Methods

How the 1 km electricity consumption data is built, what it is anchored to, how it is checked, and what it can and cannot tell you.

`Describes release 2026.10.4` Updated October 30, 2026

How the product works, for people who use data, describing release 2026.10.4. "In brief" is the quick overview; the sections after it go deeper. The Data Guide covers the files themselves; every accuracy chart is on the Countries & validation page.

In brief

- **What it is.** Annual electricity consumption for every 1 km² cell of Great Britain, France and its five overseas territories, Canada, South Korea and Australia, for 2021 to 2024 (2017 to 2025 in Australia), split into residential, non-residential and large facilities, on a gross and a net basis. The values are modeled estimates, not measurements.
- **How it is made.** Official statistics set every region's total. One model, trained on about 75,000 small areas with metered consumption, decides how each total spreads across the cells inside the region, from what satellites and censuses show about buildings, population, land use and climate. 779 smelters, steelworks, refineries, mines and other large sites are estimated one by one and placed at their own locations.
- **How good it is.** Scored on areas the model never saw, against a population split given the same official totals. The model's typical miss on a municipality's total is 101 GWh a year, against 175 GWh for the population split, on municipalities that average 260 GWh; for a country with no local data it is 160 GWh against 288 GWh. Every score below is an error: smaller is better.
- **How to use it.** Add pixels up and never compare one pixel with one meter. Where a region has an official figure, the map reproduces it. Use net to

compare with grid data and gross for energy-use questions. Great Britain also carries an experimental hourly preview on the map.

- **Beta.** Expect some errors. Check against local figures you trust before relying on the data for a decision that matters, tell us what you are using it for, and report anything that looks wrong.

What the product is

For every 1 km² cell in a covered country, the product gives annual electricity consumption in gigawatt-hours (GWh), split three ways and reported on two bases:

- **Three sectors:** residential, non-residential, and large facilities. They add up exactly to the total in every cell.
- **Two bases:** gross and net. The difference is the rooftop solar used on site and the power large facilities generate for themselves.

The values are modeled estimates, not measurements. Each region's total comes from official statistics; the model decides how that total spreads across the cells inside the region.

Country or territory	Years	Official figures the totals follow	Reproduced exactly for
Great Britain, with Northern Ireland for 2021 to 2023	2021 to 2024	DESNZ subnational electricity consumption (total final energy figures for Northern Ireland)	Every local authority district
France (metropolitan, including Corsica)	2021 to 2024	Agence ORE annual consumption	Every intercommunality (EPCI)
Guadeloupe, Martinique, French Guiana, Réunion, Mayotte	2021 to 2024	Agence ORE annual consumption	Every commune
Canada	2021 to 2024	Statistics Canada provincial supply and demand	Every province and territory; postal areas in Ontario (IESO) and municipalities in British Columbia and Quebec shape the residential split within them
South Korea	2021 to 2024	KEPCO electricity sales	The national total and KEPCO's figures for each of its 229 reporting districts
Australia	2017 to 2025, financial years	Australian Energy Statistics	Every state; distributor billing shapes the split between council areas where a distributor publishes it: Queensland, parts of Victoria, and for the residential sector from 2019 the Ausgrid area of New South Wales

"Reproduced exactly" means the map's total for that unit, on the basis the figure was published, equals the published figure (Australian states to within a fraction of a percent). Where a finer official source only shapes the split, the map's total for each unit differs from its published figure by design: by a median of about 6% for Ontario postal areas and 5% to 9% for Australian council areas (residential to non-residential) in the latest year. Australia's years are financial years ending 30 June: a year labeled 2025 runs from July 2024 to June 2025. Northern Ireland has no 2024 figure yet, so its cells are blank that year.

How it is built, in five steps

1. **Start from official statistics.** For each country we take the most detailed official consumption figures published, listed above. These set the level.
2. **Learn from metered data where it exists.** Several countries publish consumption for small areas: neighborhoods in Great Britain, IRIS zones in France, postal areas in Ontario, postal areas and council areas in parts of Australia. We train one model, across all of these places at once, on how consumption relates to things we can observe everywhere from satellite and census-derived data: built-up area and volume, population, land use, climate, and the kinds of buildings present. About 75,000 metered areas, each with several years of data, go into training, most of them in Great Britain and France.
3. **Spread each region's total across its cells.** The model predicts a consumption intensity for every cell. The cells inside each region are then scaled so they add up to the region's official figure. The model decides the pattern; the statistics decide the amount.
4. **Place large facilities one by one.** Smelters, steelworks, cement, chemical and glass plants, refineries, mines, paper mills and similar sites are identified individually and estimated per site. In the raster each site is spread over about 2 km around its location, inside its own region; the points file holds the exact location and value. The rest of the region's total is spread without them, so a factory's electricity does not land on the homes and offices around it.
5. **Check the result against areas the model never saw.** Before training we hide whole groups of neighboring districts, and we also refit the model with whole countries left out. The accuracy numbers we publish are scored on those hidden areas and held-out countries, against their published figures and against the simplest alternative: splitting the same total by population.

The three sectors

Sector	What it contains
Residential	Electricity used in homes.
Non-residential	Everything that is not a home and not a large facility: shops, offices, schools, hospitals, public buildings, light and medium industry, farms, transport, street lighting, and identified industrial sites too small to count as large facilities.
Large facilities	Sites we identify one by one that use 10 GWh or more in at least one year we have data for. A site that reaches that size counts as a facility in every year, including years it ran below it or was closed.

Total = residential + non-residential + large facilities, exactly, in every cell and on both bases. Official statistics usually publish two categories, such as domestic and non-domestic: residential corresponds to the first, and non-residential plus the large facilities those statistics count corresponds to the second. Data centers are not identified as facilities yet, because no reliable public list of them exists; their electricity sits inside non-residential.

Gross and net

Net is what the grid delivers to the premises. **Gross** is net plus the self-supply we can estimate: rooftop solar used on site, where a rooftop-solar model exists, and the power a large facility generates for itself, where we know it. Neither basis includes backup or off-grid generators, off-grid solar, demand that went unserved, or network losses: every value is consumption at the customer's meter.

Two examples. A home with rooftop solar uses 4,000 kWh a year (gross) but draws 3,000 kWh from the grid (net). A smelter with its own hydro plant uses 5 TWh a year (gross) but imports 1 TWh (net).

Where rooftop solar is modeled

Everywhere in Great Britain and Australia, and for the residential sector in Ontario within Canada. Elsewhere, including France, its overseas territories, South Korea and Canada outside Ontario, the published figure for homes and buildings is the metered one and gross is set equal to net.

▾ How the rooftop-solar estimate is made, and where to check it

The estimate starts from installed rooftop capacity, year by year (official installation statistics by area in Great Britain, the Clean Energy Regulator's

postcode data in Australia, one provincial figure in Ontario), a typical yield for the local climate, and a fixed share of that generation used on site; home batteries are not modeled, and what the panels export is in neither basis. The pixel column `no_btm_model` says whether a whole country has a model; the region file's `btm_covered` says it region by region, which is what to check in Canada. The rooftop-solar estimate is not validated; treat gross minus net as an estimate without an error band.

Large facilities on the two bases

Net is the site's grid import where we know its own generation; exports are not subtracted. Where we do not know, net equals gross and the site is flagged (`facility_basis = self_gen_unknown`), which overstates grid draw where self-generation is common and unrecorded.

Which basis to use

Net when you compare with metered or network data, or feed a model with no behind-the-meter generation. Gross for electrification planning, energy-use statistics, a national energy balance, or a model that already represents rooftop solar and on-site plants as generators (net there would count their output twice). Gross minus net, for homes and buildings, is rooftop solar used on site where it is modeled.

Large facilities

A site counts as a large facility when it uses 10 GWh or more in at least one year, roughly 1 MW running continuously. Sites come from the Climate TRACE facility inventory and from national registers. Each estimate starts from what we know about the site: its production and a typical electricity use per tonne of product for that kind of plant, a figure the company or a regulator has published, or a sizing from its reported emissions. Every fact about a site carries a source, and every estimate a confidence grade.

What this means when you read the data:

- Every site with an estimate carries a confidence grade, its operating status with the years it applies, and its sources; most also carry a low-to-high range, the spread of the figures behind the estimate. A closed site carries zero in the years after closure.

- Every estimate has a confidence grade. `high`, `medium` and `low` mean a published figure stands behind it; `unverified` means a modeled estimate with none, which the map labels "Modeled estimate". The share of each country's facility use with a published figure behind it, in the latest year: Great Britain 31%, France 25%, South Korea 13%, Canada 5%, Australia 3%.
- Sixty Australian mines have production figures but no reviewed electricity-per-tonne value yet. They appear as points marked "no electricity estimate yet" and add nothing to the raster.
- The sites inside a region's official statistics together never take more than half of that region's official non-residential figure in the raster. Where they would, each is scaled down, the point shows both the estimate and the mapped amount, and a site left with less than half of its estimate is flagged `scaled_to_fit`.

Why a region's total can differ from its official figure

Official statistics do not all count the same things. Sites those statistics leave out are still on the map, on top of the official total.

Country	What the official figures leave out	Effect on the map
Great Britain	Steelworks and aluminum smelters connected straight to the transmission grid, and the electricity any site makes for itself	A local authority holding such a site exceeds its official figure: on gross by the site's full use, on net by its grid draw
France, Martinique, South Korea	The electricity sites make for themselves	On gross, a region holding such a site exceeds its official figure by that own generation; net matches
Australia	Council-area billing leaves out the largest smelters, alumina refineries and a few mines, which are inside the state totals	The council areas holding them read far above their billing figures; the state totals still hold
Canada	Nothing: the provincial figures include everything	The facility layer stays inside them

Every facility point says how the official statistics count it (`frame`), and each country's metadata file gives the yearly amount outside them. Facility estimates are made site by site, so the hidden-area test below does not score them; instead the largest sites are compared with figures grid operators publish, and a mismatch is flagged, never used to change an estimate. Those comparisons are on the Countries & validation page.

Trained, partly trained, and anchored

Every country's totals follow official statistics. The labels say how much local small-area data the model learned from, which is what decides how good the pattern inside each region is:

- **Trained:** the model learned from consumption published for small areas here. Great Britain, metropolitan France and the five overseas territories.
- **Partly trained:** local small-area data covers part of the country. Canada, for the residential sector only: Ontario postal areas and municipalities in British Columbia and Quebec. Australia: council areas and postal areas in Queensland and Victoria, and suburbs in the Ausgrid area of New South Wales.
- **No local data (anchored):** no local small-area data. South Korea: KEPCO's sales set the national total and the split between its 229 reporting districts; the model trained elsewhere draws the pattern inside each district.

Where each applies, with where rooftop solar is modeled (the gross basis adds it; see "Gross and net" above):

Country or territory	Local training data	Rooftop solar modeled
Great Britain, with Northern Ireland for 2021 to 2023	Trained on local data	Everywhere
France (metropolitan)	Trained on local data	No model yet: gross equals net for homes and buildings
Guadeloupe, Martinique, French Guiana, Réunion, Mayotte	Trained on local data	No model yet: gross equals net for homes and buildings
Canada	Partly trained: Ontario, British Columbia and Quebec, for the residential sector only	Ontario only, residential
South Korea	No local data	No model yet: gross equals net for homes and buildings
Australia	Partly trained: Queensland, Victoria and the Ausgrid area of New South Wales	Everywhere

Every pixel says whether local data trained the model where it lies (`training_tier` in the pixel table). In an anchored or partly trained place, the held-out-country rows of the accuracy tables below are the best guide to the pattern's quality.

How accurate is it

The test

Before training we hide whole groups of neighboring districts from the model. Each hidden district is then scored two ways, with the same information: the model, and a simple population split, each spread the official total of the next level up across the hidden districts, and both are compared with the districts' published figures. The population split is a high bar, not a straw man. It is given the right totals, so the only thing left to get wrong is the pattern inside them, and that is the only thing the scores measure. Every margin below is the model's gain over a map that already has the official totals right.

▾ How the test works in detail

Each hidden district's value is produced the way an unknown place's would be. In the published files a British local authority carries its official figure exactly; the Great Britain row below says how good the model is at placing electricity in areas that size when it cannot see their figures. Neighborhoods inside a local authority are less certain than that: see the neighborhood rows and the validation page.

Great Britain, metropolitan France, Canada and Australia have hidden areas. Every country with local data is also scored by a model trained without that country at all, with the totals one level up known or with the national total only. That is the closest guide to what a newly added country would get. Because each of those refits leaves out one country, the per-country rows each rest on a single model; the rows that pool every held-out country score far more areas and are the steadier guide today.

The two scores

Every number in the tables is an error, so smaller is better: a score of 9% means that areas typically miss their published figure by 9%, not that the map is 9% accurate. Two scores are shown, side by side, for the model and for the population split:

- **RMSE, the root-mean-square error**, in GWh per area per year. Each area's miss is squared, the squares are averaged, and the square root is taken, so the largest misses weigh the most. It is the score we rank by and quote first: the model's advantage is largest in the biggest, most industrial and most unusual

areas, which a percentage score ranks no differently from the rest. Read it against the average use of the areas in that row, which the tables also give: a 101 GWh miss on municipalities that average 260 GWh is a different thing from the same miss on cities that average 5,000 GWh.

- **mdAPE, the median absolute percentage error.** Half of the scored areas are within this percentage of their published figure, and half are further off. It is easy to read, which is why we show it, but a median says nothing about how far the other half miss, and a 10% miss on a village counts the same as a 10% miss on a city. Two methods can share nearly the same mdAPE and have very different RMSE, as the Australian and the held-out-country rows show.

The model is not fitted to either score (training minimizes a different loss), so neither is a verdict on its own; where the population split is lower on a score, the tables say so. In every row both methods are given the same known totals. For the total sector the population split keeps the large facilities in place and splits the rest by population; for residential it is a plain population split.

The numbers

GWh figures are per area per year; "Average use" is the mean published figure of the scored areas.

Hidden areas, scored against the published figures of districts the model never saw:

Country and sector	Scale	Areas	Average use (GWh)	RMSE, model (GWh)	RMSE, population split (GWh)	mdAPE, model	mdAPE, population split
Great Britain, France and Australia together, total	Municipality (median 390 km²)	218	260	101	175	14%	18%
All four countries, residential	Neighborhood (median 1.4 km²)	9,945	4.6	1.9	2.7	9%	18%
Great Britain, total	Local authority district	36	954	263	483	9%	14%
France, total	Intercommunality (EPCI)	119	215	38	48	9%	13%
Australia, total	Local government area	63	112	52	53	28%	27%
Australia, residential	Local government area	253	46	4.4	8.7	12%	19%

A model trained without the country, scored on every area of that country. The first four rows pool every held-out country and territory; the rows under them are one country at a time:

Country and sector	Scale and footing	Areas	Average use (GWh)	RMSE, model (GWh)	RMSE, population split (GWh)	mdAPE, model	mdAPE, population split
All held-out countries together, total	Municipality (median 390 km²), totals one level up known	1,708	430	123	165	13%	16%
All held-out countries together, total	Municipality, national total only	1,835	406	160	288	16%	18%
All held-out countries together, residential	Municipality, national total only	1,713	152	37	151	8%	15%
All held-out countries together, residential	Neighborhood (median 1.8 km²), totals one level up known	92,074	3.7	1.6	2.0	14%	17%
Great Britain, total	Local authority district, totals one level up known	359	718	143	209	9%	11%
Great Britain, total	Local authority district, national total only	359	718	226	241	12%	14%
France, total	Intercommunality, national total only	1,220	332	134	315	16%	18%
Australia, total	Local government area, national total only	129	512	204	287	23%	23%

Reading the tables

Read the RMSE columns first, and the pooled rows before the per-country rows. On hidden municipalities, pooled across Great Britain, France and Australia, the model's typical miss on a total is 101 GWh a year against 175 GWh for a population split given the same totals, on municipalities that average 260 GWh: the model removes about two fifths of the typical error. On hidden neighborhoods the residential miss is 1.9 GWh against 2.7 GWh, on neighborhoods that average 4.6 GWh.

For a country the model has never seen, with only the national total to go on, the typical miss on a municipality's total is 160 GWh against 288 GWh for the population split, on municipalities that average 406 GWh; for the residential sector it is 37 GWh against 151 GWh, four times the model's miss. What the model buys is

far fewer large misses, in the industrial, rural and unusual places where a planning error costs most.

The mdAPE columns sit closer together (14% against 18% on hidden municipalities, 16% against 18% for a held-out country) because half of the areas are easy for any method and a median only sees that easy half. That is why RMSE, not mdAPE, is the score to compare methods on.

Each per-country row comes from a single refit without that country, and with four countries holding training data there are only four such refits, so a per-country figure can move when a country is added or a source changes; the pooled rows score every held-out area together and are the steadier guide. They will firm up country by country as more countries with local data join the training set.

Compare rows only within a scale and a test. Rows at different scales score areas of very different size, so their GWh are not comparable, and they are corrected to different known totals; the hidden-area and held-out-country tests score different sets of areas. Within one scale, bigger areas score better, which the validation page's area breakdown shows. Australia's total at the council scale is the weak case: its training areas are fewer and coarser, and on both scores a population split does as well or slightly better. The validation page shows every scale, sector and country, including the French overseas territories (national total only) and the other places where the population split scores lower, such as French non-residential at the département scale.

An outside check

On the hidden intercommunalities of France, scored on 2021 alone, the model's typical miss (RMSE) on a total is 45 GWh, against 113 GWh for the JRC EU Energy Atlas, a public product that shares national totals out with spatial proxies (115 GWh once the atlas is rescaled to the same total). By mdAPE, half of our areas are within 11% of the published figure, against 25% for the atlas (23% rescaled).

What is not validated

Large facilities (checked against grid-operator figures instead), the net basis and the rooftop-solar split, anything below the smallest unit an official figure exists for, South Korea (nothing finer than the district figures it is scaled to), and the hourly preview. Only the size of the error is measured, not its direction: whether dense places are systematically over- or under-allocated is not yet known.

What the pixels can and cannot tell you

- **Add pixels up.** A single pixel is a modeled estimate. Sum pixels over a neighborhood, a municipality, a service territory or a planning zone; the error falls as the area grows. Do not compare one pixel with one meter.
- **Official units are exact where the table above says so.** For those units the map's total is the published figure on its own basis, plus any large-site use those statistics leave out. Where a published figure only shapes the split (Ontario postal areas, Australian council areas), the map's total differs from it by a few percent. The region metadata file carries the published figure and the map's comparable sum side by side.
- **Annual energy, not peak.** The values are yearly consumption. Peak demand, what wires are sized for, needs a load shape; the experimental hourly preview is the only one we publish.
- **Year-to-year change follows the statistics.** The level of each region moves with the official figures. The pattern inside a region changes little between years, so a pixel's change is not evidence of local growth. Compare like with like: Northern Ireland is absent in 2024, and where a shaping source starts partway through a series (Ausgrid residential figures from 2019) a council's split can step that year.
- **Edges are approximate.** A cell that straddles a boundary belongs to one region; a cell on the coast may be partly sea. Over an area of tens of cells these effects are small.
- **Verify before operational use.** The data is in beta. Compare against whatever local figures you have before relying on it for a decision that matters.

Hourly profiles, experimental

Great Britain carries an hourly preview for 2024, shown on the map only: a typical day for each season and kind of day in every cell, on both bases, that sums back to the annual value. The shapes come from sampled smart-meter and standard load profiles and have not been validated. Treat them as a preview and check them against data you trust. Hourly downloads will follow once the product leaves beta.

Adding a country

The pipeline needs two things from a country: boundaries for the regions official consumption is published for, and the official figures themselves, ideally split into residential and everything else (a statistical office, the utility's sales, or distributor billing; if only a national energy balance exists, the country is labeled `A2` in the case codes below). The inputs the model observes come from global datasets and national census tables we obtain ourselves. Any consumption published for small areas (feeders, towns, postal areas) lets the model train locally and moves the country from anchored to partly trained. To request a country, email benm@rewiringamerica.org with the country and any sources you know of.

Known limitations

- All training data comes from places where nearly every building has a grid connection. The model has not been tested where many households are unconnected, and it does not know where the grid reaches.
- Rooftop solar is modeled in Great Britain, Australia and Ontario only, and is not validated anywhere.
- Data centers are not identified as facilities.
- In South Korea, electricity that industrial sites generate and use on their own premises, about 4% of the national total, is outside the utility's sales figures that set the map's level.
- Northern Ireland is covered for 2021 to 2023 only, with coarser official figures than the rest of the United Kingdom.

Appendix: case codes

Every pixel and region carries a case code such as `T3-A1`. `T` is the training data behind the cell, `A` is the official statistics the totals follow, and `+S` marks a finer official source that shapes the split within regions.

▾ Show the case codes

Code	Meaning	Applies to
`T3-A1`	Trained on local small-area data for every sector; totals from the country's own official statistics.	Great Britain, metropolitan France, the five overseas territories
`T2-A1+S`	Trained on local data for some sectors or regions; totals from the country's own official statistics; finer official data shapes the split within regions.	Canada (residential only), Australia
`T0-A1+S`	No local training data; the utility's own metered sales set the national level and each district's total; the model trained elsewhere draws the pattern inside each district.	South Korea

The codes `T0` to `T3` run from no local training data to a sector split at several granularities; `A1` is the country's own statistical office or utility and `A2` an international energy balance only (none in this release). The basis columns beside the code (`res_anchor_basis`, `non_res_anchor_basis`, `facility_basis`) are defined in the Data Guide.

Citing and license

The data is licensed CC BY 4.0. Every release has an id such as `2026.10.4` and stays published unchanged; pin the id in anything you publish, and cite the release you used:

Cite this release

Ben Moscona, Stephen Pantano, Saul Griffith (Rewiring the World), 2026. Downscaled 1 km Electricity Consumption, release 2026.10.4. CC BY 4.0. demandgrid.net

Copy citation

Data license: CC BY 4.0 (attribution required). Pin the release id in anything you publish.

The Data Guide lists the sources to attribute alongside the release: the official statistics behind each country, Climate TRACE and OpenStreetMap contributors.