Minimization of CO2 footprint by intelligent control of flexible electricity consumption

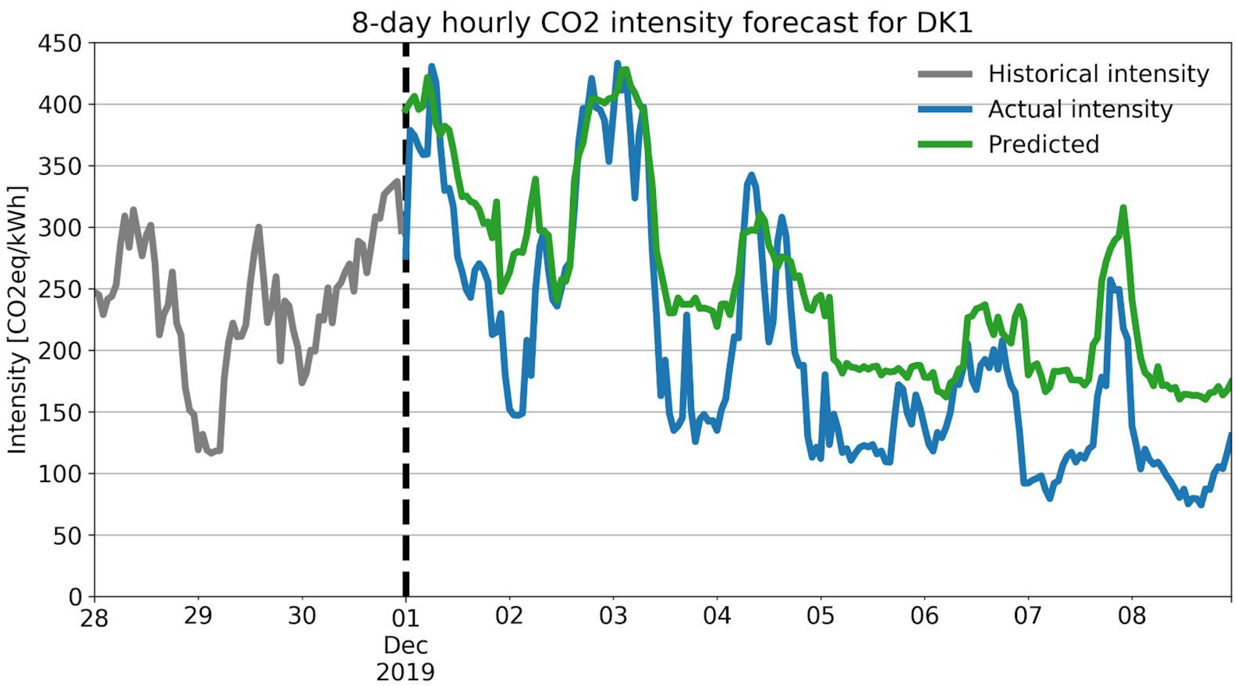

Final Report
ELFORSK Project 351-054
March 2020

Ento Labs

Resume

Planlægning af fleksibelt elforbrug er et vigtigt redskab for at mindske CO2-udledningen i elforbruget og integrere stigende mængder vedvarende energi i elsystemet. Resultatet af dette ELFORSK-projekt giver en bedre forståelse for de potentielle besparelser af kulstofemissioner og økonomiske gevinster forbundet med planlægning af fleksibelt elforbrug.

I dette projekt har vi udviklet to avancerede prognosemodeller for CO2-intensiteten i elforbruget, som styringsredskab til det fleksible elforbrug. I projektet optimerer vi driften af vores egne servere, ved at flytte server 'jobs' til tidspunkter hvor strømmen er grønnest. I vores analyse viser vi at besparelsespotentialet afhænger af forbrugets fleksibilitet og varighed. For et ugentligt job, med en to dages fleksibilitet og en varighed på seks timer, ses en 23 % CO2-besparelse, sammenlignet med at forbruge strømmen på et tilfældigt tidspunkt. Vores analyse viser desuden, at der med fordel både kan optimeres efter CO2-udledningen og spotprisen for strøm. Det er muligt at vægte de to parametre efter den ønskede effekt.

Baseret på vores 8-dages prognose for kulstofintensitet, har vi udviklet en web-API til planlægning af fleksibelt elforbrug. Denne API giver brugeren mulighed for at drage fordel af vores prognose og planlægge enhver form for fleksibel elforbrug - resultatet er altså ikke blot til optimering af serverdrift. API'en er fuldt funktionsdygtig og bruges i øjeblikket internt i Ento Labs. API'et er også tilgængeligt for interesserede parter der ønsker øget integration af vedvarende energi, eller optimal planlægning af forbrug ud fra pris eller CO2-udledning.

Timebestemte eldistributionsafgifter eller CO2-afgifter på strøm øger forretningspotentialet i intelligent planlægning af fleksibelt forbrug. Dette er især vigtigt for private forbrugere i Danmark, hvor skatter og afgifter udgør cirka to tredjedele af elprisen. I takt med øget elektrificering af transport- og varmesektoren, vil intelligent styring af denne type forbrug formentligt blive et modent forretningsområde.

Overgangen til 100 % vedvarende elektricitet i det sammenkoblede europæiske net kræver en kombination af transmissionsudvidelse, forskellige energilagringsteknologier og fleksibilitet. Dette demonstrationsprojekt har vist, hvordan man automatisk kan udnytte fleksibiliteten i elforbruget i Danmark, med minimale kapitalomkostninger for slutbrugeren, hvis ressourcen kan integreres via API. Ento Labs-platformen dækker allerede de europæiske elmarkeder.

Demonstrationsprojektet er lavet i samarbejde med Incuba A/S og Ejerforeningen Navitas A/S. Styrken består i at forstå slutbrugerens elforbrug og dermed potentialet i at aktivere fleksible ressourcer hos den enkelte forbruger. På sigt kan denne skalerbare metode også bruges i et systemperspektiv, når det samlede europæiske marked skal inkorporere en større andel vedvarende energi.

Summary

Scheduling flexible electricity consumption is an important tool for reducing CO_2 emissions of electricity consumption and integrating increasing amounts of renewable energy into the electricity system. The result of this ELFORSK project provides a better understanding of the potential savings of carbon emissions and economic benefits associated with scheduling of flexible electricity consumption.

In this project, we have developed two advanced forecasting models for CO_2 intensity of electricity consumption, as a tool to schedule flexible electricity consumption. In the project, we optimize the operation of our own servers, by moving server 'jobs' to times when the power is greenest. In our analysis we show that the savings potential depends on the flexibility and duration of the consumption. For a weekly job, with a two day flexibility and a duration of six hours, a 23 % CO_2 saving is seen, compared to consuming the electricity at a random time. Our analysis also shows that it is advantageous to optimize for both the CO_2 emission and the spot price of electricity. It is possible to weight the two parameters according to the desired effect.

Based on our 8-day forecast for carbon intensity, we have developed a web API for scheduling flexible electricity consumption. This API allows the user to take advantage of our forecast and plan any kind of flexible electricity consumption - the result is thus not just for optimizing server operation. The API is fully functional and is currently used internally in Ento Labs. The API is also available to interested parties seeking increased renewable energy integration, or optimal scheduling of consumption based on price or CO_2 emissions.

Hourly electricity tariffs or CO_2 taxes on electricity increase the business potential of intelligent scheduling of flexible consumption. This is especially important for private consumers in Denmark, where taxes and fees make up about two-thirds of the electricity price. In line with increased electrification of the transport and heating sector, intelligent control of this type of consumption will likely become a mature business area.

The transition to 100 % renewable electricity in the interconnected European grid requires a combination of transmission expansion, different energy storage technologies and flexibility. This demonstration project has shown how to utilize the flexibility of electricity consumption in Denmark automatically, with minimal capital costs for the end user if the resource can be integrated via API. The Ento Labs platform already covers the European electricity markets.

The demonstration project is a collaboration with Incuba A/S and Ejerforeningen Navitas A/S. The strength is to understand the end-user's electricity consumption and thus the potential of activating flexible resources with each consumer. In the long term, this scalable method can also be used from a systems perspective when the European market has to integrate a larger share of renewable energy.

Ento Labs

Contents

Introduction

Intelligent scheduling and control of flexible electricity consumption is a key answer to two growing concerns in the energy industry: 1) increasing intermittency in the *supply* of electricity due to increased penetration of variable renewable generation sources, and 2) increasing *demand* of electricity due to electrification of heating and transportation.

In broad terms, flexible electricity consumption aims to move electricity consumption to times where electricity supply is less constrained:

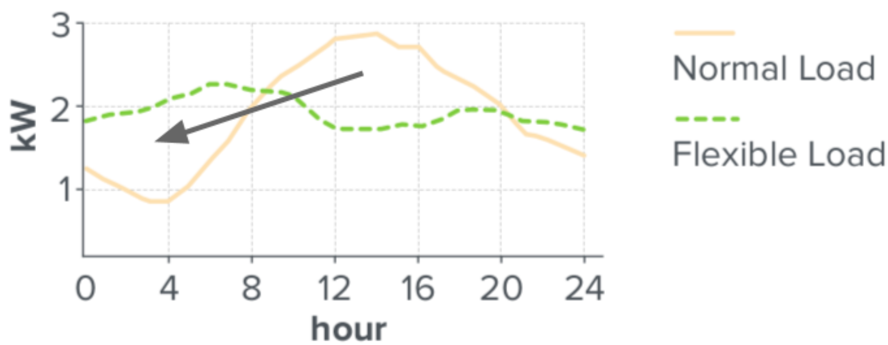

Figure 1: Demand Flexibility, Bronski et al. (2015)

In this demonstration project, we build a software tool capable of intelligent scheduling and control, enabling flexible consumption based on the CO2 emission levels of the local electricity grid in west Denmark (DK1). Building the tool on grid-level carbon emissions empower end users to lower their carbon emissions and thereby take part in the transition towards a low-carbon society.

Working with grid-level carbon emissions allows us to compare results of intelligent scheduling of different jobs. This allows us to explore the trade-off between how flexible the load is (how long a job can be postponed), the duration of the job and the carbon emissions of the job.

Grid level carbon emissions are based on “flow tracing” - a carbon accounting method, enabling calculations of the carbon intensity of electricity use on an hourly basis (Tranberg et al. 2019). Forecasting of carbon intensity for a single price area is based on weather forecasts for several representative geographical locations. Daily, weekly and monthly scheduling is based on a continuous 8-day carbon intensity forecast, developed as part of this project.

The flexible resources in this demonstration project are Ento Labs’ own servers. Intelligent scheduling of servers allows several benefits: 1) software-only integration, 2) varying levels of flexibility and load between different jobs, 3) development of capabilities in both short and long term forecasting and scheduling, and 4) a high degree of similarity to job scheduling at data centers and cloud computing/hosting companies. Carbon emissions are a growing concern for users of cloud computing in academia and the private industry (Henderson et al. 2020, Lacoste et al. 2019, Lottick et al. 2019, Schwarz et al. 2019).

This demonstration project extends Ento Labs existing energy data analytics platform. Here, electricity consumption data is analyzed in order to give residential and commercial consumers insights on how to lower their carbon emissions. The Ento Labs platform already provides users with individualized optimal solar PV dimensioning, identification of energy efficiency improvements and operational tips to improve carbon efficiency in buildings.

The combination of consumption, production and scheduling of electricity consumption in one combined platform enables end users and eventually utilities, to integrate a higher percentage of variable renewable energy in the electricity system and help accelerate the sustainable energy transition.

Renewable energy and the electricity grid

With ambitions to cut down on climate pollutants, countries worldwide are increasingly implementing intermittent renewable energy resources such as wind and solar photovoltaic (PV). The power output potential of wind and solar resources is affected by their respective geographical location. As shown below (Klein et al. 2016), the load hours of wind and solar in Europe is nearly inversely proportional between northern and southern Europe:

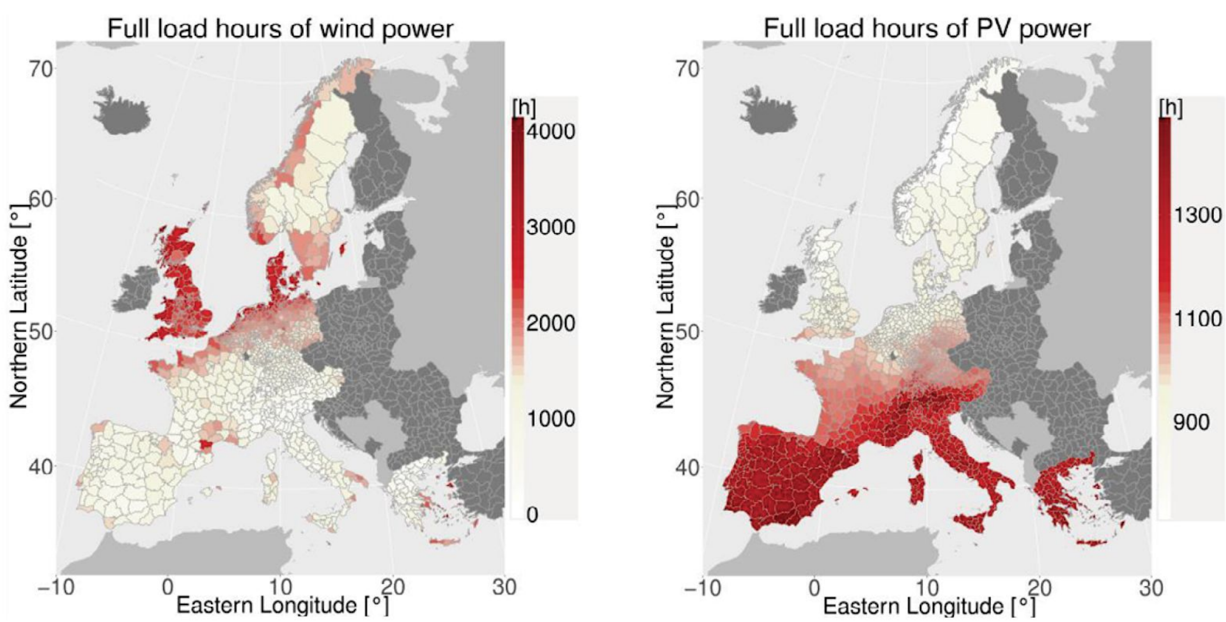

Figure 2: Klein et al (2016), Full load hours (production at rated capacity) for wind (left) and solar (right).

Given the large difference between load hours, a different mix of renewable energy supply is of preference between the European countries. The intermittency of wind and solar power resources makes it increasingly important to understand when the renewable electricity is being generated, as electricity demand must be met by the supply at all times. Renewable power generation will, depending on technology, vary widely throughout the day and the time of year as shown below, in Klein et al., (2016) scenario for 2030 for Denmark and Italy:

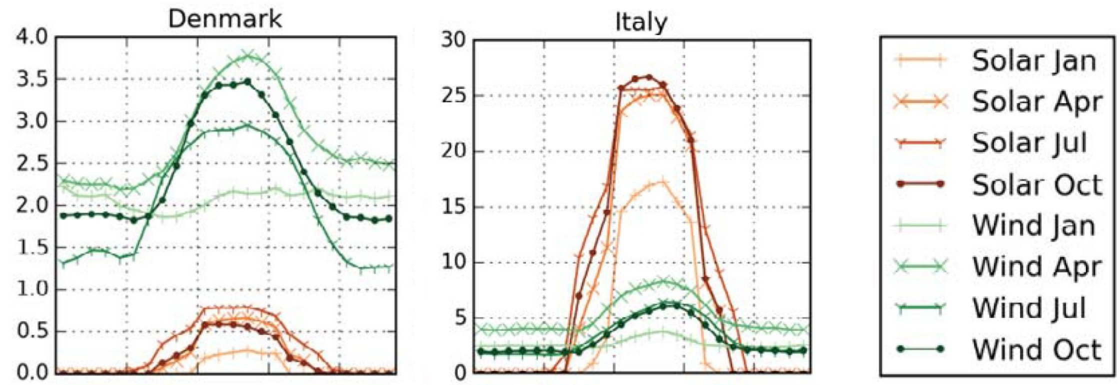

Figure 3: Klein et al (2016), Aggregated daily profiles of solar and wind generation 2030 for the months January, April, July and October (GW).

By analyzing the residual between the predicted demand and supply of renewable electricity, it becomes evident that countries with high penetration of renewable energy, such as Denmark, will have surplus electricity during mid-day in most seasons (except winter):

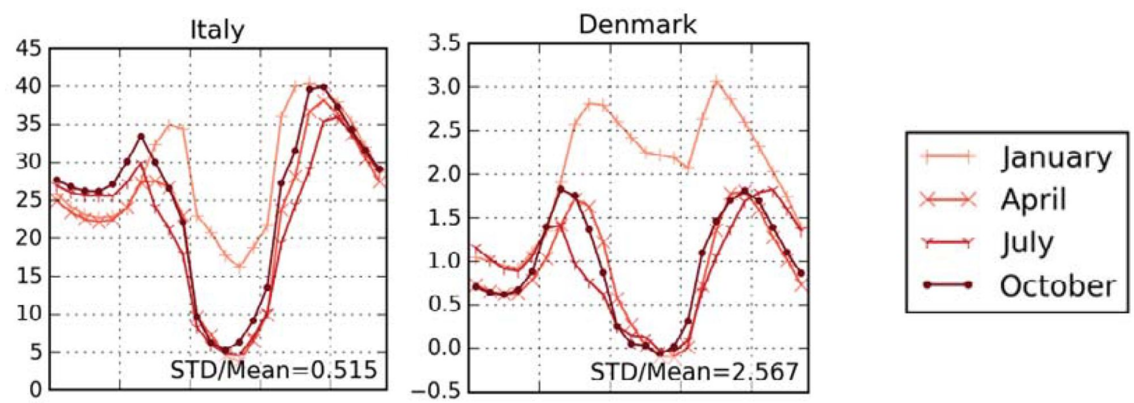

Figure 4: Aggregated daily profiles of the residual load 2030 for the months January, April, July and October. Assuming demand = 2011.

Worth noting is also the relationship between the standard deviation and the mean residual load, reflecting that countries powered by renewables have relatively wide fluctuating residual loads - proving the need for effective demand response measures (Klein et al. 2016). For Denmark, this means that the standard deviation is approximately 2.5 times larger than the mean residual consumption, throughout the year.

Mitigation strategies for handling the increasingly larger gap between supply and demand in peak hours, generally falls into two categories; 1) flexible electricity consumption, and 2) expansion of the high and low voltage grid. High-voltage grid expansions accommodates larger peak consumption demands by enabling utilization of production resources over a larger geographical area.

In this demonstration project, we have chosen to focus on building scalable technology solutions enabling flexibility in electricity consumption, due to the following reasons; 1) according to Danish Energy Association (2019) integration of renewable energy combined with further electrification of transportation can save the danish consumers around 16 billion DKK in infrastructure costs, if done ‘smart’ versus expansion of the grid. 2) As seen in Figure 2, countries within close geographical proximity generally have the same load hours, resulting in a similar renewable generation mix, resulting in concurrent generation and thus concurrent residual capacity.

Flexible electricity consumption

Flexible electricity consumption can be done in various ways, but the general idea is that an electrical load is moved to a time with less constraints. In most cases, this has been thought of as physical hardware, e.g. pre-heating a building to avoid heating at peak times of the day, charging electric vehicles at optimal times, or storing electricity in pumped hydro to use it later in the season when it’s needed more.

However, as more electricity is being consumed by data processing and data centers, flexibility in consumption of cloud or on-site computing is of increased importance.

This demonstration project focuses on this application as it entails a broad range of constraints and therefore paves the way of utilizing the developed technology, in not only this application, but also more traditional applications. In addition, a software-only application enables a high level of scalability and thus great potential social and grid-level impact.

We focus on carbon emission reduction, but also consider the trade-off between minimizing carbon emissions and the cost of electricity. It turns out that these objectives are not mutually exclusive, so in some cases one can achieve substantial carbon emissions savings and reduced cost of electricity.

For the scheduling of flexible electricity consumption to have an effect on the increasing integration of renewable energy resources in the electricity system, forecasts with sufficient horizons are required. The flexible consumer should, individually or at an aggregated level, be able to place bids in the day-ahead market to provide timely signals for the producers. Forecasts with longer horizons would enable consumers to participate in the futures markets for long-term planning and thereby reduce risks of consumers as well as producers.

Carbon intensity forecasting

Scheduling flexible electricity consumption requires a forecast of the expected carbon intensity. In this section, we present two different approaches to forecasting the carbon intensity of electricity: a short-term forecast based on decomposition methods, and an 8-day forecast based on weather. For both approaches, the carbon intensity of electricity is based on the method of Tranberg et al. (2019). The forecasts of carbon intensity should be interpreted as a carbon signal for market participants, that enables scheduling of flexible electricity consumption to minimize carbon emissions.

Short-term forecast based on decomposition methods

The forecast presented in this section has been developed as part of this ELFORSK project in collaboration with the Department of Engineering, Aarhus University. It is described in full detail in a forthcoming research paper by Bokde et al. (2020).

The purpose of this method is for a flexible electricity consumer who is willing to bid in specific hours on the day-ahead electricity market to minimize carbon emission. Consider the time framework in Figure 5 for an explanation of bidding in the day-ahead electricity market. Bids for consumption (buying electricity) in hourly intervals between 00 and 24 on the following day (D+1) must be submitted before noon today (D). Therefore, the short-term forecast of carbon intensity must be at an hourly resolution and have a horizon of at least 36 hours plus additional time for analysis and placing bids. The forecast therefore has a horizon of 48 hours.

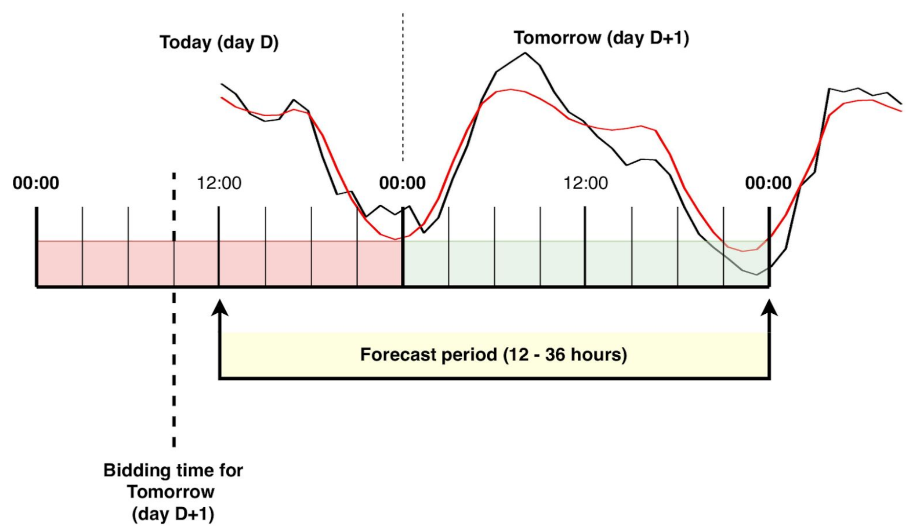

Figure 5: Time-frame for forecasting related to the day-ahead market.

The new method is based on decomposing the carbon intensity time series into three parts: seasonal, trend and random components. The seasonal component shows the patterns in the series that repeat with a fixed period e.g. daily and weekly, whereas the trend component shows the tendency of the series to increase or decrease over a long period of time, which is a smoother, general and long-term tendency. Finally, the random component, which is purely irregular in nature does not represent any pattern and is an uncontrollable and unpredictable part of the series.

The three components are forecasted individually and aggregated for the final forecast. For model selection, four models are considered: AutoRegressive Integrated Moving Average (ARIMA), feed forward neural network (FFNN), Pattern Sequence based Forecasting (PSF) and Difference Pattern Sequence based Forecasting (DPSF). It is found that the best model for predicting the seasonal component is FFNN and the best model for both the trend and random component is ARIMA.

The robustness of the new model is tested on 12 standard data sets with varying degrees of trend and seasonality against standard versions of ARIMA, FFNN, PSF and DPSF without decomposition. This is done to verify the quality of the new method and guarantee that it is not just a good fit for a specific artefact in the carbon intensity time series. It is found that the new method outperforms the traditional approaches in 7 out of 12 cases.

The figure below shows a 48-hour forecast of the carbon intensity of electricity in Denmark (red) together with the realized values (black). In this example, the consumer is looking to plan four hours of electricity consumption. According to the forecast this is scheduled in the intervals 16-17 and 18-21 on the following day. The average intensity for these four hours is 109 gCO2/kWh. Comparing this to consuming at 4 random hours throughout the day this intelligent scheduling results in a reduction of CO2 emissions of 19 %.

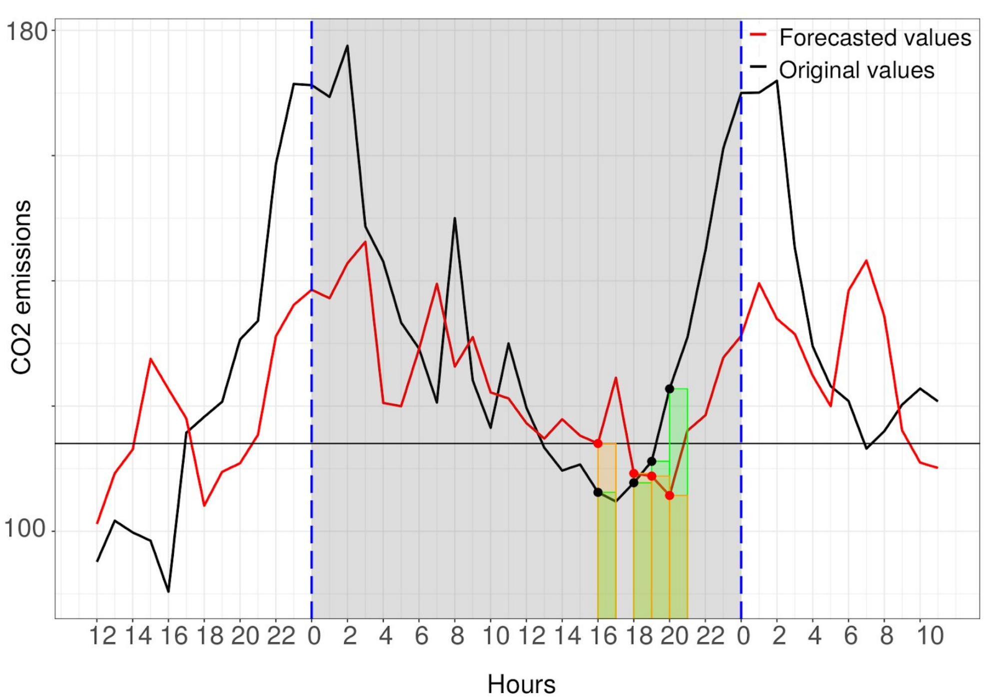

Figure 6: Comparison of forecasted (red) and actual (black) carbon intensity for Denmark.

8-day forecast based on weather

The forecast presented in this section has been developed solely by Ento Labs as part of this project.

The approach of this method is different from the one described above. The first method relied only on decompositions of the input time series. This means it does not take any external information into account. The non-parametric machine learning method presented in this section is able to learn its own representation of the inherent patterns of the carbon intensity time series, such as seasonality and trend from time variables, as well as utilize external weather data.

We use two weather data sets: one consisting of historical reanalysis weather data used to train our model and the other consisting of weather forecasts used for predicting the carbon intensity of electricity.

The model is trained on historical reanalysis weather data from ERA5[1] covering two full years from 2018 to 2019. Initially, we selected a large number of potential weather locations throughout Europe and measured their individual impact on the model accuracy through feature importance analysis. We found that despite the interconnectedness of the European electricity network, the weather conditions from locations far from the price area of interest are of little importance. In the end, we limited the input to 10 weather locations to train the model for the DK-1 and DK-2 price zones. These locations are illustrated in Figure 6. Offshore locations are close to existing wind farms. The six weather variables used for each location are: snow depth, air pressure, cloud cover, humidity, wind bearing, and wind speed.

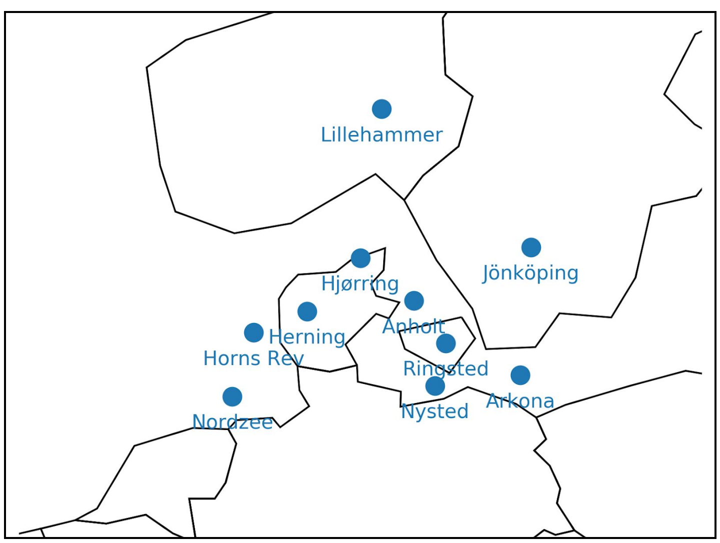

[1]ERA5 hourly data on single levels from 1979 to present:
https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-single-levels?tab=overview

Figure 7: Selected weather locations for training and forecasting the carbon intensity of electricity in Denmark.

For predicting the carbon intensity with a horizon of 8 days, as weather input we use forecasts from the Global Forecast System[2] provided by the National Oceanic and Atmospheric Administration. The forecasts are updated four times per day with 3-hourly resolution and a 192 hours (8 days) horizon. We interpolate the forecasts to hourly resolution before predicting the carbon intensity.

A note on temperature as input: we observed poor quality of forecasted temperatures for offshore locations. Based on this, we chose to exclude the use of temperature as input for all locations. It could have been included for onshore sites, but had very little effect on the forecast quality.

As an example, we show a historical forecast for the first week of december 2019 in Figure 8. The dashed line shows the time of forecasting. The grey line shows historical values for the carbon intensity and the blue line shows the realised carbon intensity in the 8 days following the time of the forecast. The green line shows the predicted carbon intensity.

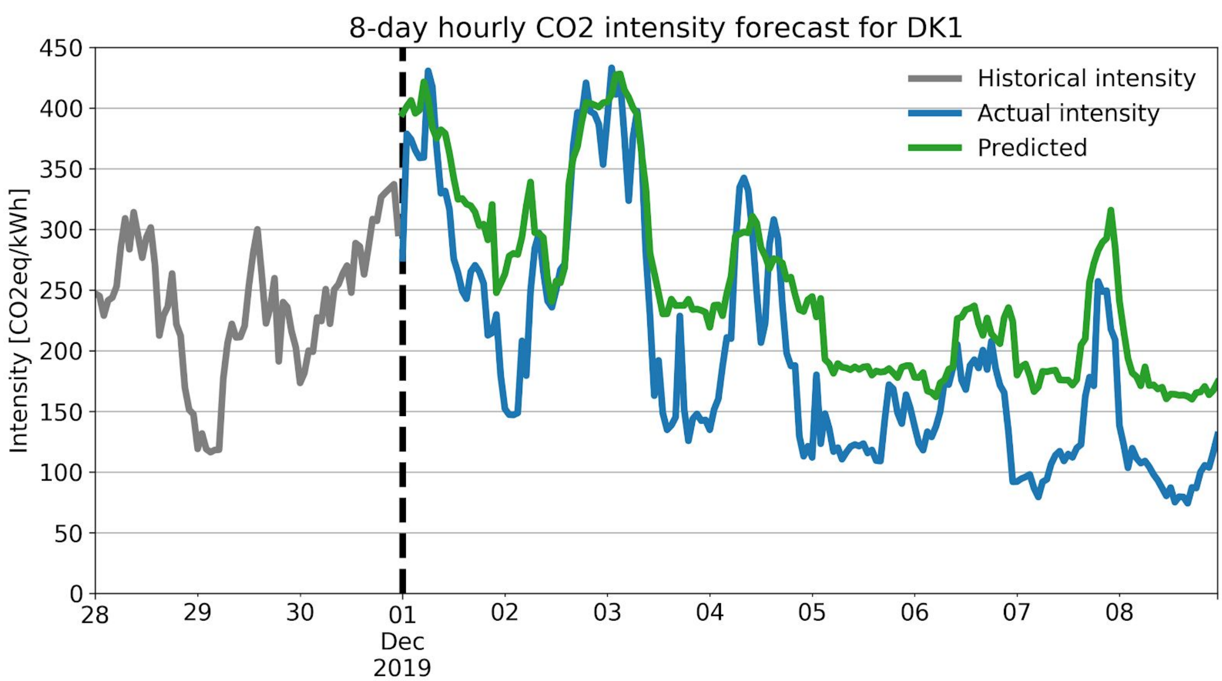

Figure 8: Historical 8-day forecast (green) for the first week of december 2019. Actual values are shown in grey and blue.

Another set of examples shows the overlap of forecasts from three consecutive days overlaid on the realized intensity. These are shown for three days in January 2020 in the figure below. Each of the three forecasts are based on GFS weather forecasts for the respective day.

[2]Global Forecast System (GFS):
https://www.ncdc.noaa.gov/data-access/model-data/model-datasets/global-forcast-system-gfs

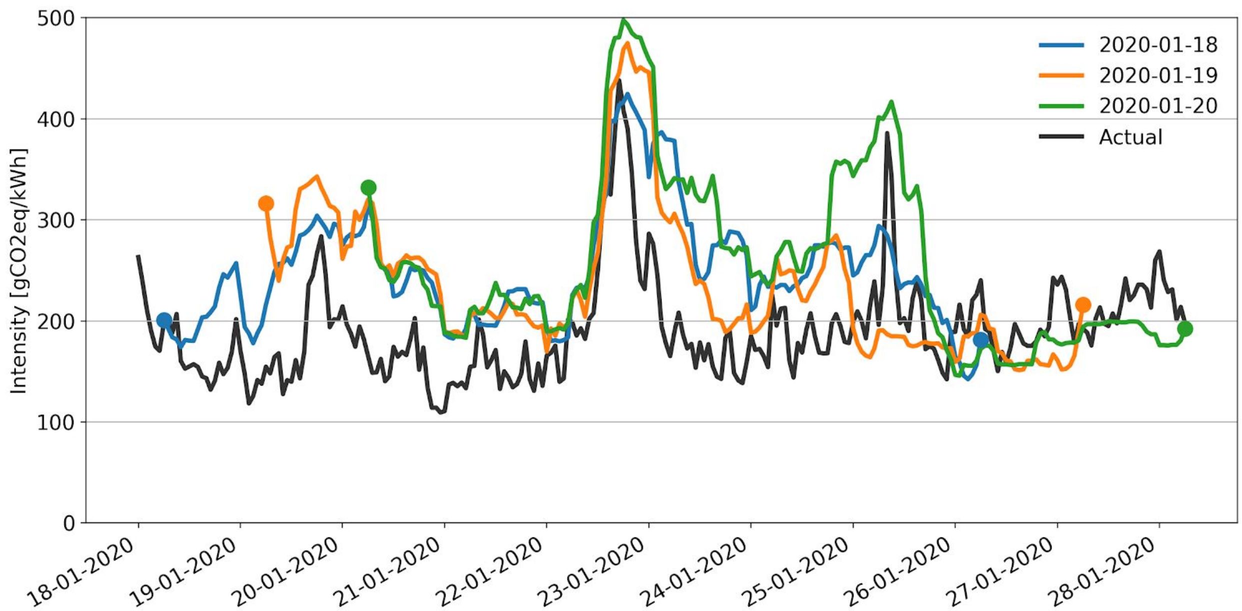

Figure 9: Overlaps of three consecutive 8-day forecasts of carbon intensity. Actual values are shown in black.

In both examples above it is clear that this forecast is not perfectly accurate. This is due to the long horizon of 8 days and that the input weather forecasts are not perfect. Note also that there is not a perfect mapping between the weather and the carbon intensity in the electricity system. When it comes to scheduling of flexible demand, however, the most important aspect of this type of forecast is not the hourly accuracy, but rather the ability of predicting the time of local minima and maxima in the intensity time series. For the flexible consumer it is paramount to know the time of local minimum to move their flexible demand. The actual hourly value of the carbon intensity is of lesser importance. With this perspective in mind our forecast seems to be performing well in both examples, especially when taking into account the hourly resolution and horizon of 192 hours.

A forecast with a horizon of 8 days has a much larger potential impact on scheduling flexible electricity consumption compared to the forecast with a horizon of only 24 or 48 hours. Combined with the convenient availability of 8-day weather forecasts, we choose to use the 8-day forecast for the study of carbon emissions savings in the following sections.

Carbon emissions savings

In this section we first investigate the theoretical potential for carbon emissions savings based on a perfect carbon intensity forecast. Secondly, we show results of actual emissions savings obtained by our 8-day carbon intensity forecast. In all cases we further investigate the trade-off between minimizing carbon emissions and the cost of electricity.

Theoretical potential of flexible demand

Daily job

The first example is a daily computing job with a duration of 3 hours - this could be retraining the weights of an ML model based on incoming data every day. We compare emissions between running the job fixed at 02 in the morning and optimizing to find the best 3 hours between midnight and 6 in the morning.

Figure 10 shows simulation results for the daily job through 2018 and 2019. Emissions from the fixed run at 02 are shown in blue and the optimized in green. The dashed lines show the average for each of the two approaches. For these specifications, on average, we find a 9.6 % reduction in emissions when selecting the best three hours between 0-6 every night instead of always running the job at 02.

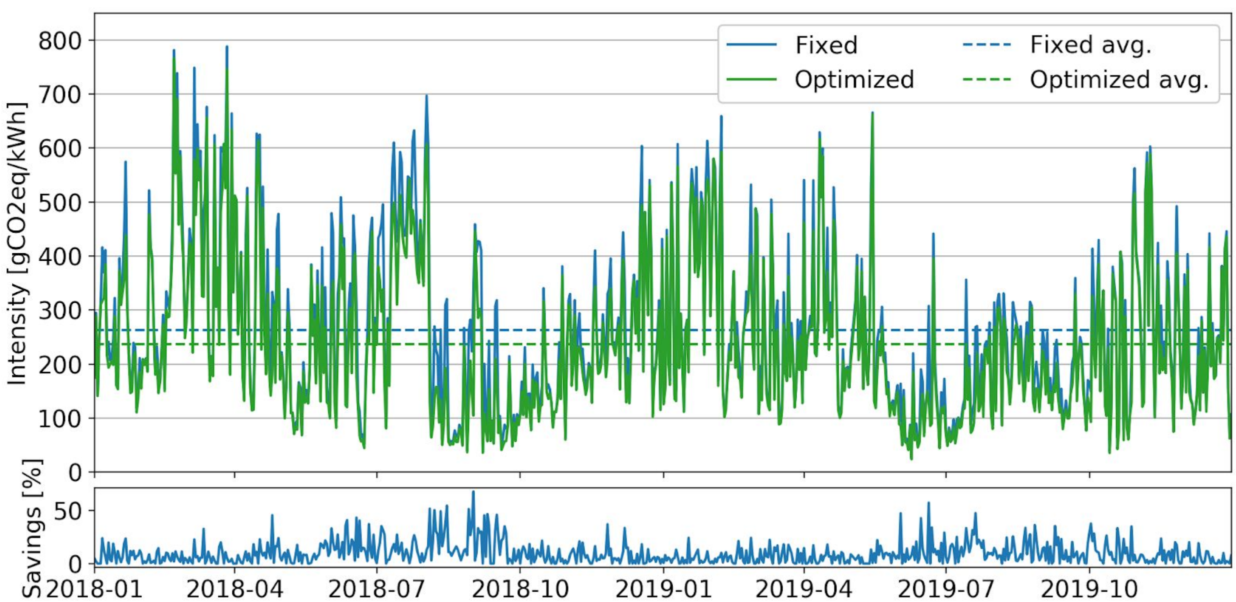

Figure 10: Comparison of carbon intensity for fixed (blue) and scheduled (green) consumption at daily frequency and resulting savings (bottom).

Monthly job

The second example is a monthly computing job with a duration of 6 hours - this could be a larger job related to invoicing or performing model selection based on incoming data for the entire month. We compare emissions between running the job fixed at midnight at the end of the month and optimizing to find the best 6 hours plus/minus 3.5 days from midnight at the end of the month.

Similar to the daily job, Figure 11 shows simulation results for the monthly job through 2018 and 2019. Emissions from the fixed run at midnight are shown in blue and the optimized in green. The dashed lines show the average for each of the two approaches. For these specifications, on average, we find a 48 % reduction in emissions when selecting the best six hours in a 7 day window centered around the end of the month instead of always running the job at midnight at the end of the month.

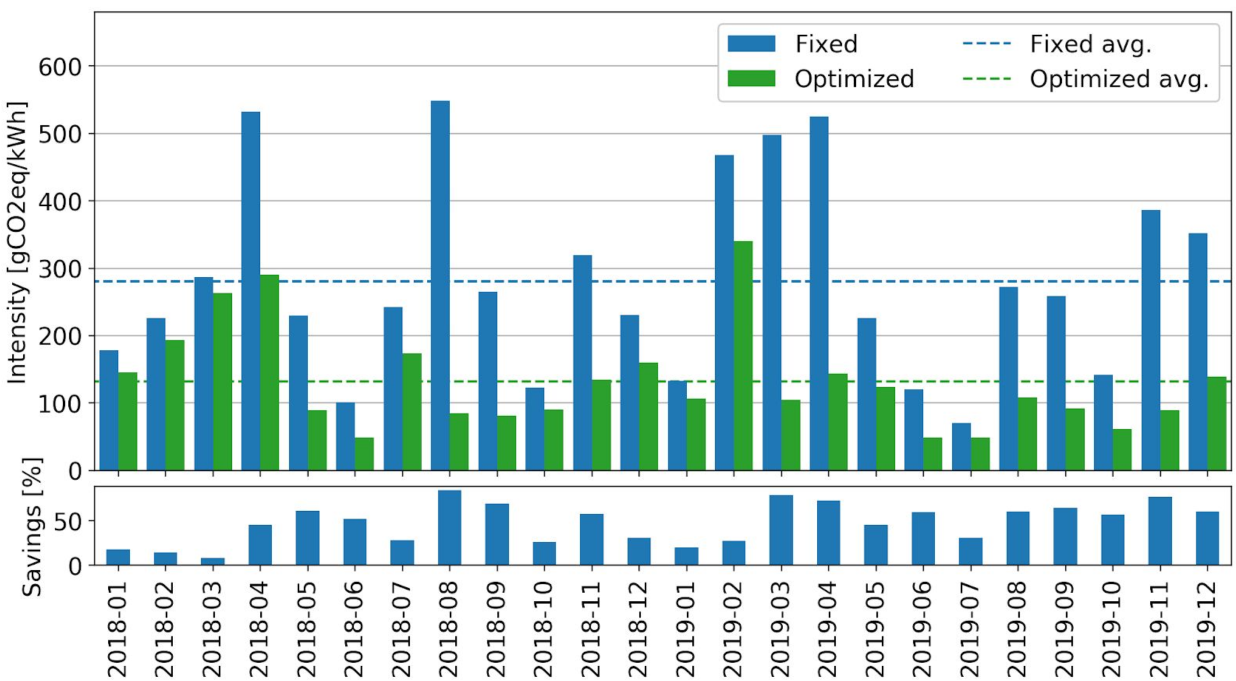

Figure 11: Comparison of carbon intensity for fixed (blue) and scheduled (green) consumption at monthly frequency and resulting savings (bottom).

Trade-off between emissions and price

The optimization objective of the previous two examples was solely the carbon intensity of electricity. In this section, we investigate the trade-off between optimizing for lowest emissions and optimizing for the lowest cost. Tranberg et al. (2020) showed a strong, but also time-varying, negative correlation between wind power production and electricity price in the West Denmark price zone. Based on this insight there should be potential for a joint optimization

objective for the CO_2 intensity and the electricity price to reduce both the emissions and the cost of a flexible computing job.
The figure below shows the average reduction in CO_2 emissions (blue) and the cost (green) for the monthly job described above for all of 2018 and 2019. The first axis shows how much weight is assigned to the price in the scheduling optimization: 0 % price weight means optimizing 100 % for the carbon intensity and 100 % price weight means 0 % carbon intensity. The first blue dot to the left corresponds to the difference between the two dashed lines in the previous figure: an average of 48 % reduction of the carbon emissions when the optimization objective is exclusively the carbon intensity. In this case the cost savings are just 0.67 %. On the other hand, if the optimization objective is exclusively the price, the average cost savings reach 33 % whereas the emissions savings are just 0.94 %. It seems there is a favourable weight between CO_2 and price around 50 % since the price savings flatten at higher price weight. At this point, the emissions savings are 32 % and the cost savings are 28 %.

Interestingly, for a more price conscious consumer, a price weight of up to 80 % would result in cost savings of 31 % and still give 23 % emissions savings.

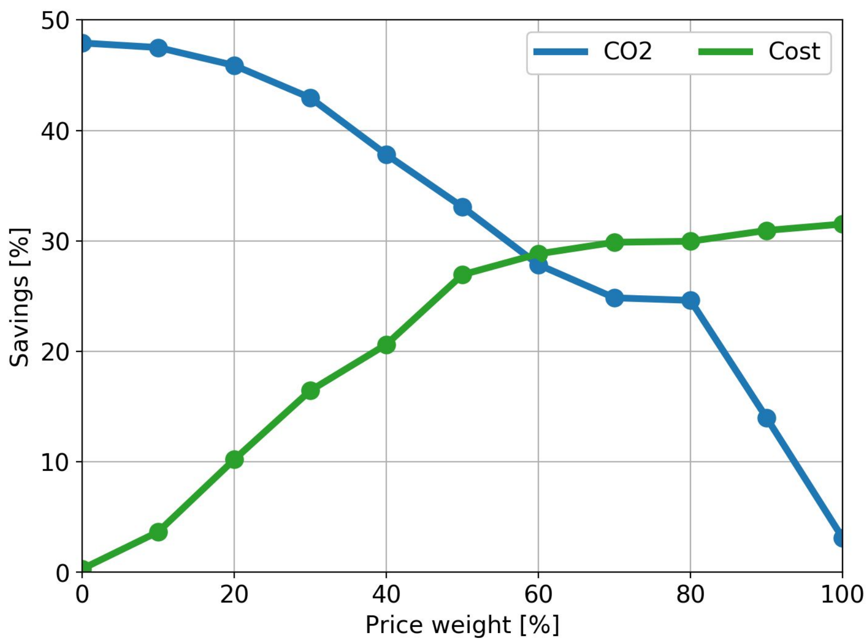

Figure 12:

Actual savings from flexible demand

The results in this section are based on the 8-day carbon intensity forecast described earlier. We have obtained daily historical weather forecasts for the most recent year, which allow us to produce historical carbon intensity forecasts and thereby calculate actual carbon emissions savings from flexible electricity consumption in this time frame.

Note that since this is an actual forecast, sometimes the obtained savings can be negative, which is not the case for the perfect forecast used to determine the theoretical potential in the previous section. However, despite a very small number of negative cases, on average we find substantial carbon emissions savings.

We show three examples with increasing levels of flexibility: a daily job with a window of 16 hours, a weekly job with a window of 48 hour and a monthly job with a window of 7 days. In all cases we investigate the trade-off between minimizing carbon emissions and the cost of electricity.

Daily

The first example is a daily computing job with a duration of 6 hours - this could be retraining the weights of an ML model based on incoming data every day. We compare emissions between running the job fixed at midnight and optimizing to find the best 6 hours between 16 and 8 in the morning.

The figure below shows simulation results for the daily job for the last year. Emissions from the fixed run at midnight are shown in blue and the optimized in green. The dashed lines show the average for each of the two approaches. For these specifications, on average, we find a 7 % reduction in emissions when selecting the best 6 hours between 16-8 every night instead of always running the job at midnight. Compared to running the job at a random time during the window, the emission reductions are 10 %.

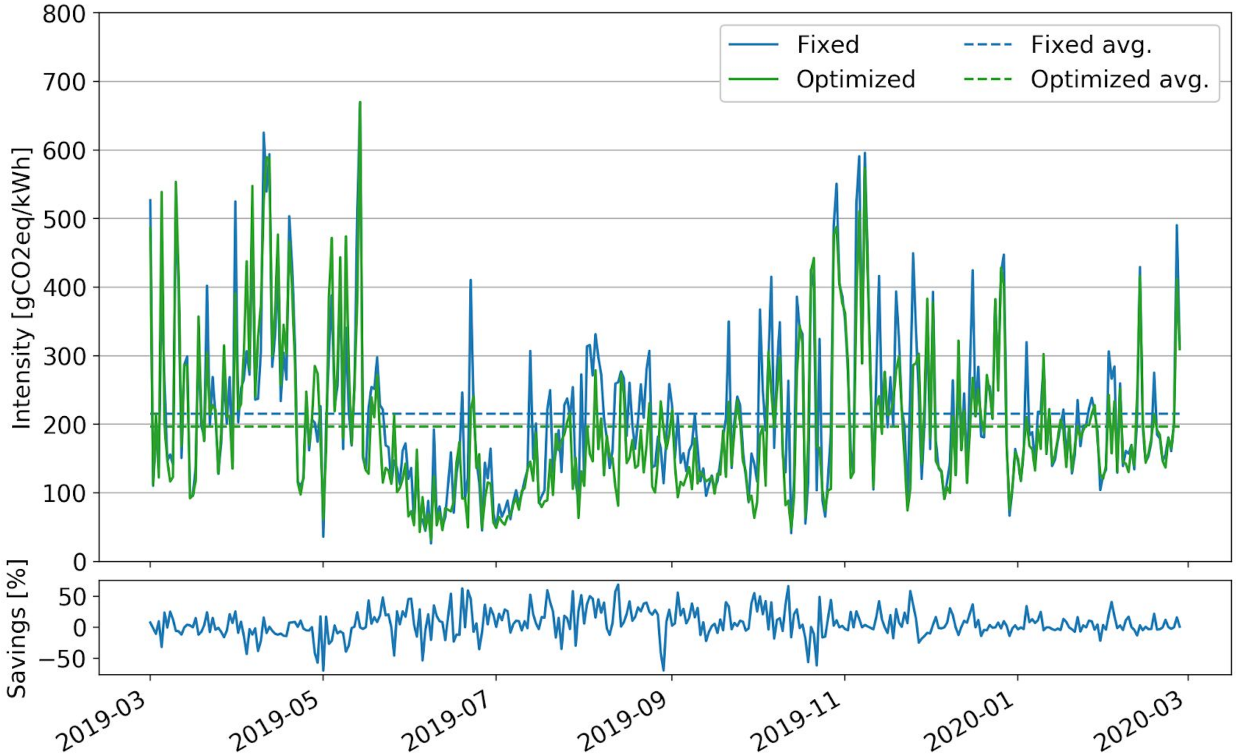

Figure 13: Comparison of carbon intensity for fixed (blue) and scheduled (green) consumption at daily frequency and resulting savings (bottom).

The optimization objective in the example above was solely the carbon intensity of electricity. We now investigate the trade-off between minimizing emissions and minimizing costs. The figure below shows the average reduction in CO2 emissions (blue) and the cost (green) for the daily job described above. The first axis shows how much weight is assigned to the price in the scheduling optimization: 0 % price weight means optimizing 100 % for the carbon intensity and 100 % means 0 % carbon intensity. The first blue dot to the left corresponds to the difference between the two dashed lines in the previous figure: an average of 7 % reduction of the carbon emissions when the optimization objective is exclusively the carbon intensity. In this case the cost savings are negative 19 %, which means the cost is increasing when optimizing for lower emission. On the other hand, if the optimization objective is exclusively the price, the average cost savings are 2.2 % whereas the emissions savings are just 0.34 %.

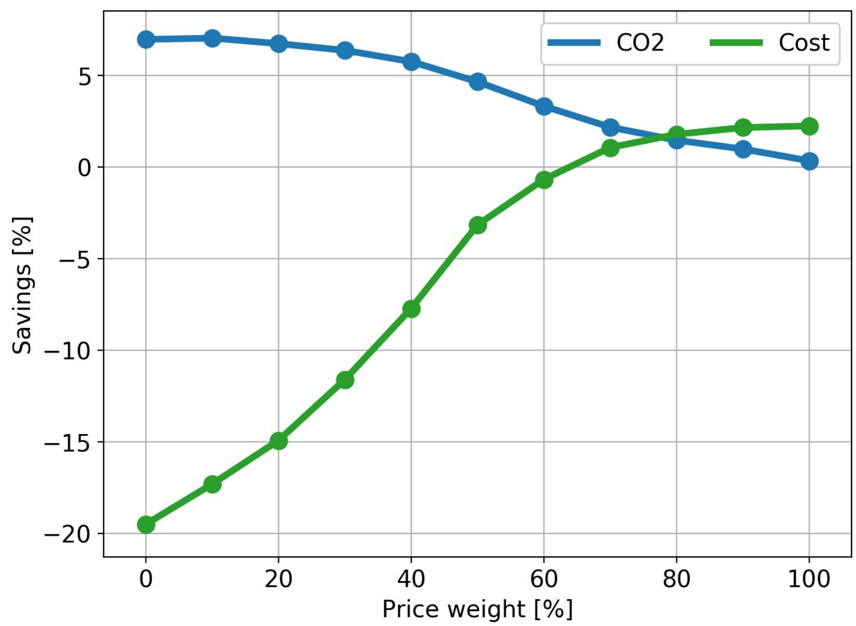

Figure 14: Average reduction in CO2 emissions (blue) and the cost (green) for a daily job. Negative reduction in costs means it becomes more expensive.

Weekly

The second example is a weekly computing job with a duration of 6 hours. Similar to the daily job, this could be retraining the weights of an ML model. We compare emissions between running the job fixed at midnight at the end of the week and optimizing to find the best 6 hours during the first two days of the following week.

The figure below shows simulation results for the weekly job for the last year. Emissions from the fixed run at midnight are shown in blue and the optimized in green. The dashed lines show the average for each of the two approaches. For these specifications, on average, we find a 11 % reduction in emissions when selecting the best 6 hours instead of always running the job at midnight. Compared to running the job at a random time during the window, the emission reductions are 23 %.

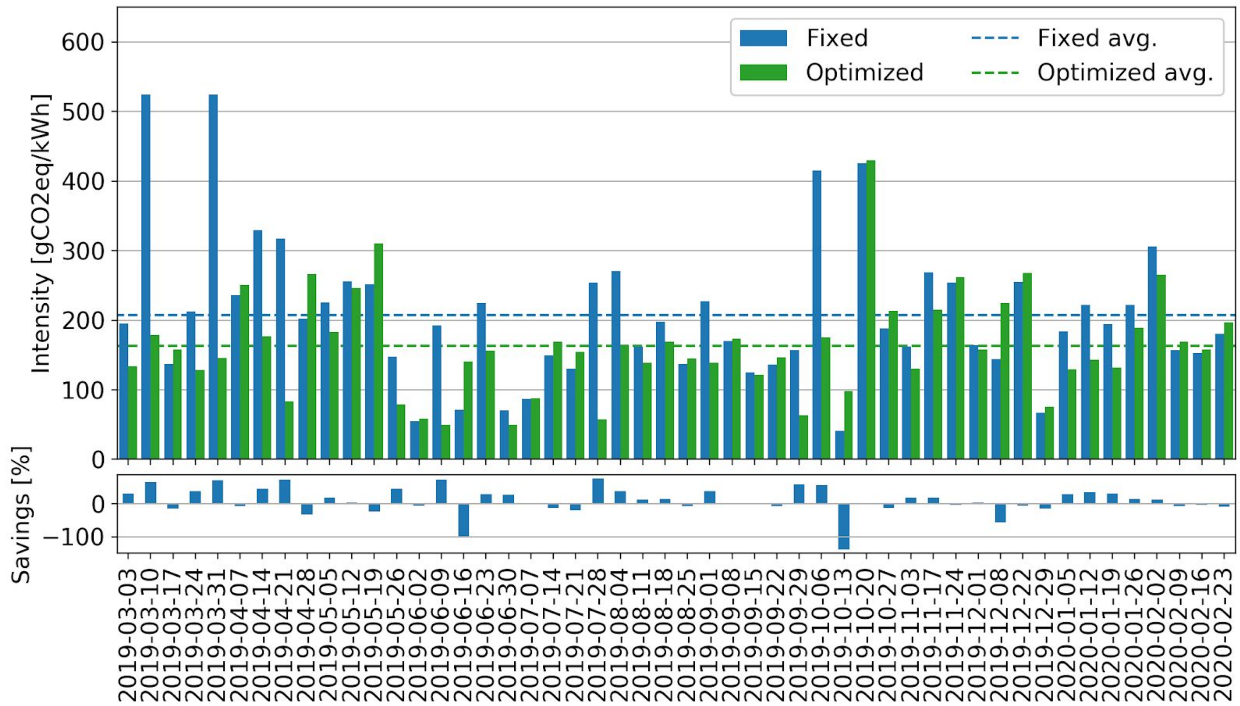

Figure 15: Comparison of carbon intensity for fixed (blue) and scheduled (green) consumption at weekly frequency and resulting savings (bottom).

As for the daily job above, we investigate the trade-off between minimizing emissions and minimizing costs. The figure below shows the average reduction in CO2 emissions (blue) and the cost (green) for the daily job described above. Again, the first blue dot to the left corresponds to the difference between the two dashed lines in the previous figure: an average of 11 % reduction of the carbon emissions when the optimization objective is exclusively the carbon intensity. In this case the cost savings are negative 25 %, which means the cost is increasing when optimizing for lower emission. On the other hand, if the optimization objective is exclusively the price, the average cost savings are 1.6 % whereas the emissions savings are 8.3 %.

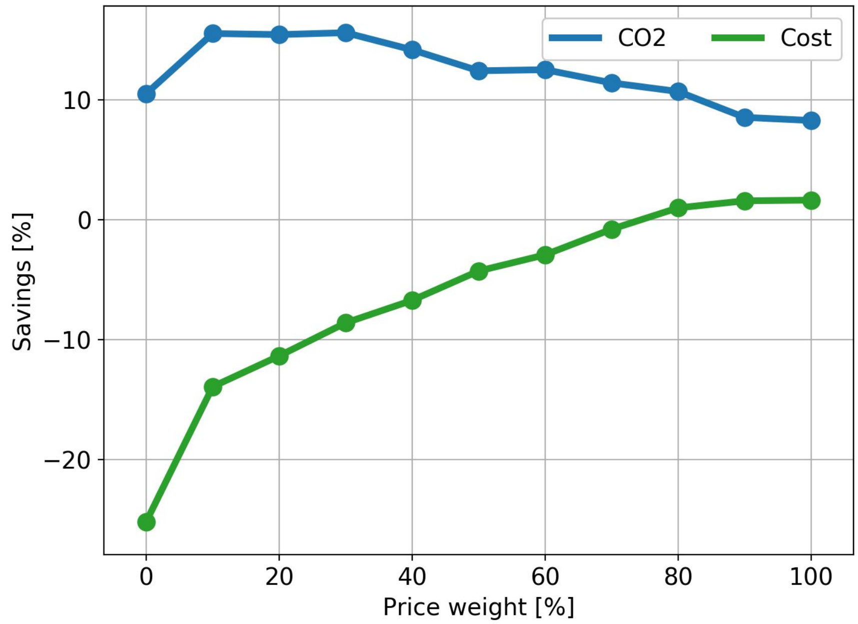

Figure 16: Average reduction in CO2 emissions (blue) and the cost (green) for a weekly job. Negative reduction in costs means it becomes more expensive.

Monthly

The third example is a monthly computing job with a duration of 6 hours. Similar to the monthly job in the previous section, this could be a job related to invoicing or performing model selection. We compare emissions between running the job fixed at midnight at the end of the month and optimizing to find the best 6 hours during the first week days of the following month.

The figure below shows simulation results for the weekly job for the last year. Emissions from the fixed run at midnight are shown in blue and the optimized in green. The dashed lines show the average for each of the two approaches. For these specifications, on average, we find a 34 % reduction in emissions when selecting the best 6 hours instead of always running the job at midnight. Compared to running the job at a random time during the window, the emission reductions are 31 %.

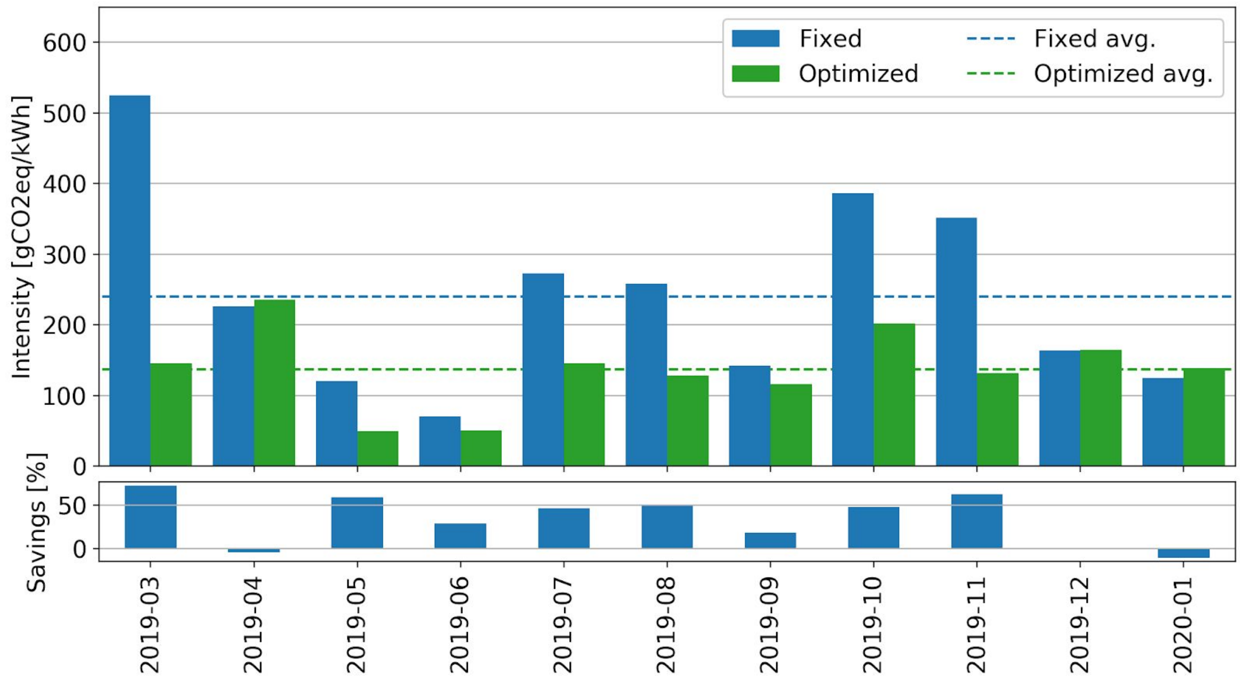

Figure 17: Comparison of carbon intensity for fixed (blue) and scheduled (green) consumption at monthly frequency and resulting savings (bottom).

As for the weekly and daily jobs above, we investigate the trade-off between minimizing emissions and minimizing costs. The figure below shows the average reduction in CO2 emissions (blue) and the cost (green) for the daily job described above. Again, the first blue dot to the left corresponds to the difference between the two dashed lines in the previous figure: an average of 34 % reduction of the carbon emissions when the optimization objective is exclusively the carbon intensity. In this case the cost savings are 0.77 %, which means the cost is increasing when optimizing for lower emission. On the other hand, if the optimization objective is exclusively the price, the average cost savings are 42 % whereas the emissions savings are 17 %.

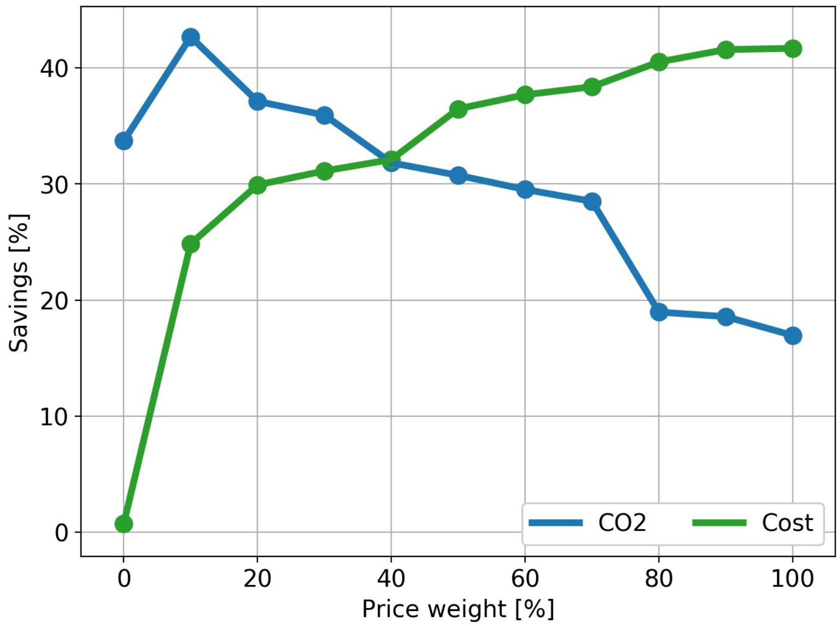

Figure 18: Average reduction in CO2 emissions (blue) and the cost (green) for a daily job.

A note on cost savings

We saw in the case studies above that minimizing the carbon emissions can result in increased costs (negative savings) for the daily and weekly jobs. To understand this, let's look at the daily profiles for the spot price and carbon intensity.

The figures below shows the daily profile for spot price (left) and carbon intensity (right). The horizontal green lines mark the median and the boxes show the interval from the 25th percentile to the 75th percentile. The whiskers show the 10th and 90th percentile. There are occurrences of negative prices, but this happens only about 1.5 % of the time. Both series display a morning and evening peak that is more evident for the price than the intensity.

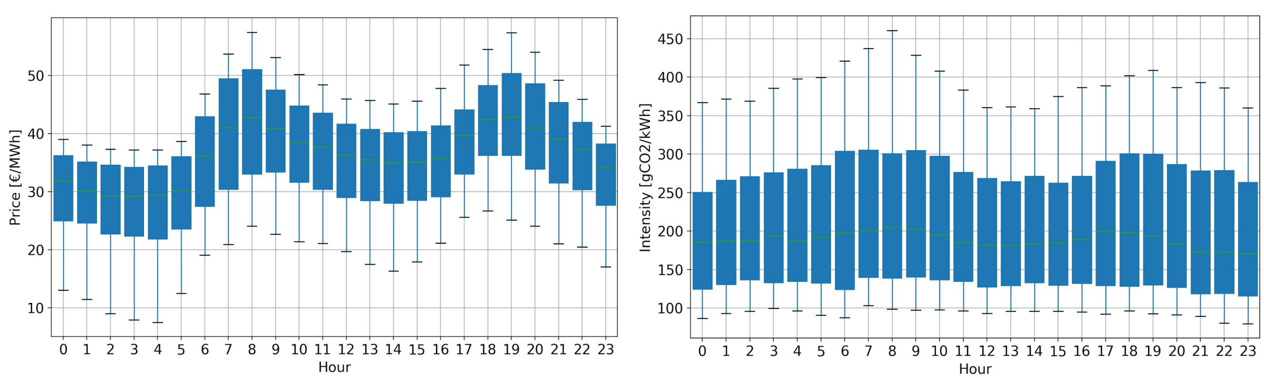

Figure 19: Daily profile for spot price (left) and carbon intensity (right). Green bar marks median. Boxes show 25th to 75th percentiles. Whiskers show 10th and 90th percentiles.

From these figures we see a clear daily profile for the spot price and to a lesser extent for the intensity. This means that the hourly variation in the intensity is more chaotic than the price.

The dependence between the two is shown in the scatter plot below. The correlation between the price and the intensity is 0.35.

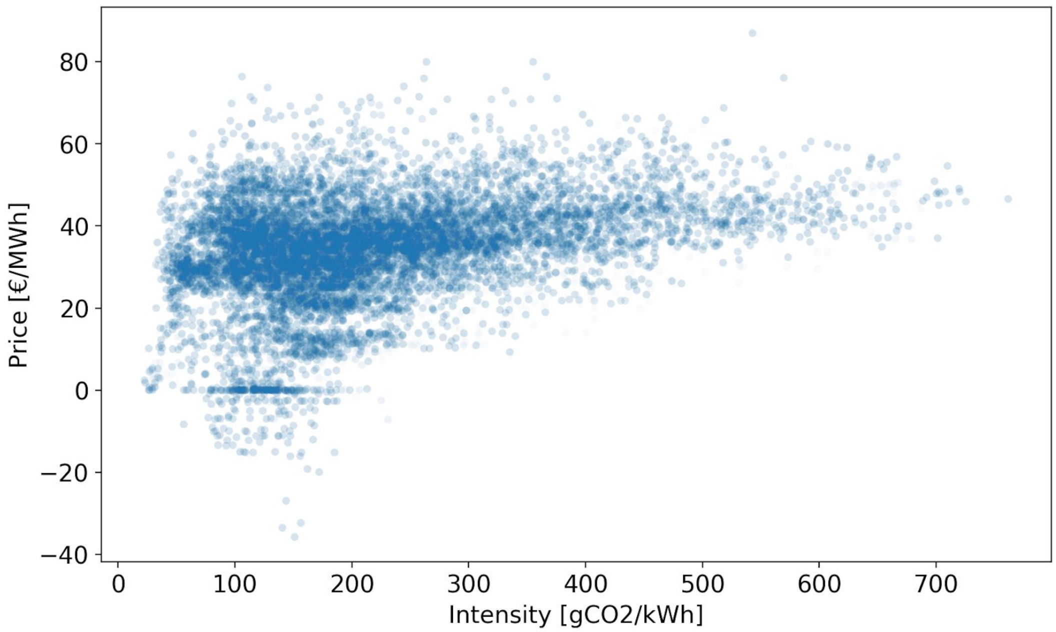

Figure 20: Scatter plot that shows the dependence between carbon intensity and the spot price.

The relatively low correlation between the spot price and the carbon intensity is what causes the patterns we see in the case studies, when varying the weight of our optimization objective: the maximum cost savings and emissions savings occur at different weights.

The monthly job was able to obtain cost savings of up to 42%. The daily and weekly jobs only got to 2.2 % and 1.6 %, respectively. These results are caused by selecting the fixed reference hour to be midnight. From the daily profile in Figure 19, we can see that midnight (hour zero) is generally among the cheapest hours in the day. Had we instead selected a random hour within the window of flexibility, the results would have looked more favourable. This would have a similar effect on the carbon emissions savings as the carbon intensity of the midnight hour is among the lowest in the day. Because of this, the results presented here should be seen as a conservative estimate. From this insight we confirm the general rule of thumb that it’s better to consume during the night - away from the daily peaks.

Scheduling API

Prior to this project, Ento Labs has developed a platform[3] that combines several sources of data for individual geographical locations (metering points) to produce individualized forecasts of electricity consumption, self-generation from solar panels, as well as the share of renewables in the grid.

Ento labs specializes in optimizing carbon efficiency related to energy usage. With increasing levels of renewable energy in the electricity grid, scheduling of flexible demand becomes an increasingly important part of increasing carbon efficiency. In this project, we have built an extension to the existing platform in the form of an API to determine the minimal CO2 emissions and optimize the scheduling of flexible electricity demand.

Based on the 8-day carbon intensity forecast, we have developed a web API for scheduling flexible demand be it computing jobs or any other intensive electricity consumption source. The scheduling API uses our forecasts to suggest the time that minimizes the carbon emissions of a flexible job within a given time window and expected duration of the job. The scheduling API has a horizon of at least 7 days depending on the update frequency of the underlying weather forecasts.

This API allows potentially anyone to benefit from our forecast and schedule flexible demand provided they have the necessary hardware installed to control the flexible resources e.g. charging of electric vehicles. For computing jobs, no additional hardware is required, which is why we have chosen this for the demonstration project.

We briefly present the use of our API in the sections below. For interest in using this API for business or research purposes, please get in touch with us at contact@ento.ai.

Request

The scheduling endpoint is accessible with an API token at `/schedule` (relative to the base API path, only available to pilot and beta customers at the moment.

Location and time information is automatically estimated from the request. There are four additional parameters:

1. `duration`: the expected duration of the job (default: 1h).
2. `window`: the time window within to find for the least CO2eq emitting timeslot (default: 24h).
3. `start`: adjust the start of the window (default: now).

[3] https://ento.ai

4. `end`: adjust the end of the window (default: none).

Response

The API returns a JSON formatted response with six keys:

1. `duration`: the duration used in seconds.
2. `window`: the window used in seconds.
3. `start`: the window start timestamp used.
4. `end`: the window end timestamp used.
5. `best`: the start time of the forecasted, least CO2eq emitting timeslot.
6. `savings_pct`: the forecasted CO2eq savings compared to now in percent.

Examples

Example 1: `duration` and `window`

```
$ curl -G -H "Authorization: Token XYZ" \
    --data-urlencode "duration=1h" \
    --data-urlencode "window=24h" \
    https://api.ento.ai/v1/schedule

{
  "duration": 3600,
  "window": 86400,
  "start": "2020-03-04T08:00:00Z",
  "end": "2020-03-05T08:00:00Z",
  "best": "2020-03-04T15:00:00Z",
  "savings_pct": 26
}
```

Example 2: `duration`, `start`, and `end`

```
$ curl -G -H "Authorization: Token XYZ" \
    --data-urlencode "duration=6h" \
    --data-urlencode "start=2020-03-05T00:00:00Z" \
    --data-urlencode "end=2020-03-12T00:00:00Z" \
    https://api.ento.ai/v1/schedule

{
  "duration": 21600,
  "window": 604800,
  "start": "2020-03-05T00:00:00Z",
  "end": "2020-03-12T00:00:00Z",
  "best": "2020-03-10T19:00:00Z",
  "savings_pct": 54
}
```

Conclusion

Optimal scheduling of flexible electricity consumption is an important part of reaching carbon efficiency. With increasing levels of variable renewable energy sources in the future, it becomes increasingly important to schedule flexible demand accordingly. The outcome of this ELFORSK project provides a better understanding of the potential savings of carbon emissions as well as costs associated with scheduling flexible electricity consumption. This in turn provides a better understanding of the potential business cases related to flexible demand.

During this project, we have developed and tested two forecasting models for the carbon intensity of electricity. Through several examples, we show the savings potential of carbon emissions and costs associated with scheduling computing jobs of varying frequency and level of flexibility. The results are firstly a confirmation of the general rule of thumb that it’s best to consume electricity during the night - away from the daily peaks in price and carbon emission intensity. Secondly, the savings potential increases when the relation between the duration and flexibility window of the flexible demand increases. The larger the window or the shorter the job, the more potential there is for reducing carbon emissions and cost.

Based on our 8-day carbon intensity forecast, we have developed a web API for scheduling flexible demand be it long-running computing jobs or any other intensive electricity consumption source. This API allows potentially anyone to benefit from our forecast and schedule any flexible demand - not limited to computing. The API is fully functional and currently used internally within Ento Labs. It’s also made available in an early version for interested parties for business and research purposes.

Understanding the electricity consumption of end-users, identifying what part of their consumption is flexible, and enabling autonomous systems to schedule the flexible consumption at scale becomes increasingly important, as our electricity grid becomes increasingly dependent on variable renewable energy sources.

During the demonstration project, Ento Labs has engaged with several market participants. From these engagements we have learned that there is interest from business to schedule flexible demand based on signals such as price and carbon intensity or even provide ancillary services to the electricity grid. The current regulation for ancillary services requires a minimum bid size, which means that only large consumers can participate individually. However, aggregating flexibility among a portfolio of smaller consumers and even households can overcome this requirement.

Looking into the near future, a CO2 tax or tariffs and electricity taxes that vary with the hourly carbon intensity further increases the business case of intelligently scheduling flexible consumption. This is especially important for household consumers in Denmark, where tariffs and taxes account for almost two thirds of the electricity price[4].

The transition to 100 % renewable electricity in the interconnected European electricity network requires a combination of transmission expansion, different storage technologies and flexibility. This demonstration project has shown how to automatically utilize flexibility in electricity consumption in Denmark requiring little to no capital cost for the flexible consumer. As next steps, we are engaging in additional real-world use cases and developing the forecasting and scheduling frameworks further. In addition, we are integrating this technology further into our carbon efficiency products and investigating automated flexible demand detection algorithms. The aim is to provide the easiest way for energy users to become as carbon efficient as possible, accelerate the transition to a sustainable energy system and help limit global warming.

[4] https://ec.europa.eu/eurostat/statistics-explained/index.php/Electricity_price_statistics#Electricity_prices_for_household_consumers

References

Bokde, N., Tranberg, B., Andresen, G. B. (2020). Short-term CO2 emissions forecasting based on decomposition approaches. *In Preparation*.

Bronski, P., Dyson, M., Lehrman, M., Morris, J., Touati, H. (2015). The Economics of Demand Flexibility: How "Flexiwatts" Create Quantifiable Value for Customers and the Grid. Rocky Mountain Institute.

Danish Energy Association. (2019). Elbilerne kommer – det er et politisk valg, hvordan vi sikrer, at elnettet bliver klart.

Henderson, P., Hu, J., Romoff, J., Brunskill, E., Jurafsky D., Pineau J. (2020). Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning. *Working Paper*. arXiv:2002.05651

Klein, K., Killinger, S., Kalz, D., Herkel, S., Salom, J., & Cubi, E. (2016). Comparison of the Past and Future Residual Load in Fifteen Countries and Requirements to Grid-Supportive Building Operation. *Proceedings of EuroSun2016*. doi:10.18086/eurosun.2016.09.07

Lacoste, A., Luccioni A., Schmidt, V., Dandres, T. (2019). Quantifying the Carbon Emissions of Machine Learning. *Preprint*. arXiv:1910.09700

Lottick, K.,Susai, S.,Friedler, S. A., Wilson J. P. (2019). Energy Usage Reports: Environmental awareness as part of algorithmic accountability. *NeurIPS 2019*. arXiv:1911.08354

Schwartz, R., Dodge, J., Smith, N. A., Etzioni O. (2019). Green AI. *Preprint*. arXiv:1907.10597

Tranberg, B., Corradi, O., Lajoie, B., Gibon, T., Staffell, I., & Andresen, G. B. (2019). Real-time carbon accounting method for the European electricity markets. *Energy Strategy Reviews, 26*, 100367. doi:10.1016/j.esr.2019.100367

Tranberg, B., Hansen, R. T., & Catania, L. (2020). Managing volumetric risk of long-term power purchase agreements. *Energy Economics, 85*, 104567. doi:10.1016/j.eneco.2019.104567