the office hours.

Table I summarizes the error measures for the nine installations when using the PVLIB to model the gross production, which is an essential part of the physical approach. The average RMSEn for the PVLIB and persistence models are 11 % and 19 % respectively. Note that the error does vary much between different type of installations.

The daily cycles of the mean RMSEn when using PVLIB to model the gross consumption are illustrated in Fig. 3. The error corresponding to a persistence model is included for comparison. The error is shown as a function of forecast length. All forecasts are initialized at 00 UTC and the PVLIB based approach outperforms persistence for all forecast lengths.

We then turn to the prediction of the net load, now also including the statistical approaches. Table II summarizes the error measures for the nine installations when using the four different models; PVLIB, linear, ANN and persistence. The errors for both PVLIB and persistence are larger when predicting the net load compared to the errors for the gross PV forecasts. The reason is that the physical model now also needs information about the hourly gross consumption. Assuming persistence for the gross consumption turned out to be problematic. The r^2 values for such predictions varied between 0.10 and 0.83 for the nine sites. This shows that the gross consumption was less stable over time than we had thought.

The statistical models all perform better compared to the physical PVLIB approach. The ANN performs best on average with an r^2 value of 0.76 followed by the linear model (0.71), PVLIB (0.62) and persistence (0.40). The same order goes for the RMSEn where again the ANN is best with a mean value of 11 % followed by the linear model (12 %), PVLIB (16 %) and persistence (20 %). Note that there is a difference in the r^2 values between different sites for the net load predictions. The correlation for the office site is notably smaller for all the models even though only weekdays are included in the data.

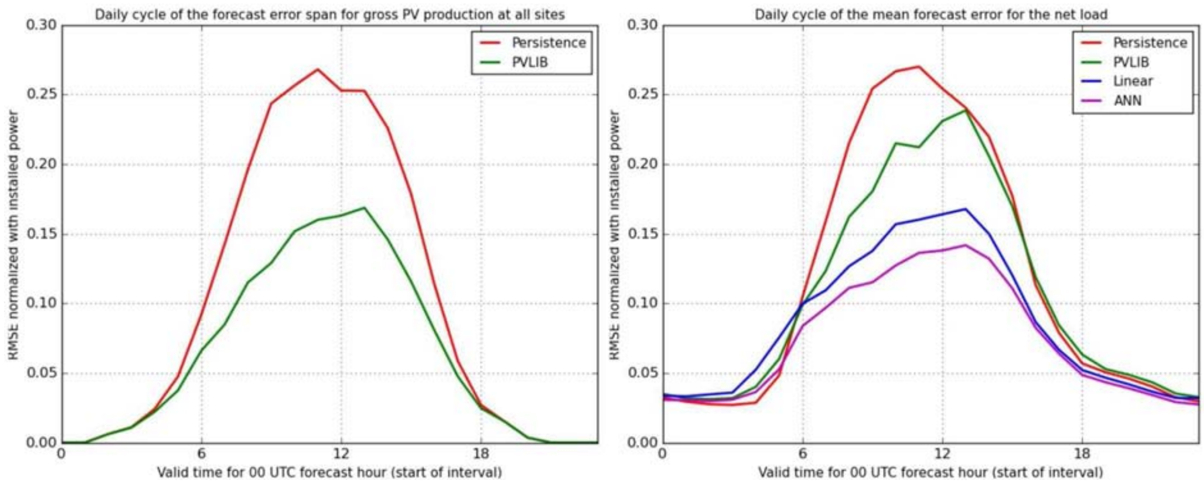

Figure 3: Left: Daily mean RMSEn cycle for the gross PV predictions (persistence: red, PVLIB: green). Right: Daily mean RMSEn cycle for net load predictions (persistence: red, PVLIB: green, linear model: blue, ANN: magenta).

The daily cycles of the mean RMSEn, when using PVLIB and the statistical models to forecast the net load, are illustrated in the right panel of Fig. 3. Again, the error is shown as a function of forecast length and all forecasts are initialized at 00 UTC. The ANN performs best, followed by the linear model, PVLIB and persistence. For the linear and ANN models the error grows somewhat towards the afternoon. This can be explained by the NWP forecast deteriorating with the length of the forecast. In the afternoon the performance of the PVLIB models becomes very similar to that of the persistence model. The reason for this is likely due to its dependence on the persistence forecast for the gross consumption. The error in the latter is high in the late afternoon when there is a peak in the consumption.

5 DISCUSSION AND CONCLUSION

Solar radiation is fluctuating in time and space and is non-trivial to predict. When a day-ahead forecasts is needed NWP models offer the best information [3]. However, in order to turn the solar radiation into a PV production forecast information about the geography and geometry at the installation is needed.

In this paper we studied how one physical and three statistical models performed on the task to predict the net load at a single site for the coming day. The input consisted of information about measured net production and net consumption from the previous day along with a NWP forecast for the next 24 hours. In a real situation the forecast has to be available well before the electricity market closes at about midday and hence the forecast horizon needs to be stretched to 42 hours. Such considerations will be the subject of further studies along with up-scaling of the forecast to an area of a regional electricity network.

The results show that it is possible to estimate the parameters in a physical model (PVLIB) based only on the daily cycle of the mean hourly net production and gross consumption. Information about the latter can be assumed to be available from load measurements during a corresponding time period prior to the PV installation. However, the pattern of consumption may change after the installation has come into place [9], e.g. by the purchase of an electric car or by considering the electricity produced BTM to be available free of charge. On the other hand, our modelling approach using PVLIB opens up for the possibility to include a model for the daily cycle of the average hourly consumption pattern, e.g. by doing a principal component analysis of a data base of consumption patterns and describe the pattern with a few eigenvectors in the optimization along with the PVLIB model parameters. In this way both the BTM production and consumption could be included in the optimization.

The physical model based on PVLIB shows good results when checked against the gross PV production. This means that both the NWP and the PVLIB models perform well. However, here the puropose was to predict the net load and for this purpose the physical model was outperformed by the black box models (linear and ANN). A likely reason for this is that the physical approach calls for an additional model for the hourly gross consumption. The gross consumption was predicted using persistence but this assumption turned out to be too weak.

Using an off the shelf ANN from TensorFlow reduced the overall RMSE, normalized with installed power, with only one percentage unit compared to the linear model. However, looking at the error with respect to forecast length reveals that the ANN actually provides a better prediction. Earlier studies have shown that the choice (e.g. the ANN) between different non-linear models is not critical [7].

Even better performance could perhaps be achieved if a recurrent structure is tried, exploiting the correlation in time between the foretasted values. One could also think of using a hybrid approach where the gross production is modelled with PVLIB and the gross consumption is modelled with an ANN. What to include in the input vector is another question for further investigations. Here we picked information we thought was reasonable. No evaluation was made regarding how useful different input parameters were for the prediction. Future work should also look at using input from probabilistic NWP forecasts. This should be a way to describe and account for uncertainties in the solar radiation forecasts.

To conclude, we have shown that a physical model

Table II: Model performance in terms of mean r 2 values and RMSEn for net load using PVLIB, a linear model, an ANN and persistence at the different sites (mean value, H: household, A: apartment block, O: office).

		Mean	H1	H2	H3	H4	H5	A1	A2	A3	O1
PVLIB	r^2	0.62	0.63	0.66	0.59	0.70	0.67	0.55	0.67	0.68	0.41
Linear	r^2	0.71	0.69	0.78	0.74	0.78	0.78	0.60	0.76	0.77	0.49
ANN	r^2	0.76	0.74	0.81	0.77	0.82	0.81	0.65	0.82	0.81	0.63
Persistence	r^2	0.40	0.44	0.44	0.42	0.47	0.45	0.31	0.45	0.46	0.18
PVLIB	RMSEn [%]	16	16	17	16	15	16	15	16	16	16
Linear	RMSEn [%]	12	13	13	12	11	12	12	12	12	12
ANN	RMSEn [%]	11	12	12	11	10	11	11	10	11	10
Persistence	RMSEn [%]	20	20	23	19	19	20	18	20	20	18

can be estimated and used to predict the gross PV production in the absence of BTM data. However, it was outperformed by statistical black box models such as linear regression and ANN when it comes to forecasting the net load.

ACKNOWLEDGEMENT

The authors greatly acknowledge the Swedish Energy Agency that funded this work via grant 43231-1 as part of the national research program “El från solen”.

REFERENCES

[1] Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mane, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viegas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

[2] Lisa Bengtsson, Ulf Andrae, Trygve Aspelien, Yurii Batrak, Javier Calvo, Wim de Rooy, Emily Gleeson, Bent Hansen-Sass, Mariken Homleid, Mariano Hortal, Karl-Ivar Ivarsson, Geert Lenderink, Sami Niemelä, Kristian Pagh Nielsen, Jeanette Onvlee, Laura Rontu, Patrick Samuelsson, Daniel Santos Mu noz, Alvaro Subias, Sander Tijm, Velle Toll, Xiaohua Yang, and Morten Odegaard Koltzow. The HARMONIE-AROME Model Configuration in the ALADIN-HIRLAM NWP System. Monthly Weather Review, 145(5):1919–1935, 2017.

[3] D. K. Chaturvedi and Isha. Solar power forecasting: A review. International Journal of Computer Applications, 145(6):28–50, Jul 2016.

[4] Lorenz Elke, Scheidsteger Thomas, Hurka Johannes, Heinemann Detlev, and Kurz Christian. Regional PV power prediction for improved grid integration. Progress in Photovoltaics: Research and Applications, 19(7):757–771, 2011.

[5] W. F. Holmgren, R. W. Andrews, A. T. Lorenzo, and J. S. Stein. Pvlib python 2015. In 2015 IEEE 42nd Photovoltaic Specialist Conference (PVSC), pages 1–5, June 2015.

[6] Soteris A. Kalogirou. Artificial neural networks in renewable energy systems applications: a review. Renewable and Sustainable Energy Reviews, 5(4):373 – 401, 2001.

[7] Luis Martn, Luis F. Zarzalejo, Jess Polo, Ana Navarro, Ruth Marchante, and Marco Cony. Prediction of global solar irradiance based on time series analysis: Application to solar thermal power plants energy production planning. Solar Energy, 84(10):1772 – 1781, 2010.

[8] E. J. Mlawer, S. J. Taubman, P. D. Brown, M. J. Iacono, and S. A. Clough. Radiative transfer for inhomogeneous atmospheres: RRTM, a validated correlated-k model for the longwave. Journal of Geophysical Research: Atmospheres, 102(D14):16663–16682, 1997.

[9] F. A. Monforte, C. Fordham, J. Blanco, S. Barsun, A. Kankiewicz, and B. Norris. Improving Short–Term Load Forecasts by Incorporating Solar PV Generation. Technical Report CEC-500-2017-031, California Energy Commission, 2017.

[10] J-J. Morcrette, H. W. Barker, J. N. S. Cole, M. J. Iacono, and R. Pincus. Impact of a New Radiation Package, McRad, in the ECMWF Integrated Forecasting System. Monthly Weather Review, 136(12):4773–4798, 2008.