Home-energy forecasting for storage optimization

Elena Popovici, PhD
Senior Computer Scientist

neurio

neurio

Neurio: Smarter Energy Management

Industry-Leading Hardware
Over 150,000 installations
Easiest sensor to install <15 minutes
Custom CT design gets into more panels
Real-time granular data

Intelligent Software
Sophisticated analytics
Highly engaging frontend user experience
API offering

Neurio makes hardware and software for electricity management, currently primarily targeted at home-level management.

The Neurio sensor goes into your breaker panel and then you can use an app either on your computer or on your phone to see your electricity consumption and your generation, if you have a generation source like solar panels.

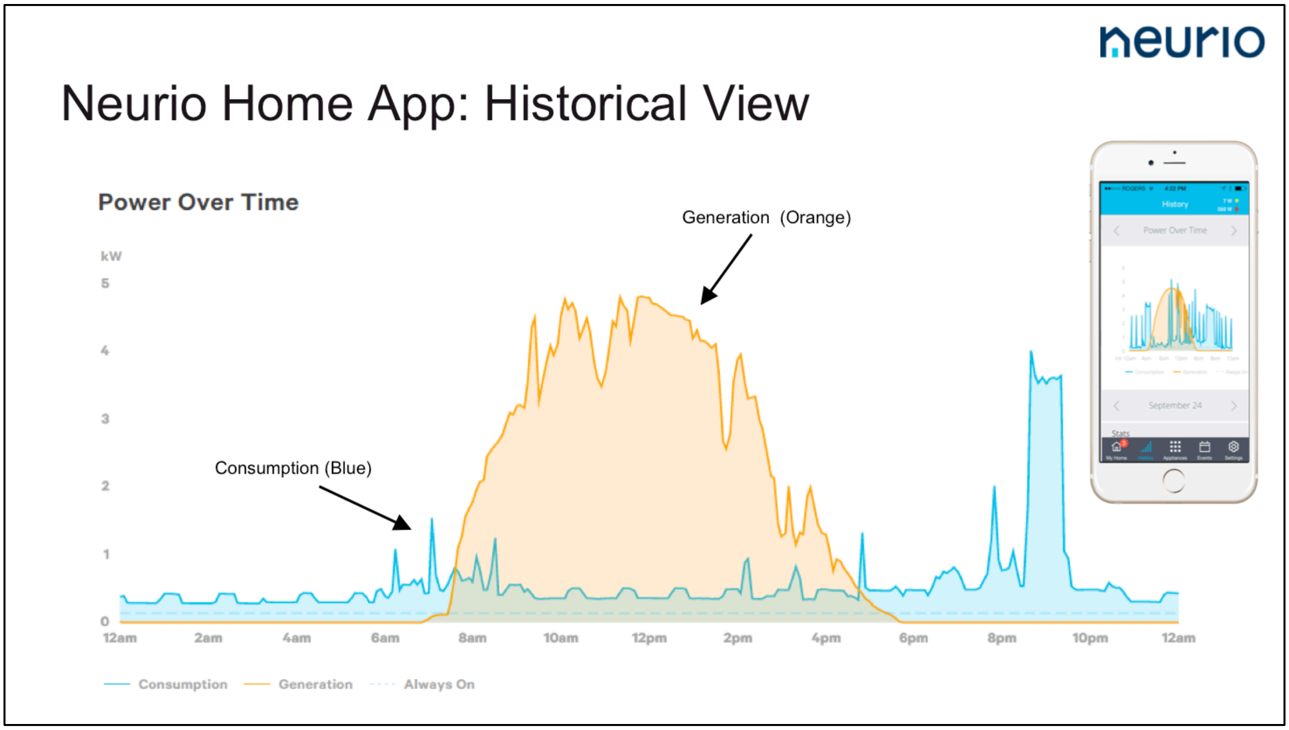

It’s similar to a smart meter, except the data is available instantaneously and with secondly granularity, unlike, say, the BC Hydro website which shows you hourly aggregates with one day delay.

This knowledge may empower the user to change their consumption patterns, in fact Neurio was initially called energy-aware.
Being aware of when and how much energy one uses is a necessary tool for behavior change, but the main driver for people is usually cost.
So utility companies try to encourage behavior change through pricing structures.

Utility Rates

- Tariff Components
  - Tiered pricing
    - E.g. BC: in each 2 month period, 8.29 cents per kWh for first 1,350 kWh and 12.43 cents per kWh for everything over that threshold
  - Time-of-use pricing
    - E.g. **CA**, AZ, HI
  - Demand charges
    - E.g. AZ: $17.438 per kW for highest hourly consumption over entire month
  - Combinations of the above, seasonal differences, minimum charges and more!

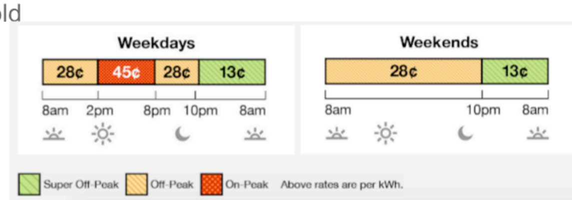

- Rules
  - Allow/disallow selling solar to grid
  - Allow/disallow charging battery from grid
  - Allow/disallow selling from battery to grid

=> Custom battery operation strategies

An electricity bill can have multiple components.

For instance, here in BC our billing tariff is really simple, we only have tiered pricing: in each 2 month period, one rate per kWh for first 1,350 kWh and a higher rate for everything over that threshold.
But in other regions electricity billing is a lot more complicated. For instance, it can involve different rates for different times of day, such as in California, shown on the right.

Tariffs can also include a charge for the peak consumption over the month. For instance, in Arizona this can be a steep $17/KWh.
There can also be seasonal differences, minimum charges and combinations of all of these components.
There are also rules restricting selling and buying of energy from the grid.

For a human to change their habits in a way that reduces their bill may be challenging, both because of complexity of the tariff and for constraints such as “can’t do laundry while not at home” and “need a comfortable temperature while at home”.

This is where Neurio saw an opportunity and has built custom, parameterized, battery operation strategies for a number of tariffs. Because when you have a battery, you can essentially shift electricity around.

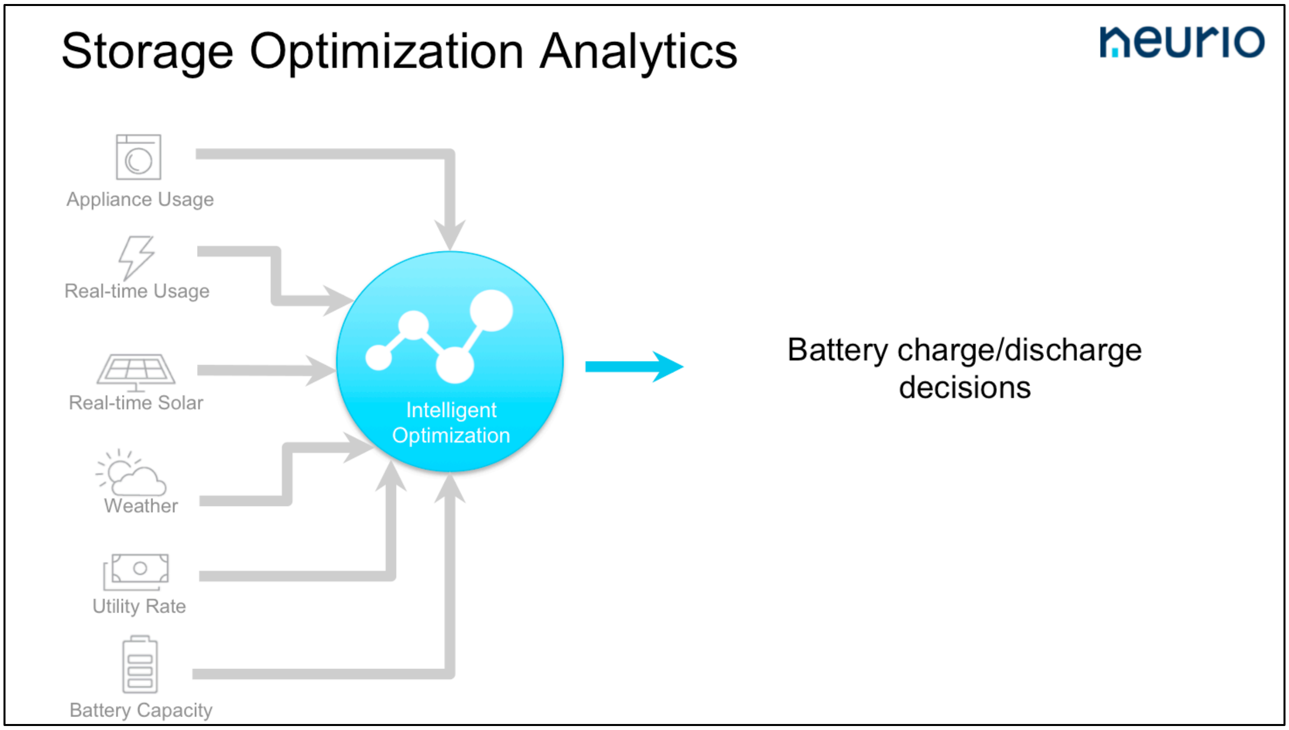

These strategies are hand-designed optimization heuristics that take in information about real-time electricity usage, power generation from solar, weather, the billing tariff of the utility company, and the battery capacity, and make decisions about when and how much to charge/discharge the battery so as to result in as low a bill as possible.

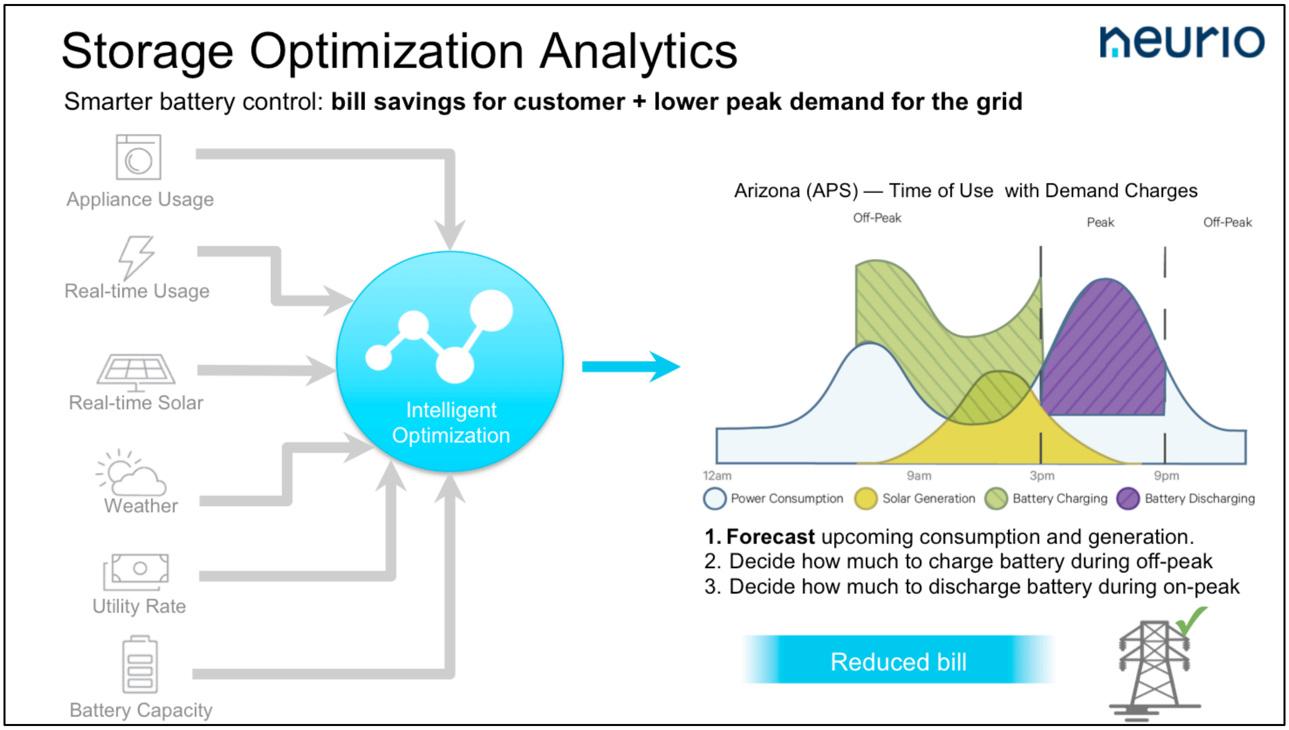

For instance, in Arizona, which has two time-of-use periods, peak and off-peak, as well as demand charges during peak, it is obviously beneficial to reduce the amount drawn from grid during peak by discharging the battery, and charge the battery during the off-peak, both from excess solar and, if allowed, from the grid as well.

This saves the consumer money AND it's also good for the grid, because it reduces the demand for electricity during the peak period, which means reducing the need for expensive infrastructure whose sole purpose is to cover that peak.

The devil is in the details. To determine how much and at what rate to charge and discharge the battery, it is very useful to be able to forecast both the upcoming consumption and the upcoming generation from solar panels, and it is this FORECASTING piece that my work at Neurio has been focused on.

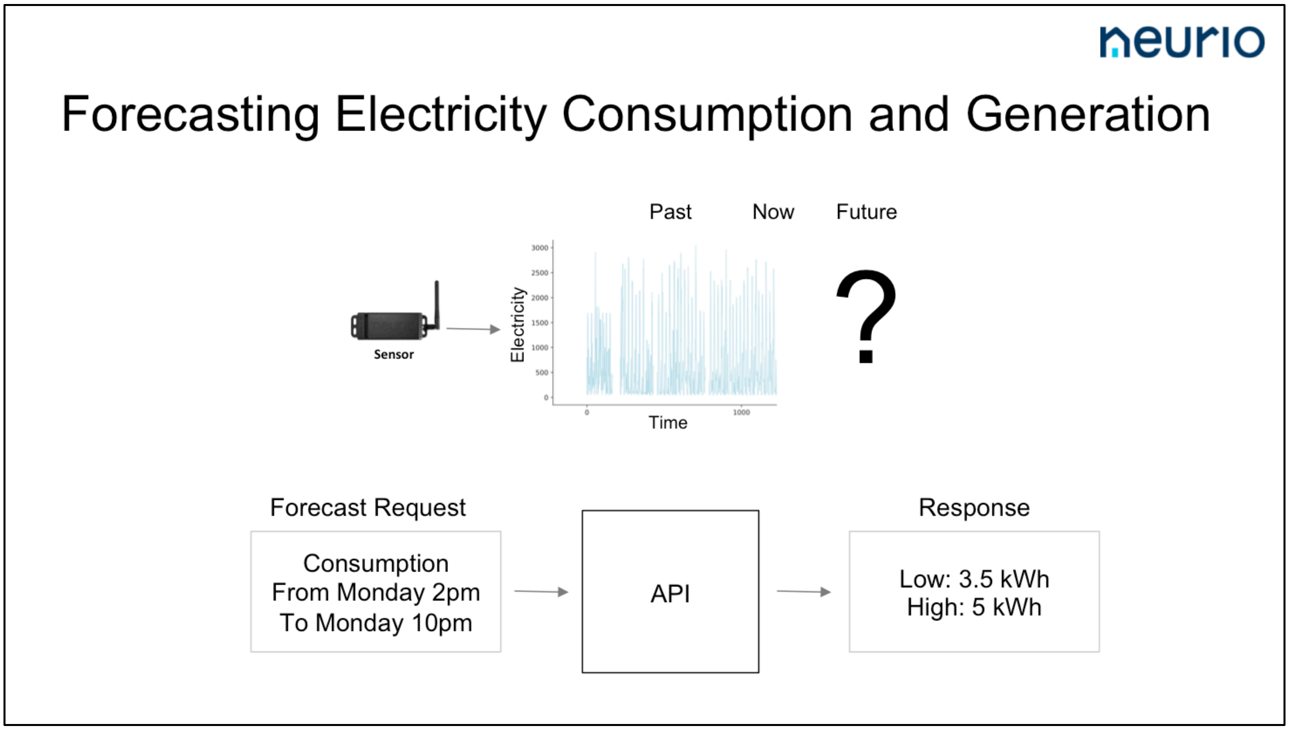

How do we do forecasting?

Over time, the sensor gives us a time series of consumption and/or generation, and we would like to learn from the past in order to predict the future.
So we're building an API that can take in requests such as "tell me what consumption will be this upcoming Monday from 2pm to 10pm" and responds with a low and a high estimate of consumption for that time period.

Inside Forecasting

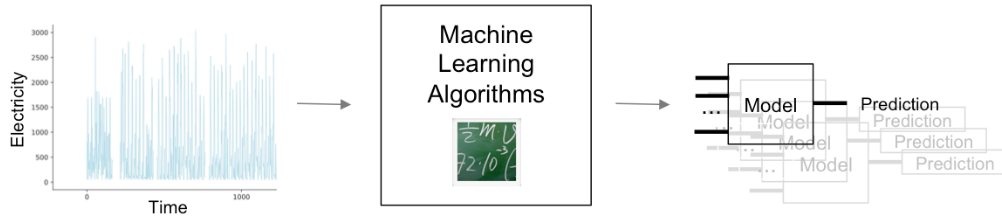

- Many decisions:
  - Which type of ML algorithm? (e.g. neural nets, SVMs)
  - Which parameter values for that algorithm?
  - Which “features” to use as inputs? (e.g. prior consumption, day of week, month of year, etc.)
  - Which part of the data to use for training?
  - ...

To do that, we use supervised machine learning algorithms, that take in a time series of historic data, and spit out so-called “models” or **predictors**, where a predictor is a box that takes inputs such as day-of-week, consumption during the past 12 hours, etc. and spits out a prediction for, say, consumption over the next 6 hours; or maybe it takes in weather forecasts and spits out a prediction for solar generation for the next day, etc.

Building machine learning algorithms involves many decisions, as listed here.

We use a multitude of approaches to make these decisions.

For instance, for the highlighted types of decisions we use a combination of hard-coding based on experience, tuning and automation, as follows.

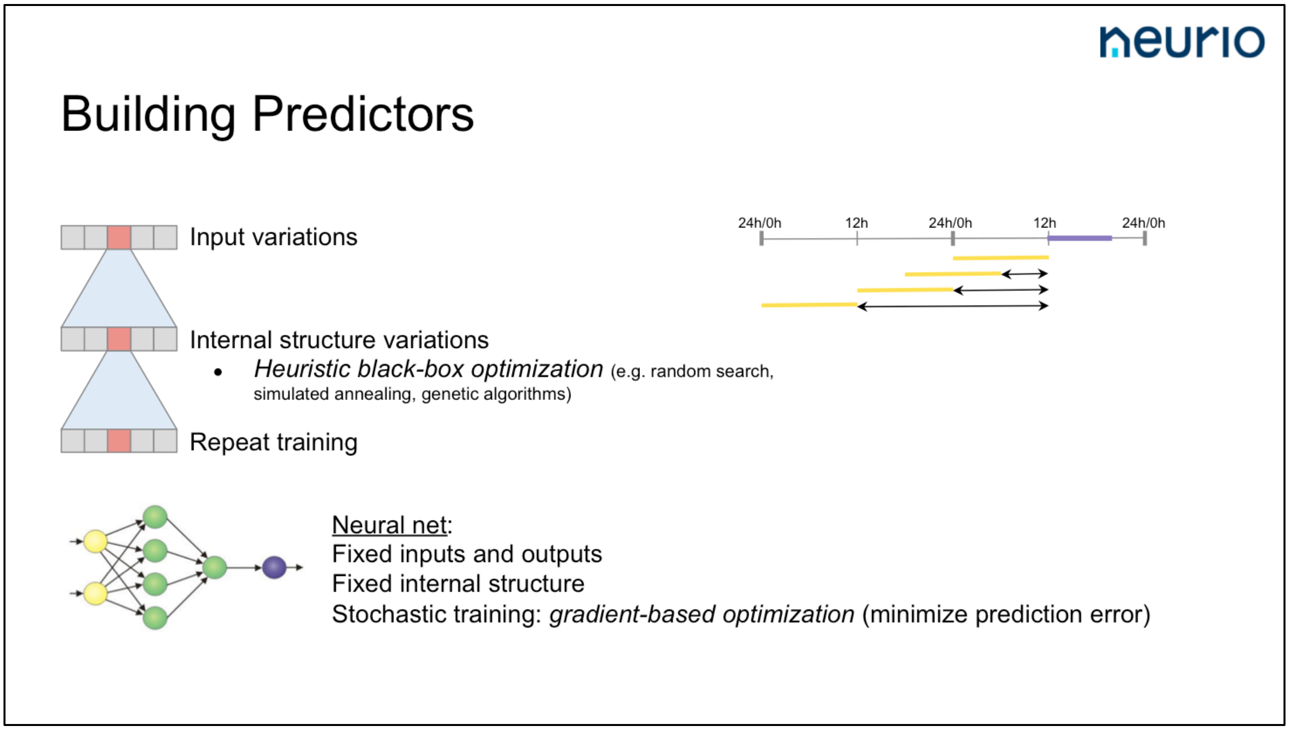

Since we’re dealing with a regression problem (meaning we need to predict numeric quantities) and we know the problem is not linear in nature, we decided to go with neural networks.

For those not familiar with neural networks, they’re essentially directed graphs performing computations in the nodes using weights associated with the edges. A given network has fixed input and output nodes and a fixed internal structure, meaning number of layers, numbers of nodes in each layer and computation functions inside the nodes. The weights on the edges are determined via a so-called “training” process that aims to get the network to fit a set of input-output pairs called the “training data”. The inputs are run through the network, an output is computed and compared with that expected from the data in order to compute a prediction error. The error function has a closed form whose derivative in the weights can be computed, thus error minimization is done via **gradient-based optimization**. This training process is stochastic in nature, therefore the first thing we do is to repeat it multiple times to get better estimates of the success of training.

Then, since there is no a-priori good network structure, we try multiple such structures. This process can be automated using **heuristic black-box optimization** techniques.

All of this hierarchical process produces one predictor for a given type of inputs. Then we go on to repeat the whole process for different input types. We do this for the following reason: suppose we want to predict consumption for the immediate next 6 hours. Then we can use models that have as inputs previous consumption up to this very moment. But if we are trying to predict today consumption for a certain 6 hours tomorrow, then we need models that have a temporal gap between the consumption inputs and the outputs. We call this gap an offset and we build multiple models with different offsets, then at prediction time we determine which ones we can actually use, given how far into the future we are trying to predict.

Inside Forecasting

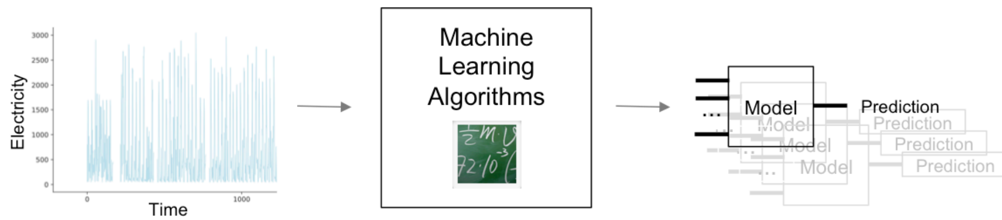

- Many decisions:
  - Which type of ML algorithm? (e.g. neural nets, SVMs)
  - Which parameter values for that algorithm?
  - Which “features” to use as inputs? (e.g. day of week, month of year, prior consumption, etc.)
  - Which part of the data to use for training?
  - ...

What about the other types of decisions? For instance, in addition to historic consumption, should we use as inputs things such as the day of the week or the time of the year? And which historic data should we train with? The last week? The last month? The last quarter?

For these kinds of decisions, we did some offline **exploratory data analysis** that I will detail in the remainder of this talk.
In particular, we wanted to know if an unsupervised machine learning technique like clustering can help answer those questions, and of course, data visualization is always a big part of such exploration.

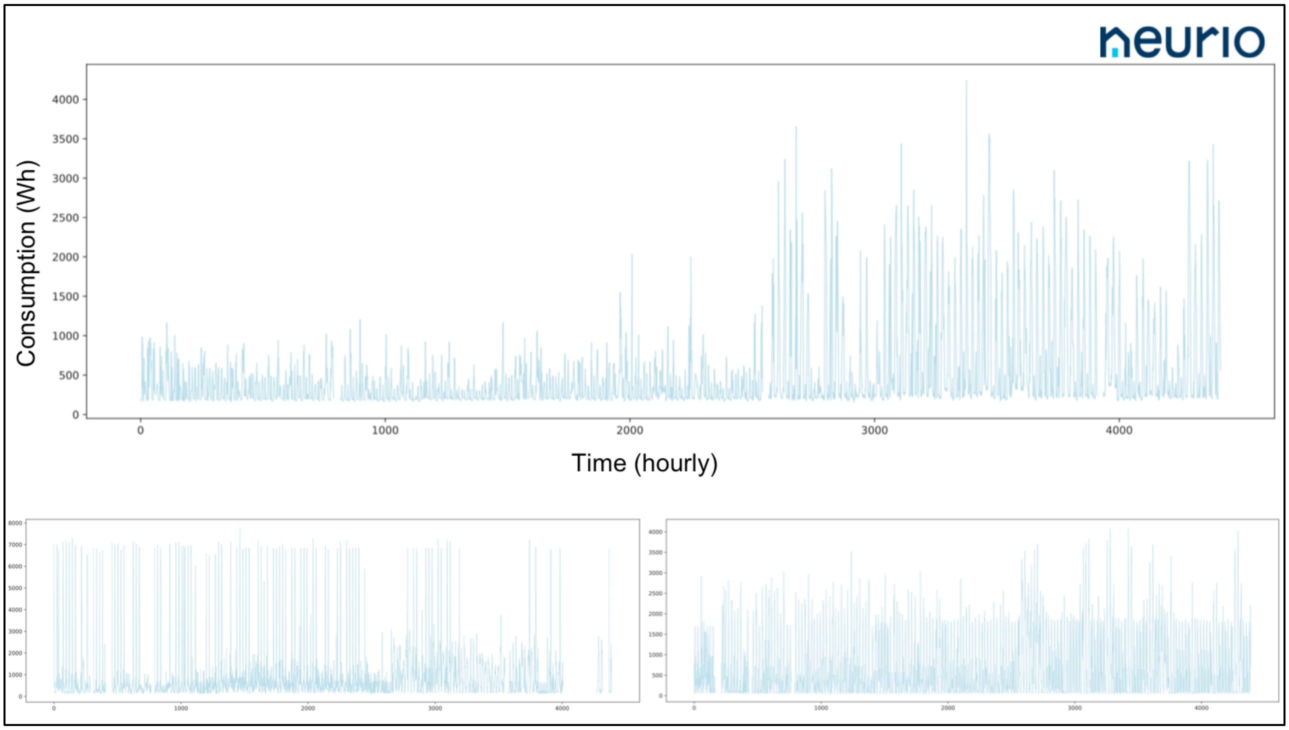

We're going to focus on consumption only, and we start with a time series of hourly aggregates for one home, spanning several months. Here are three example homes.

In the top one, it's fairly obvious by eye that this home had an increase in consumption about two thirds of the way through the period. But we cannot tell much more from this plot.

The bottom left one one has some interesting spikes that became less common towards the end of the period.

The bottom right is one where I, for one, find it very difficult to describe what, if anything, is changing over time.

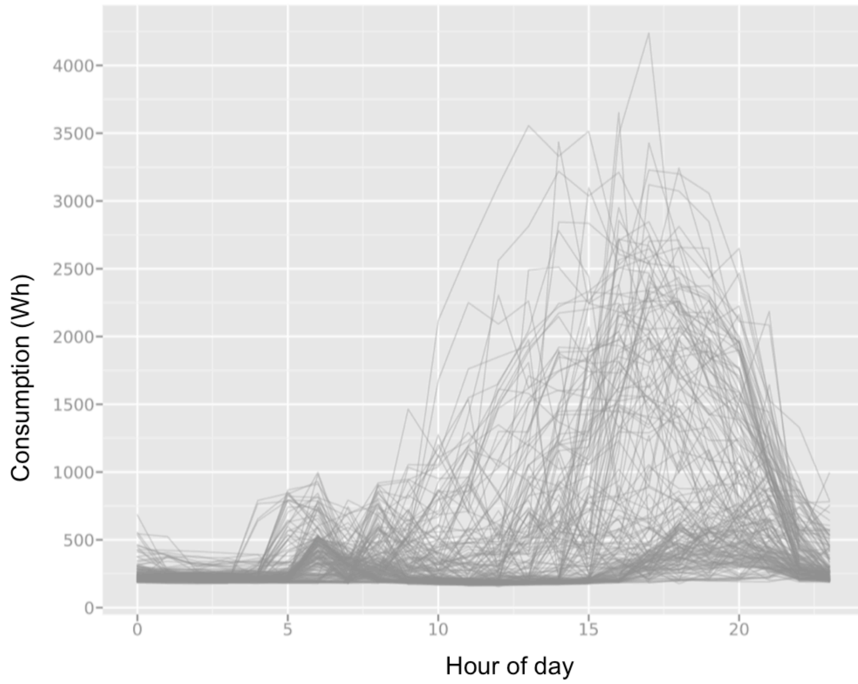

Let’s focus on the first example home and try a different way to look at the data. Instead of plotting the days in sequence, we overlap them. The X-axis is now hour of day and each line represents a different day.

We start to notice some patterns in this plot. Most days have a spike in the morning around 6am and then there’s another increase in the afternoon/evening, but for some days it’s small and for others is really large. So maybe there’s two types of days, but it’s not super obvious.

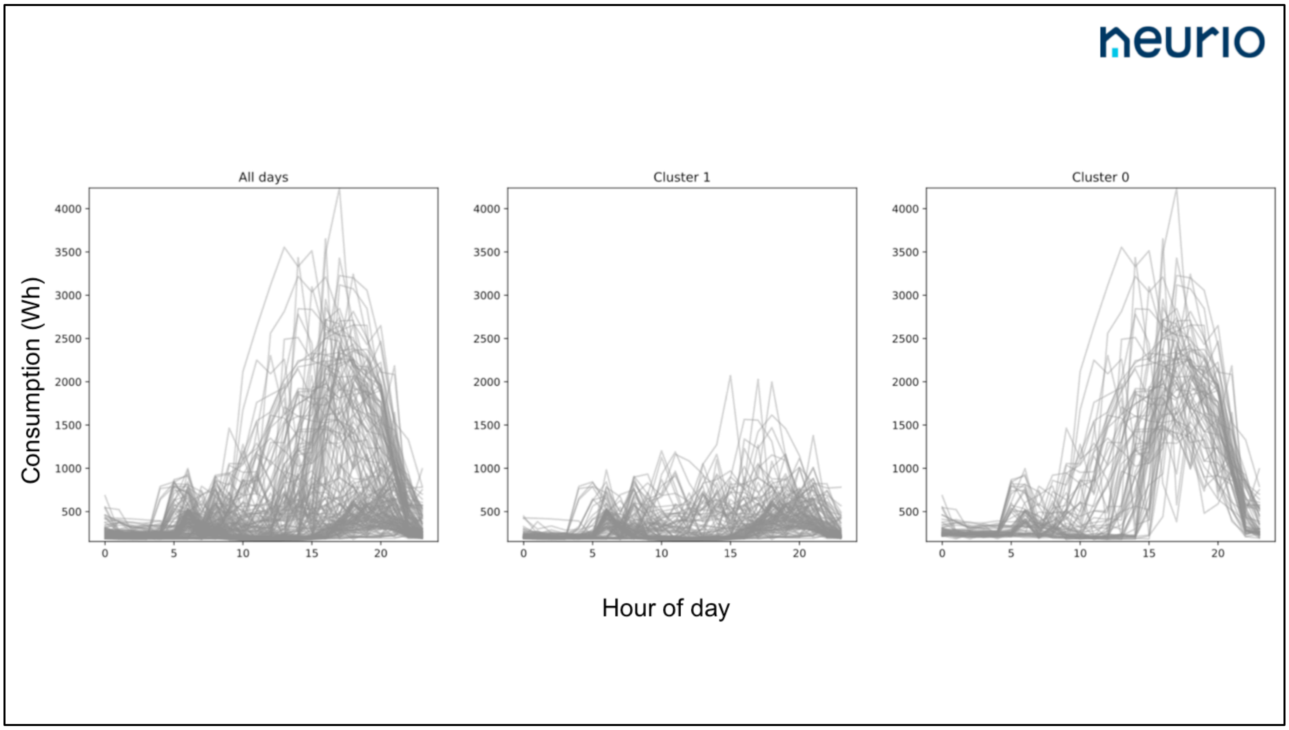

So the next thing I did was to feed this data through a clustering algorithm and ask it to find 2 clusters. Here's what it produced.
On the left we still have the original plot, with all the days, and then the next two panels are subsets of days that the algorithm placed in different clusters. It's now very clear that there are indeed (at least) two different kinds of days, one w/ a high afternoon peak, one without.
What might these differences have to do with?

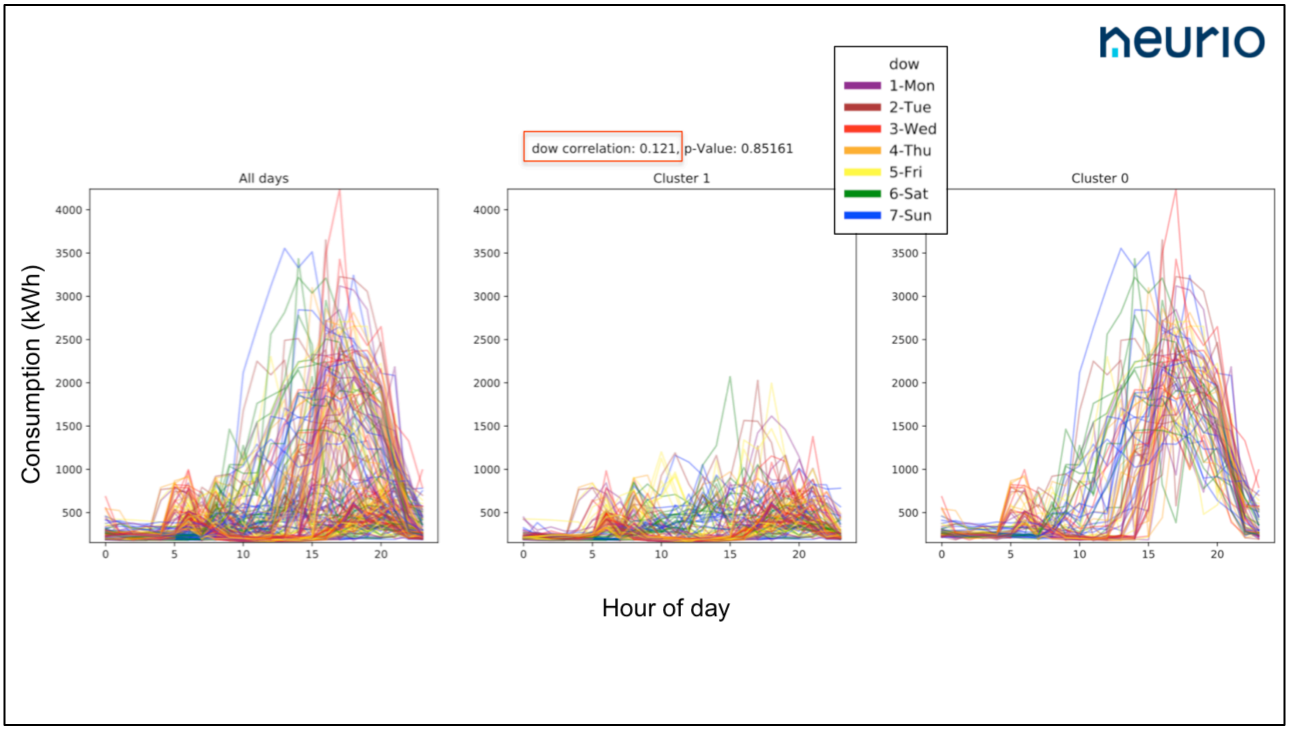

To answer that, I took the same plots and colored the lines by day of week. You can see that both clusters mix all colors, so day of week is not what differentiates them. In fact, I also computed a Spearman correlation between the cluster each day was placed in and the day of week, and the correlation was very low.

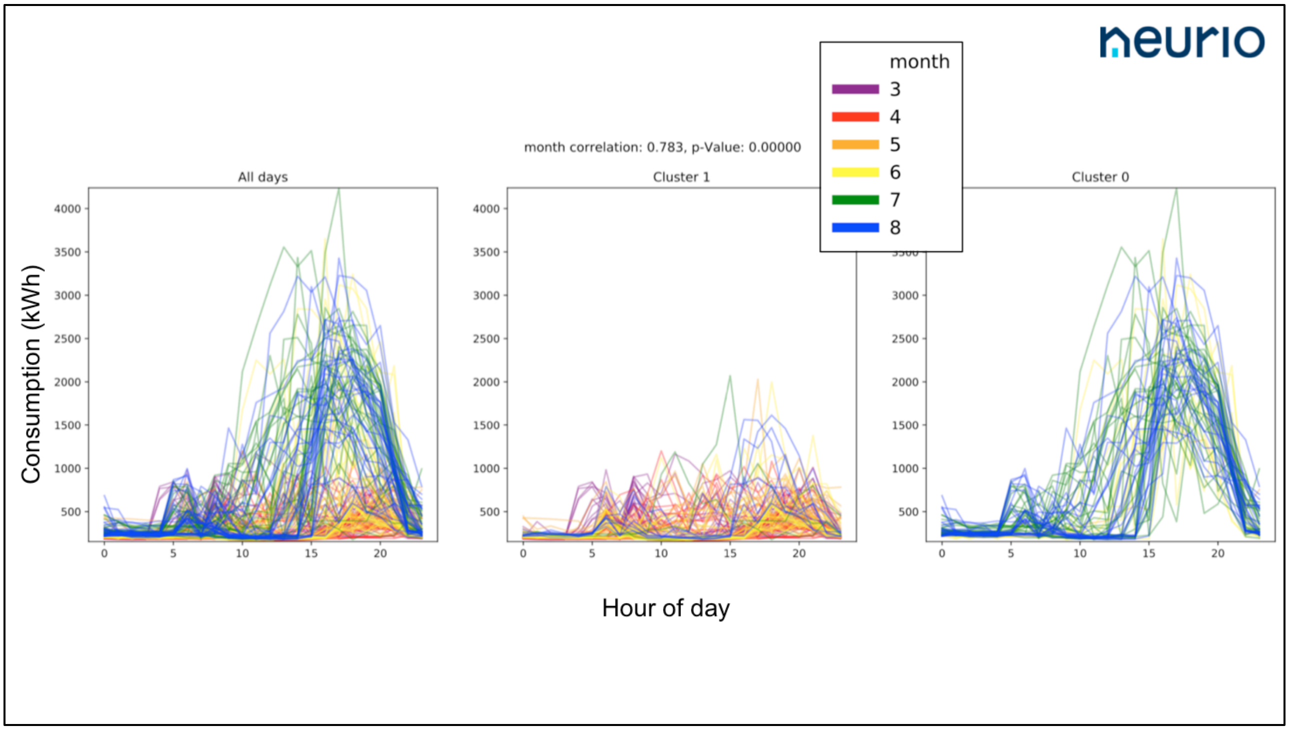

When we color the same plots by month of year, they look quantitatively very different. Most of the blue and green, which are July and August, are in one cluster and the rest of the colors are in the other cluster.

So the shift in behavior that we saw in the original time series chart we can now pin-point to being due largely to a change in afternoon/evening behavior for that home.

How about a way to see both day-of-week relevance and month relevance?

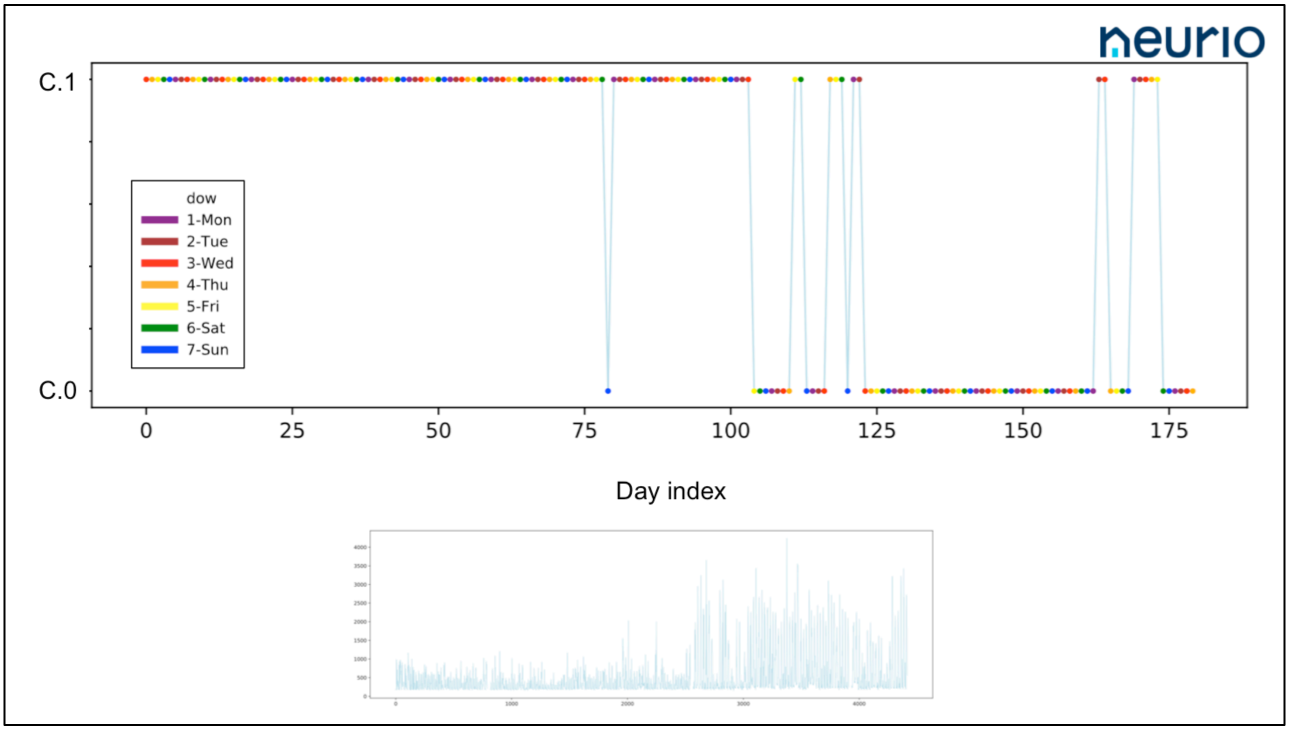

Here it is. X-axis is time and each point is a day. Y-axis denotes which of the two clusters the day was assigned to. The color of the dot denotes the day of the week. So we can see in a single plot that the clustering is not related to the day of week, but it is related with the time of year.

From a computational standpoint, determining where a change-point has occurred is much easier to do with this daily time series of clusters than with the original hourly energy time series.

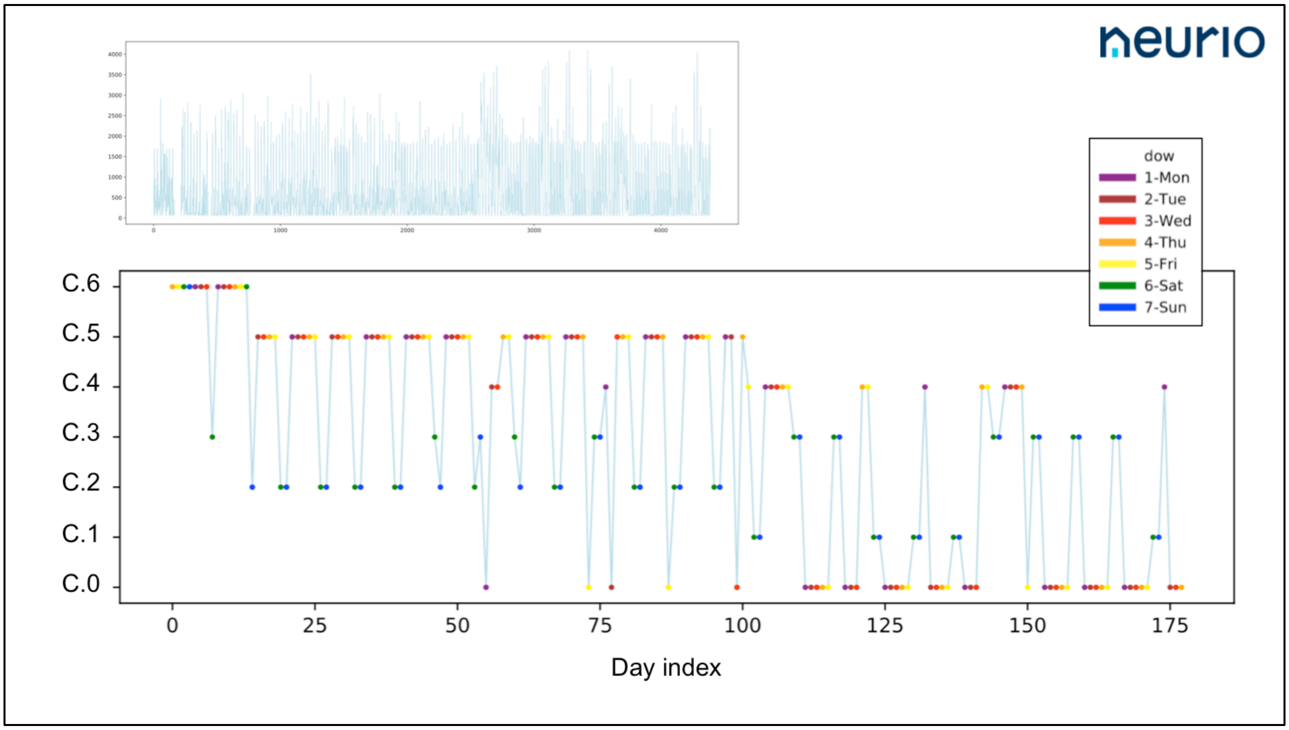

The number of clusters is a parameter of the clustering algorithm. Here's an example where we set that parameter to 7 for sample home #3. We can see from the chart that some clusters contain only weekend days and some clusters contain only week days. But we can also see a time of year effect, where about two thirds of the way through the period both the workday clusters change (from 5 and 6 to 0 and 4) and weekend clusters change (from mostly 2 to 1 and 3).

None of this was distinguishable in the original hourly plot.

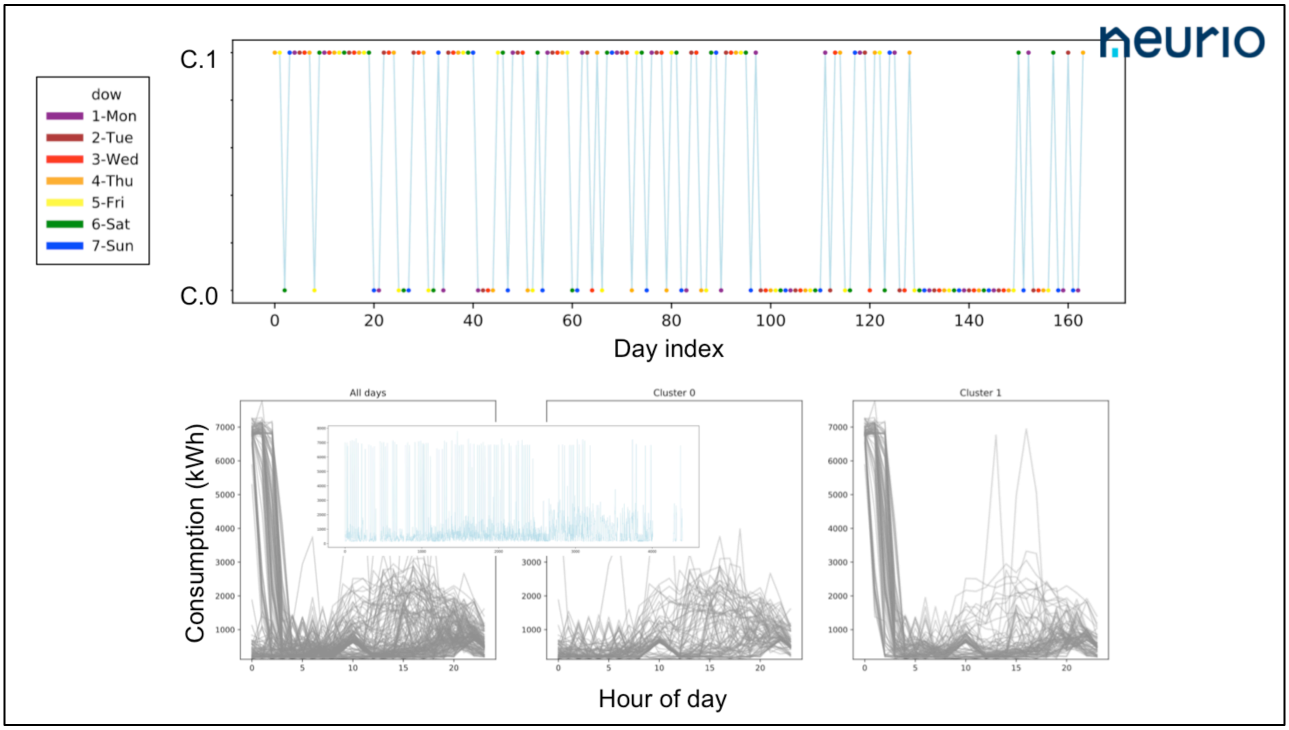

And here is an example, where the two clusters found don't appear to have anything to do with either the day of week or the time of year.
When we look at the shape of the day profiles in those two clusters, we can see it identified which days have a peak at midnight and which don't. And by the way, this type of peak corresponds to charging an electric vehicle.
This, by the way, is for sample home #2.

Exploratory Data Analysis Summary

- Can clustering help? YES!
  - Time-series of clusters => easier change-point detection => which data to use to train models
  - Cluster-feature correlation => which features to include in models

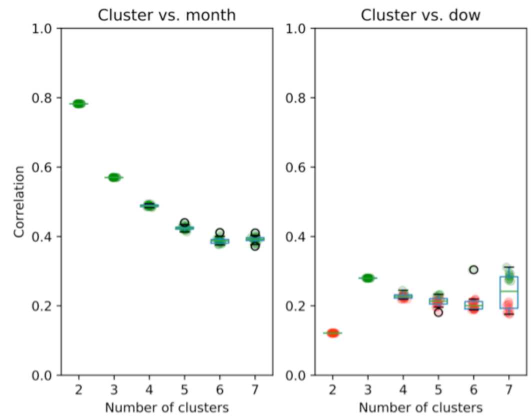

- Other uses
  - User-facing visuals (inclusion in app)

So, can clustering help us answer the questions we set out at the beginning? Yes, it can!

Transforming the data from an hourly energy time series to a clusters time series can help us determine whether there was a shift in behavior and therefore we may need to retrain our model with new data.

Computing correlations of clusters to either day of week or month of year can tell us whether we should include those features into our model. In fact, here is another visualization summarizing correlations with either day of week or month of year for different numbers of clusters. Green means the correlation was statistically significant, red means it wasn't.

And data visualization played a big role in helping draw all these conclusions.

Moreover, the exploratory visualizations produced in the process give us a starting point for designing visuals for including in the app to give users deeper insights into their electricity usage.

neurio

Conclusions

- Hand-designed optimization heuristics
- Supervised and unsupervised machine learning
- Gradient and black-box optimization
- Exploratory data analysis
- Visualization

→

Solar

+

Battery

=

Grid-friendly

By using a variety of computational techniques we can ensure that a sustainable energy source like solar can be made grid friendly in spite of the fact that it is intermittent and its generation is misaligned with consumption needs. And all of this could be adapted to work for other sustainable energy sources beyond solar and other storage systems beyond batteries.