

A global model of hourly space heating and cooling demand at multiple spatial scales

Received: 7 July 2022

Iain Staffell ¹✉, Stefan Pfenninger ² & Nathan Johnson¹

Accepted: 1 August 2023

Published online: 14 September 2023

 Check for updates

Accurate modelling of the weather's temporal and spatial impacts on building energy demand is critical to decarbonizing energy systems. Here we introduce a customizable model for hourly heating and cooling demand applicable globally at all spatial scales. We validate against demand from ~5,000 buildings and 43 regions across four continents. The model requires limited data inputs and shows better agreement with measured demand than existing models. We use it first to demonstrate that a 1 °C reduction in thermostat settings across all buildings could reduce Europe's gas consumption by 240 TWh yr⁻¹, approximately one-sixth of historical imports from Russia. Second, we show that service demand for cooling is increasing by up to 5% per year in some regions due to climate change, and 5 billion people experience >100 additional cooling degree days per year when compared with a generation ago. The model and underlying data are freely accessible to promote further research.

Buildings account for 30% of global final energy demand and CO₂ emissions, with heating responsible for around half^{1,2}. Although space cooling is less widely adopted than space heating, the warming climate, population growth centred on the tropics, and rising affluence mean that demand for space cooling is growing rapidly³. Air conditioners and electric fans account for 10% of global electricity demand⁴. Unless actions are taken to address the efficiency of cooling equipment, energy demand for space cooling could more than triple by 2050³⁻⁵. Accurate modelling of the spatial and temporal impacts of weather on energy supply and demand is critical to decarbonizing energy systems around the world. Heating and cooling demand depend on population density, weather, the building stock and occupant behaviours, which vary over space and time. Various tools are available for modelling spatial and temporal impacts of weather on the variability of renewable energy supply^{6,7}, yet there are few tools available to study heating and cooling demand with the same global scope and consistency. The most relevant existing tools⁸⁻¹⁰ have restricted geographical scope (only Europe or the United States), and some only consider heating, neglecting the growing demand for space cooling¹¹.

Short-term measures to reduce energy consumption have seen heightened political interest since Russia's 2022 invasion of Ukraine, which spurred many nations to cut their dependence on Russian energy

imports to improve energy security¹²⁻¹⁵. One proposal for reducing energy consumption is to lower thermostat set-point temperatures in buildings, yet little is known about how savings vary over space and time¹³. Reducing thermostat set points also has clear benefits for consumers (lower energy bills) and the climate (faster emission savings when compared with retrofit)^{4,13,16}. At the same time, more than 5 million deaths per year (nearly 10% of global mortality) are associated with excess cold (4.6 million) and heat (0.5 million)¹⁷. A successful global energy transition will depend in part on understanding where, when and how much energy is required to provide universal thermal comfort¹⁸.

This paper develops a generalized framework for modelling hourly space heating, cooling and total energy demand at national, regional and individual-building scales using openly available meteorological reanalysis data. It validates simulated demand profiles against metered electricity and gas demand data from 43 regions and ~5,000 individual buildings worldwide. It then explores two applications of this model: the impact of reducing thermostat set points on gas demand in individual buildings and at national scales, and the influence of changing regional climates on demand for cooling over the past 40 years. Finally, it launches an interactive front end for accessing these data, which allows users to perform custom simulations of heating, cooling and total demand anywhere in the world.

¹Centre for Environmental Policy, Imperial College London, London, UK. ²Faculty of Technology, Policy and Management, Delft University of Technology, Delft, The Netherlands. ✉e-mail: i.staffell@imperial.ac.uk

Modelling building energy demand

The most common approaches for modelling energy demand in buildings can be divided into four categories: physical, statistical, surrogate and degree-day models. Supplementary Table 1 compares each approach and our own model—the Demand.ninja—against four key criteria (see also Supplementary Notes 1–4).

Physical (white-box) models use energy balance equations and hundreds of input parameters to simulate a building's energy demand. Such a detailed approach is invaluable for optimizing building performance during design stages or when retrofitting. However, programming physical models is time consuming and data intensive, and requires considerable knowledge of the tool, which makes these models difficult to use and to scale beyond individual buildings. A growing body of evidence shows that large discrepancies between simulated and measured energy demand can arise when these models are used predictively without calibration^{19–24}. Calibrating the output of physical models against metered energy demand data improves performance, as demonstrated by NREL's ResStock and ComStock models for example¹⁰. These models use novel sampling techniques to allow scaling beyond the individual-building level, with disaggregation by building type and end-use, which allows them to predict changes in energy use for different combinations of building upgrades. The main limitation is that they require extensive and detailed data for calibration to provide accurate results, hence calibration is limited to the United States. Extending this calibration to a global scope would require equally detailed data for other regions that are currently unavailable.

Purely statistical (black-box) models derive mathematical relationships between measurements of energy demand and other variables, which are used to predict demand^{25–28}. These models can provide fast and accurate results but are only valid within the scope of their training data, and thus are difficult to generalize or customize to fit other contexts. Surrogate models are statistical models trained using input and output data from physical models^{29,30}. With a sufficiently large and generalized training set, a single surrogate model could in theory (though not yet in practice) be used to quickly simulate energy demand across an exhaustive array of building designs and climate zones^{29,31}. Such a model would be both generalizable and customizable but would still require extensive input data and knowledge to be used correctly. If set up incorrectly and without calibration, surrogate models would suffer from the same potential inaccuracies as the physical models upon which they are based^{29,32}.

Degree-day models provide a simpler approach for modelling energy demand, but can only disaggregate demand into total heating and cooling, rather than multiple individual end-uses³³. Space heating and cooling demand are strongly correlated with outdoor temperature, which is openly available at high resolutions (for example hourly in 0.5° grid cells). Outdoor temperatures are translated into degree days, which measure by how much (in °C) and for how long (in days) the outdoor air temperature is higher or lower than a balance point temperature—the range of temperatures at which a building's internal heat gains counterbalance external losses, such that it requires neither heating nor cooling³³. Degree days are then used to derive synthetic heating and cooling demand profiles by calibrating the model using measured energy demand data. The Hotmaps Project⁹ estimates daily heating and cooling demand across 28 European countries using temperature as the only meteorological input, and When2Heat⁸ estimates hourly heating profiles for 16 European countries using temperature and wind speeds. However, the most basic degree-day models “seem to represent the observed data equally well”³⁴ as When2Heat, despite its additional processing steps. While degree-day models have broad applications and low data requirements, they have limited scope for customization (beyond temperature thresholds) and are not widely validated—only at national level in two to four countries, or against individual buildings within a single country (which is critical for generalizing the model globally)^{8,9}.

Here, we present the Demand.ninja, a grey-box model of final energy demand (that is, the amount of energy as gas or electricity that is consumed in buildings) for space heating and cooling. It derives from a degree-day model but incorporates additional weather variables and statistical parameters that account for interactions between buildings and the local climate to improve the model's ability to reproduce measured demand. We calibrate the model and validate its performance by simulating total electricity and gas demand at national, regional and building levels across the world, and develop representative but customizable parameters that allow the model to be used in other contexts. The model's workflow is divided into nine stages, illustrated in Fig. 1a, and described in Methods.

First, we calculate the building-adjusted internal temperature (BAIT) at each location (step 1 in Fig. 1a). BAIT builds on existing thermal indices to represent the temperature that a person would feel inside a building that is neither heated nor cooled (Supplementary Note 5). It is derived from daily average external air temperature, solar irradiance, wind speed and humidity. It uses several fitted coefficients to parameterize how the internal temperature of a building is influenced by the local climate (that is, solar gains, ventilation and humidity). Figure 1b shows the relationship between BAIT and raw temperature in Great Britain, as a function of solar irradiance and wind speed. BAIT is then smoothed over the previous two days to account for the thermal inertia of insulated buildings (step 2). These adjustments explain how, at the same external temperature, the internal temperature of a well insulated building with large south-facing windows will be greater than that of a poorly insulated, draughty building with small windows. When it is hot and cooling is required, raw, unsmoothed temperatures are blended back into BAIT (step 3) to account for temporary ventilation (for example, windows being opened at night), which flushes heat accumulated in the building from the previous day, reflecting the fact that raw temperature is a better predictor of cooling demand³³.

Heating and cooling degree days (HDDs and CDDs) are then calculated using BAIT and balance-point temperature thresholds derived from it (step 4). If generating aggregated national or regional data, HDDs and CDDs in each grid cell are aggregated across several calculations using a range of input parameters to reflect the diversity across buildings and occupants, and the resulting degree days are then averaged across grid cells, weighted according to population density to account for higher demand in more populous regions (step 5). Next, HDDs and CDDs are translated into final energy demand using three coefficients (step 6). The baseline power coefficient measures temperature-independent energy demand, and heating and cooling power coefficients measure the additional demand per HDD/CDD per building (at the building level) or per person (at the national or regional scale). They account for differences in the size of buildings, the efficiency of building envelopes and heating and cooling systems. Large, poorly insulated buildings with inefficient equipment will yield higher power coefficients.

Power coefficients, balance-point temperature thresholds, and coefficients for building–weather interactions can all be specified by a user for exploratory analysis. We also optimize these at specific locations by calibrating the model using measured electricity and/or gas demand data at the national, regional and individual-building levels (steps 7 and 8). The optimal parameters are averaged across regions to give generic values for each parameter to allow the model to be applied globally.

Finally, as the above stages are computed at daily resolution, we derive globally representative diurnal profiles of heating and cooling demand (Extended Data Fig. 1), which are applied to give hourly time series of demand for regions and countries (step 9).

Figure 2 illustrates the process of modelling hourly electricity demand in New York State, United States. When calculated using optimal threshold temperatures, New York has almost twice as many HDDs (980)

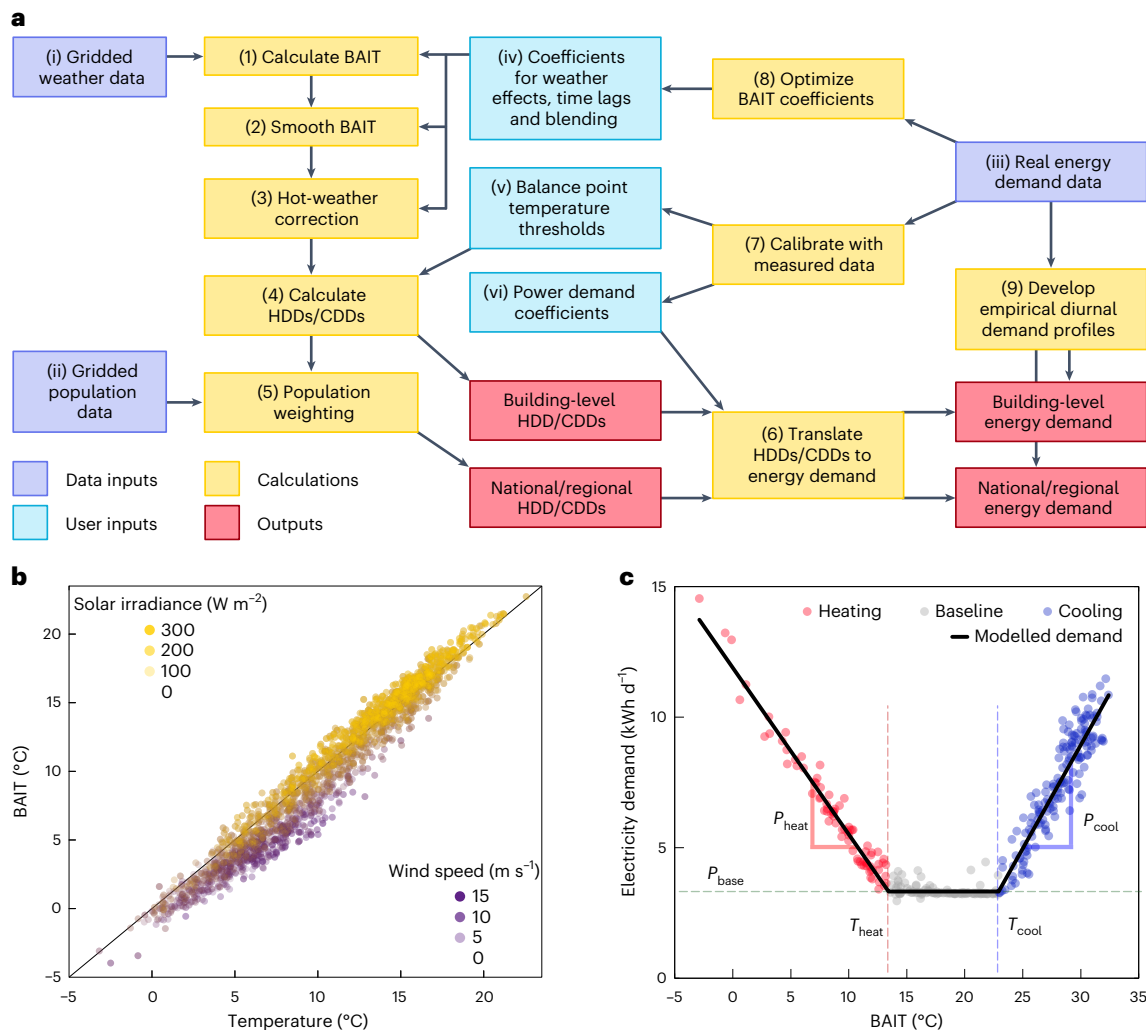


Fig. 1 | Illustration of the Demand.ninja method. a, The nine stages of the model's workflow. Raw data inputs (labelled i–iii) are shown in blue, derived or user-specified inputs (labelled iv–vi) in turquoise, calculation stages (labelled 1–9) in yellow and outputs from the model in red. **b**, How solar irradiance and wind speed affect the relationship between raw temperature and BAIT in Great Britain from 2015 to 2019, relating to step (1) and input (iv) in the model's workflow. Higher wind speeds reduce BAIT below air temperature, and higher solar irradiance increases BAIT above air temperature. A diagonal line with a 1:1 relationship is shown for reference. **c**, The relationship between BAIT and electricity demand using an exemplary home in Austin, Texas, United States,

relating to steps (6) and (7) and inputs (v) and (vi) in the workflow. Individual data points show daily-average metered demand in the home, with colours relating to how they are classified by the model. The black line shows the Demand.ninja regression (that is, modelled demand) with change points for when heating and cooling are required ($n = 364$). The lines labelled T_{heat} and T_{cool} refer to balance point temperature thresholds ($^{\circ}\text{C}$), line P_{base} refers to baseline or temperature-independent energy demand (kWh) and slopes P_{heat} and P_{cool} are the coefficients for heating and cooling demand ($\text{kWh d}^{-1} \text{ } ^{\circ}\text{C}^{-1}$), which show how much demand increases with falling/rising BAIT.

as CDDs (530) per year; however, heating demand is met only partially using electricity (13% of households), whereas all cooling demand is met using electricity³⁵, so hot days have a far greater impact on electricity demand than cold days. This is seen in Fig. 2e, where the slope on the right is steeper than that on the left. Figure 2f–h shows that, using only weather data and five optimized parameters, the Demand.ninja models the state's total electricity demand accurately.

Validation against historical data and other models

We validate the Demand.ninja at the national and regional levels using energy demand data covering 27 countries (43 regions), and at the individual-building level using smart-meter datasets covering 4,773 buildings across six countries (Supplementary Tables 2 and 3). Optimal model parameters and performance for each dataset are summarized in Supplementary Table 4 and Supplementary Figs. 43 and 44.

Figure 3 shows the model's ability to represent daily natural gas demand in the United Kingdom, and how this improves as elements of the BAIT process are incorporated. Applying this process reduces the root mean squared error (RMSE) by three-quarters relative to a basic degree-day model, allowing national-scale gas demand to be modelled to within ± 2.7 GW (2.1%). The largest improvements are a 33% reduction in RMSE by moving from national average to gridded temperature inputs (Fig. 3b relative to Fig. 3a), another 33% by accounting for solar gains (Fig. 3f relative to Fig. 3e) and 16% by accounting for diversity in buildings and occupants through aggregating multiple simulations (Fig. 3g relative to Fig. 3f). The latter two are not included in previous studies^{8–10}.

Figure 4 explores the error on modelled daily energy demand from the Demand.ninja. Figure 4a,c shows that model error is stable over time in both countries/regions and individual buildings respectively. At the national/regional level, the average normalized root mean squared error (NRMSE) was 1.9% across Europe ($R^2 = 0.93$) and

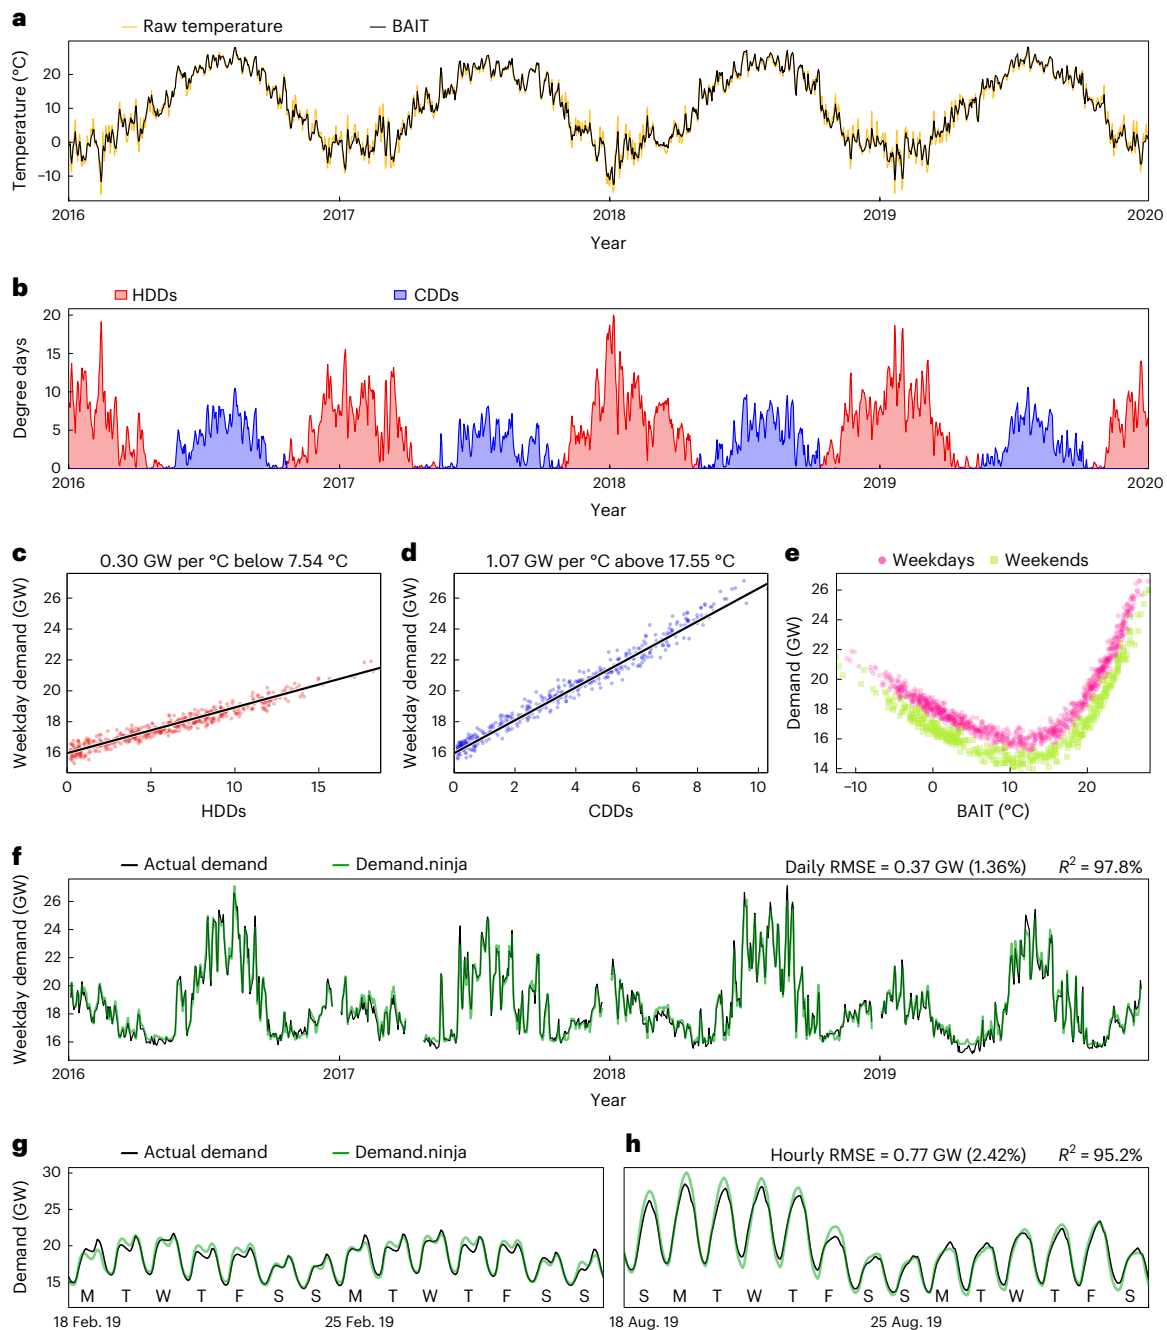


Fig. 2 | The process of converting from temperature to BAIT, to degree days, to energy demand. Measured electricity demand covering all end uses in New York State (NYISO) from ref. 77 is used as a case study. Supplementary Figs. 1–42 give plots for all other regions. **a**, The daily temperature averaged across the state (weighted by population density) and the corresponding BAIT. BAIT follows temperature but has lower daily variability due to temporal smoothing, and is generally higher than temperature during spring and lower in autumn months due to differences in solar irradiance relative to temperature. **b**, The corresponding HDDs and CDDs using the optimal temperature thresholds derived from metered electricity demand. **c,d**, The relationship between these degree days and daily metered electricity demand, with points showing demand from individual weekdays and lines showing the derived linear regressions ($n = 1,392$) used to model energy demand. Panel headings give the threshold

temperatures and power coefficients for heating and cooling. **e**, The relationship between metered daily electricity demand and BAIT, with the tightness of the fit indicating the model's capability to predict demand. **f**, Comparison of historical daily electricity demand with the model estimate covering the entire span of metered data. Demand is measured in terms of average power demand through the day in gigawatts. Gaps are shown where metered data are either missing or unrepresentative due to national holidays (for example, Christmas). **c,d,f** focus only on working days to remove noise from demand being lower on weekends and holidays for socioeconomic reasons. **g,h**, Comparison of historical and modelled demand at hourly resolution over two fortnights covering winter (**g**) and summer (**h**) months. Inset letters denote days of the week. Legends in **f** and **h** give statistical measures of the fit quality covering the entire multi-year period.

2.1% across the United States, Australia and Japan ($R^2 = 0.94$). Across all regions, estimated demand for nine days out of ten is within $\pm 2.5\%$ of the true value. Model error is similarly stable, but much larger when simulating demand from individual buildings. NRMSE rises to 7.6%

($R^2 = 0.82$), with 90% of observations within $\pm 8.1\%$. Buildings are more difficult to model due to highly individual usage and behavioural patterns (for example people going on holiday or controlling their heating and cooling system in a non-uniform way across years). Modelling

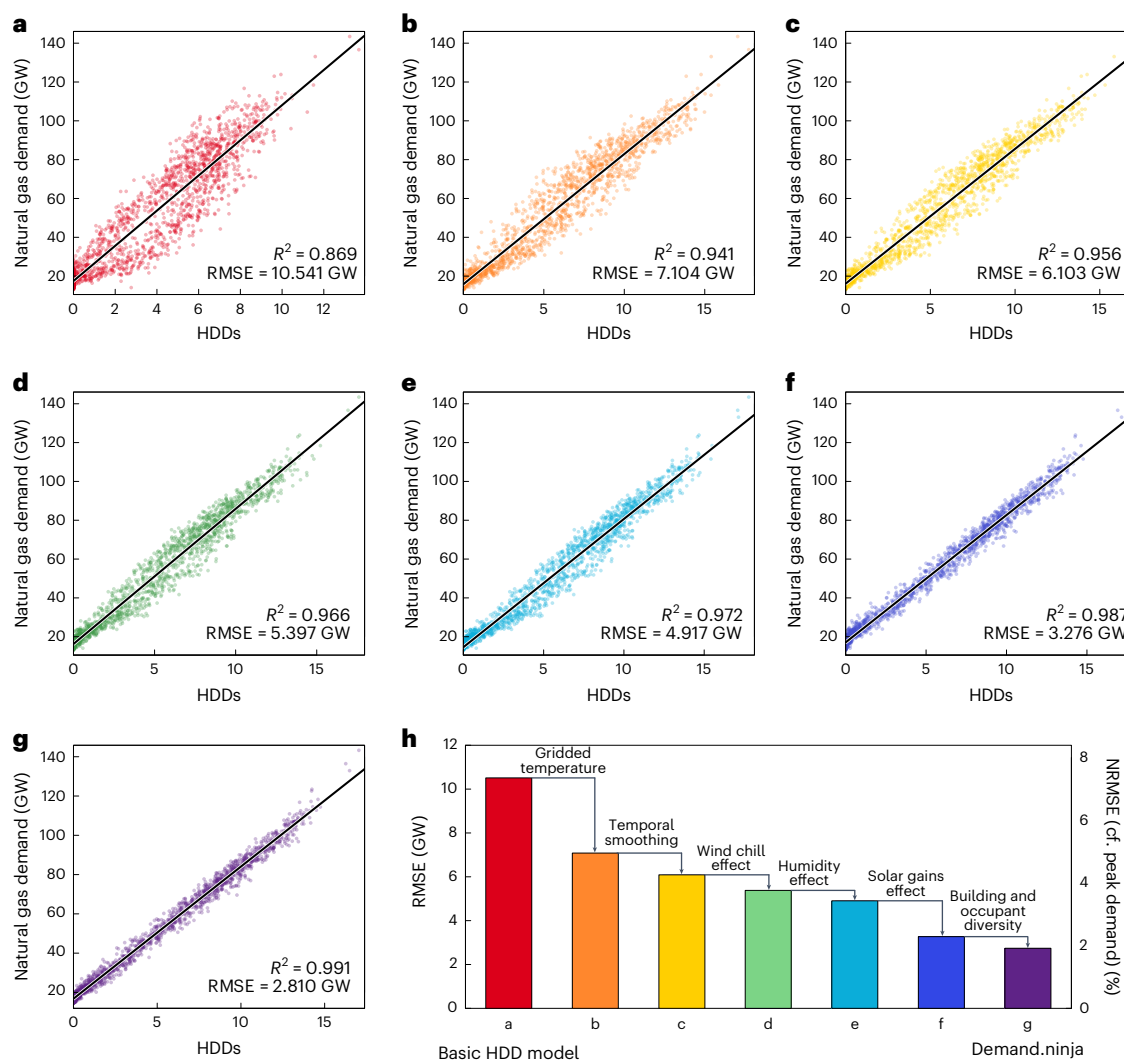


Fig. 3 | Ability of the Demand.ninja to model UK natural gas demand at daily resolution. **a–g**, The correlation between modelled HDDs and measured gas demand covering all end uses excluding power generation from ref. 78 over the period 2015–2019. Points show demand on individual days and lines show the derived linear regressions ($n = 1,827$). Each panel shows incrementally more sophisticated models, starting with the most common elements and moving to the most novel. **a**, A basic degree-day model using the national average temperature. **b**, Degree days calculated using gridded temperatures that have been population weighted. **c**, The addition of smoothing temperatures over the preceding two days. **d**, The addition of wind chill (higher wind speeds reducing the temperature index). **e**, The addition of humidity effect (greater humidity

reducing the temperature index when it is cold or increasing it when it is hot—note that this has a more influential impact in hotter climates). **f**, The addition of solar gains (greater irradiance increasing the temperature index). **g**, The superposition of multiple simulations with stochastically varied parameters to reflect the diversity of building construction and occupant behaviour. **h**, Summary of the improvement in modelling quality when adding each element of the Demand.ninja process. Each panel uses the optimal heating temperature threshold for the given model, so that the improvements shown relate to which elements of the BAIT process are included/excluded, rather than miscalibration of the model. Supplementary Fig. 45 shows the summary of model improvements (**h**) for electricity demand in four regions for comparison.

aggregate regions benefits from the diversity effect, where random uncorrelated errors in individual buildings are (partially) cancelled out when modelling a large population, as seen for example with modelling the output of wind farms³⁶.

Figure 4b,d shows that model error is also stable across the spectrum of external temperatures, implying that Demand.ninja has comparable performance across heating, cooling, and temperate conditions when day–night temperature swings give two-way temperature flow between the building and its environment. Aggregated across all building datasets, the model's relative error is larger for peak heating needs (when it is coldest) than for peak cooling needs. This is possibly an artefact of the heterogeneous building data used: a greater percentage of profiles with cooling demand were from larger commercial buildings, which are modelled with greater accuracy than the small single-family households that formed the vast majority of heating profiles.

Figure 4e and Extended Data Fig. 2 summarize the error on daily electricity and gas demand across countries/regions, compared with other models in the literature (Methods). Across European countries, the errors on Hotmaps' demand estimates are 1.63 ± 0.35 times larger than for the Demand.ninja when it is calibrated to each region⁹, while When2Heat's errors are 1.70 ± 0.44 times larger¹¹. Similarly, across US electricity markets, NREL's ResStock and ComStock models give an error 1.65 ± 0.46 times larger than the calibrated Demand.ninja model¹⁰.

When the Demand.ninja is used with generic inputs (the same 'global average' temperature thresholds and BAIT coefficients drawn from our training data applied to every region), these improvements are 1.46 ± 0.32 relative to Hotmaps, 1.47 ± 0.37 relative to When2Heat and 1.31 ± 0.39 relative to NREL. Across the 53 regions considered, the generic Demand.ninja model had lower RMSE than other models in all except one region (1.84% versus 1.82% in Texas).

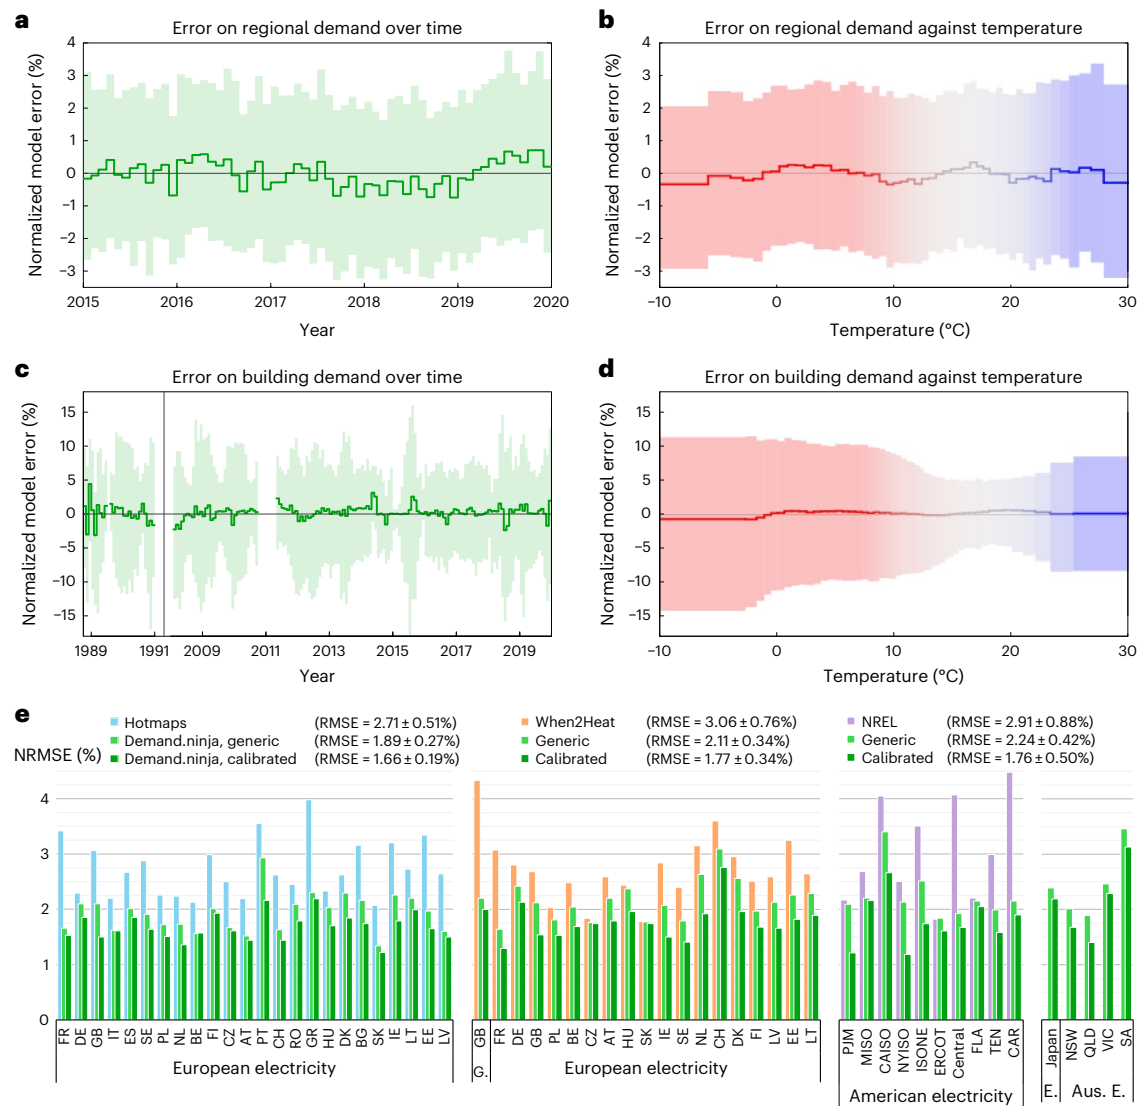


Fig. 4 | Quality of modelling daily energy demand. **a, b**, Summary of the error (Demand.ninja minus metered demand, normalized by maximum demand within each region) across 65,000 observations of daily electricity demand from 43 countries and regions. **a**, The error over time: the central line shows the median (P_{50}) of all days in each month across all regions; the shaded area shows the 5th and 95th percentiles ($P_5 - P_{95}$). **b**, Summary of the error as a function of outdoor air temperature, showing the same percentiles as in **a**. Observations are grouped into 50 bins with equal numbers of observations; hence, they have variable widths. Bins are coloured according to the average number of HDDs (red) and CDDs (blue) in the observations they contain, giving an indication of balance point temperatures. See Supplementary Figs. 46–88 for individual plots of each region. **c, d**, Summary of the error across 1.4 million observations of daily electricity and gas demand from 4,773 buildings, as a function of time (**c**) and external temperature (**d**). The error shows less temporal consistency in **c** than in **a**, as **c** comprises several datasets with different spatial and temporal coverages and building characteristics (for

example, residential versus commercial), whereas the regional data in **a** have broadly the same time spans and identical scopes. See Supplementary Figs. 89–99 for individual plots of each dataset. **e**, Summary of the NRMSE between measured demand and simulations from Demand.ninja and other models. Demand.ninja is shown with temperature thresholds and BAIT parameters calibrated for each region, and with generic parameters that are common across all regions (specified in Methods). The time frame varies between models (2010 for Hotmaps⁹, 2016–2019 for When2Heat¹¹ and 2018 for NREL¹⁰), so Demand.ninja was compared with each separately. European countries are represented by their ISO2 code and the United States is divided into power markets and Australia into states, each ordered by annual energy demand. All bars relate to electricity demand (shortened to E.) except one, which is for gas demand (shortened to G.). Extended Data Fig. 2 visualizes the relative improvement across datasets, and Supplementary Tables 5–8 give the underlying data and expand all acronyms used. Supplementary Figures 100–153 show the time-series comparison in each region.

Comparing the performance of the Demand.ninja model with different levels of customization reveals the relative impact of model design and calibration. Taking the example of When2Heat, using Demand.ninja with generic parameters reduces RMSE by 0.95 percentage points (from 3.06% to 2.11%), and the calibration of all parameters reduces this by a further 0.34 percentage points (to 1.77%). These reductions have a ratio of 74:26. Averaging across the three datasets, we find that model design provides 76% of the improvement. In a further trial (not shown in Fig. 4e), where temperature thresholds were calibrated for each region but generic BAIT parameters were retained, we find

that calibrated temperature thresholds provide 8% and calibrated BAIT parameters the remaining 16%.

One possible use-case for the Demand.ninja is to model energy demand for countries, regions or buildings where metered demand is unavailable, meaning that calibration of model parameters is not possible. To explore the legitimacy of this idea, Fig. 5 shows validation of the model's simulations when using only global generic parameters derived from our training dataset. Figure 5a–d shows simulated electricity demand in four countries from outside this training dataset. The average NRMSE across the four regions is 2.42% ($R^2 = 0.94$), only

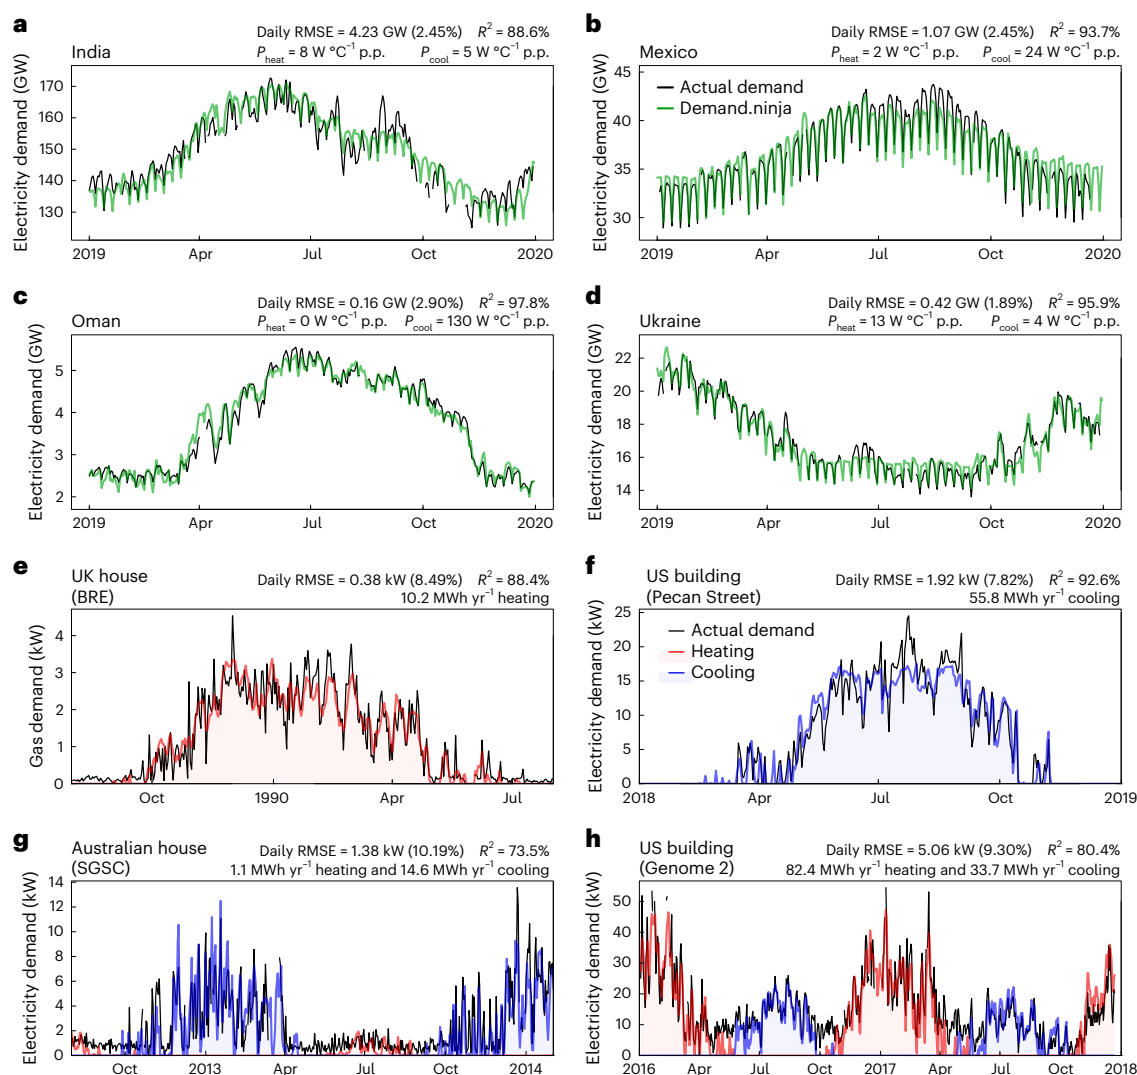


Fig. 5 | Validation of the Demand.ninja model using generic parameters without calibration. **a–d**, Simulations of national electricity demand for countries that were not included within the training dataset for calibration (that is, out-of-sample validation). **e–h**, Simulations of individual-building demand where the only specified input was the annual energy demand for heating and/or cooling. **a–d** are validated against national electricity demand data⁷⁹, for India (**a**), Mexico (**b**), Oman (**c**) and Ukraine (**d**), using the default BAIT parameters and temperature thresholds in all cases, so the only free variables were the three power coefficients used to establish the amount of non-thermal demand and scale heating and cooling degrees to demand. These parameters are given in the figure legends, represented per person (p.p.). **e–h** show simulations of randomly

selected individual-building demand from four datasets, spanning different decades and continents: Milton Keynes, United Kingdom (**e**), Texas, United States (**f**), New South Wales, Australia (**g**), and Washington DC, United States (**h**). The coloured lines and shaded areas show modelled heating and cooling demand. The default BAIT parameters and temperature thresholds were used and P_{base} was fixed at zero, giving an indication of model performance when given only the location of a building and its annual demand for heating and cooling, which are given in the legend. Supplementary Figs. 154–164 show simulations for a further ten randomly sampled buildings from each of the 11 building datasets used in this study. In all cases, national or building energy demand estimates could be improved if bespoke parameters were used.

marginally worse than the average across the regions within the training set. The results show that Demand.ninja performs consistently well without calibration across a variety of climatic zones.

Figure 5e–h shows simulated energy demand in four randomly selected buildings from our training dataset, where the only specified input was the annual demand for heating and/or cooling. Even with this very limited input, which could be gathered from building energy performance certificates or annual energy bills, the Demand.ninja with only generic parameters can accurately simulate daily building energy demand across different continents.

Reduced gas demand from lowering building temperature

We model the impact of changing thermostat set-point temperatures on gas demand in individual buildings and across Europe and the United

States by decreasing heating threshold temperatures throughout the year. Figure 6 shows the modelled impact of reducing the thermostat set point by 1 °C on gas demand in an individual house and summarized across 2,062 UK houses³⁷.

Between October and May, the individual house consistently saves 0.26 kW (6.4 kWh d⁻¹) because the daily average BAIT is at least 1 °C below the household's heating threshold and thus the maximum gas saving is realized each day. In contrast, between June and September, the daily average BAIT falls below the heating threshold only sporadically, and so gas is consumed (and saved) only on isolated cold days. This house sits centrally among the wider population of houses, with an annual saving of 10% of total gas demand (1,866 kWh yr⁻¹). The percentage saving is normally distributed across the population of houses—most UK households are expected to save 8.1–11.1% (± 1 s.d.) of their annual gas demand by turning the thermostat down by 1 °C.

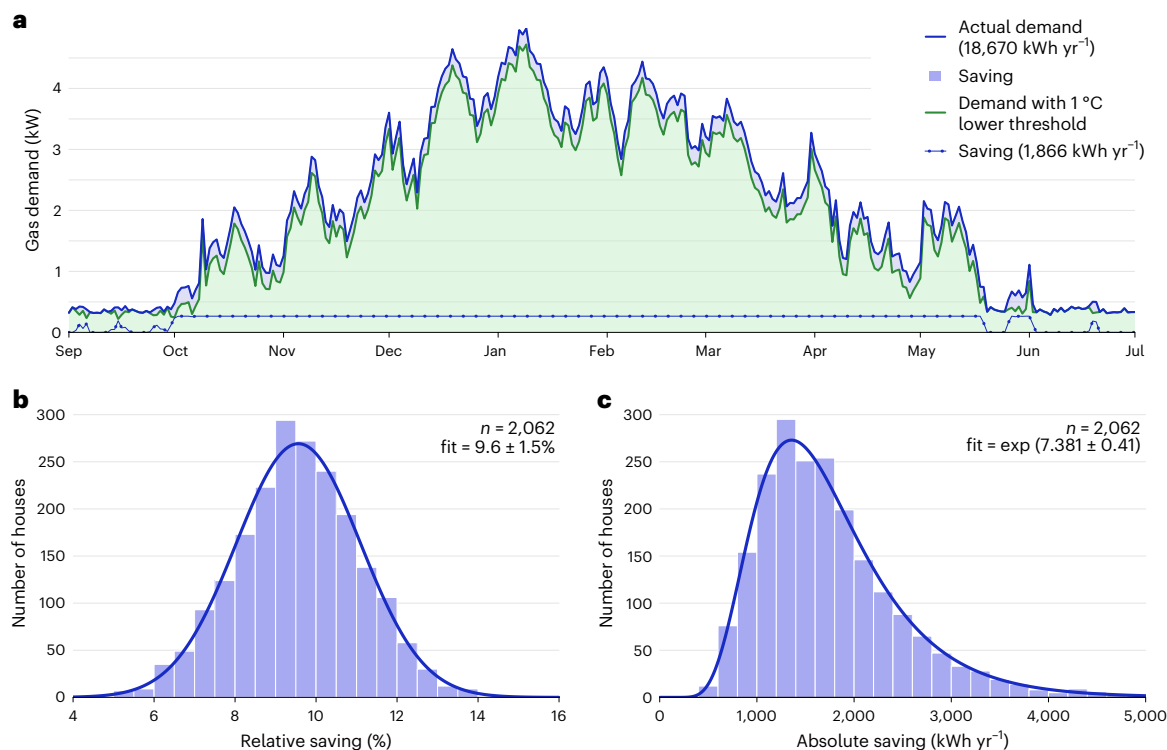


Fig. 6 | The reduction in natural gas demand from turning down the thermostat by 1 °C in individual houses. **a**, The change in daily gas demand over one year (2009–2010) in a randomly selected British house from the Energy Demand Research Project dataset³⁷. Measured demand is represented by the solid blue line. Modelled demand with a 1 °C lower heating threshold is represented by the green line and the shaded green area beneath it. The saving

is shown as the shaded blue area and the dotted blue line. **b, c**, Summaries of the gas saving across the whole population of houses with metered gas demand from that dataset, relative to total gas demand in each house (**b**) and in absolute terms (**c**). Bars show the frequency distribution of savings, and lines give the fits to a normal and log-normal distribution, respectively. See Supplementary Table 3 for a summary of the Energy Demand Research Project dataset.

Absolute savings instead follow a log-normal distribution in line with the underlying distribution of household energy demand (that is, an average house consumes 15 MWh of heat, and similar numbers consume less than 7.5 MWh and more than 30 MWh). The median saving is 1,605 kWh yr⁻¹, while the mean saving is higher at 1,748 kWh yr⁻¹ with the s.d. a factor of 1.51 ($e^{0.410}$) lower or higher than the mean (1,065–2,419 kWh yr⁻¹). With the 2022 residential retail tariff for gas of 11.4 p kWh⁻¹ (ref. 38), the average annual saving is equal to £199 per household.

Figure 7 shows modelled monthly gas demand in the United States and Europe, covering all end-uses (residential, commercial and industrial sectors), and excluding transformation (power generation and vehicle fuel), hereafter referred to as ‘total demand’. In both cases, the NRMSE between modelled and measured gas demand is 2.1–2.2% (Supplementary Figs. 165 and 166).

We estimate that lowering thermostat set points by 1 °C would achieve an annual saving of 240 TWh (817 billion cubic feet, bcf) of natural gas in Europe and 161 TWh (550 bcf) in the United States—4.4% and 3.4% of total demand respectively. This equates to an annual emission reduction of 49 MtCO₂ in Europe and 33 MtCO₂ in the United States (1.5% and 0.7% of energy-related CO₂ emissions respectively)³⁹. The wholesale price of gas averaged €93.00 MWh⁻¹ in Europe and \$16.21 MWh⁻¹ (\$4.75 per million British thermal units) in the United States over the heating season of 2021–2022 (Methods), and so these savings equate to €22 billion in Europe and \$2.6 billion in the United States. In the United States it is possible to assess the impact per sector, and we find annual savings of 101 TWh in residential buildings (7.1%).

The savings are relatively insensitive to BAIT, as they depend on the length of the heating season (how many days are below the heating threshold) rather than how severe or mild the peak of winter is. Despite year-to-year variations in gas demand (for example, compare 2012 and

2014 winter seasons in the United States), the annual saving varies by just ±9 TWh in Europe and ±4 TWh in the United States (±4% and 3%). When modelling a 2 °C reduction in the heating threshold temperature, the annual gas saving roughly doubles to 464 TWh (1,582 bcf) in Europe and 312 TWh (1,066 bcf) in the United States (8.6% and 6.7% of total demand).

The geographic split in per capita gas savings and their cost implications are shown in Fig. 7c,d. Annual per capita savings are 392 ± 237 kWh ($4.2 \pm 1.4\%$ of total demand) across European countries and 472 ± 226 kWh ($3.8 \pm 1.2\%$) across US states, with cost savings of €36.45 ± 22.05 and \$7.65 ± 3.65 respectively. The greatest absolute savings are in the United Kingdom and Germany (49 and 40 TWh) and in California and New York (15 and 12 TWh), while the greatest per capita savings are in the Netherlands (1,004 kWh) and Michigan (937 kWh).

Per capita gas and economic savings across regions are influenced largely by two factors: the length of the heating season and the heating power coefficient. More northerly regions with longer heating seasons generate larger savings with all else equal, because more days fall below the threshold temperature and are impacted by the intervention. Similarly, regions with higher heating power coefficients (that is regions with larger or less efficient buildings, or a larger share of gas-based heating) generate larger savings, because each unit of delivered heat avoided saves more gas. The latter factor explains why New Mexico with a heating season of 144 days but a heating power coefficient of 106 W °C⁻¹ per capita realizes a greater saving than Washington state with a heating season of 196 days but a heating power coefficient of 72 W °C⁻¹ per capita. Similarly, this explains why the United Kingdom (with a heating season of 190 days) realizes a greater saving than Finland (240 days), as its heating power coefficient is 110 W °C⁻¹ per capita, compared with just 35 W °C⁻¹ per capita.

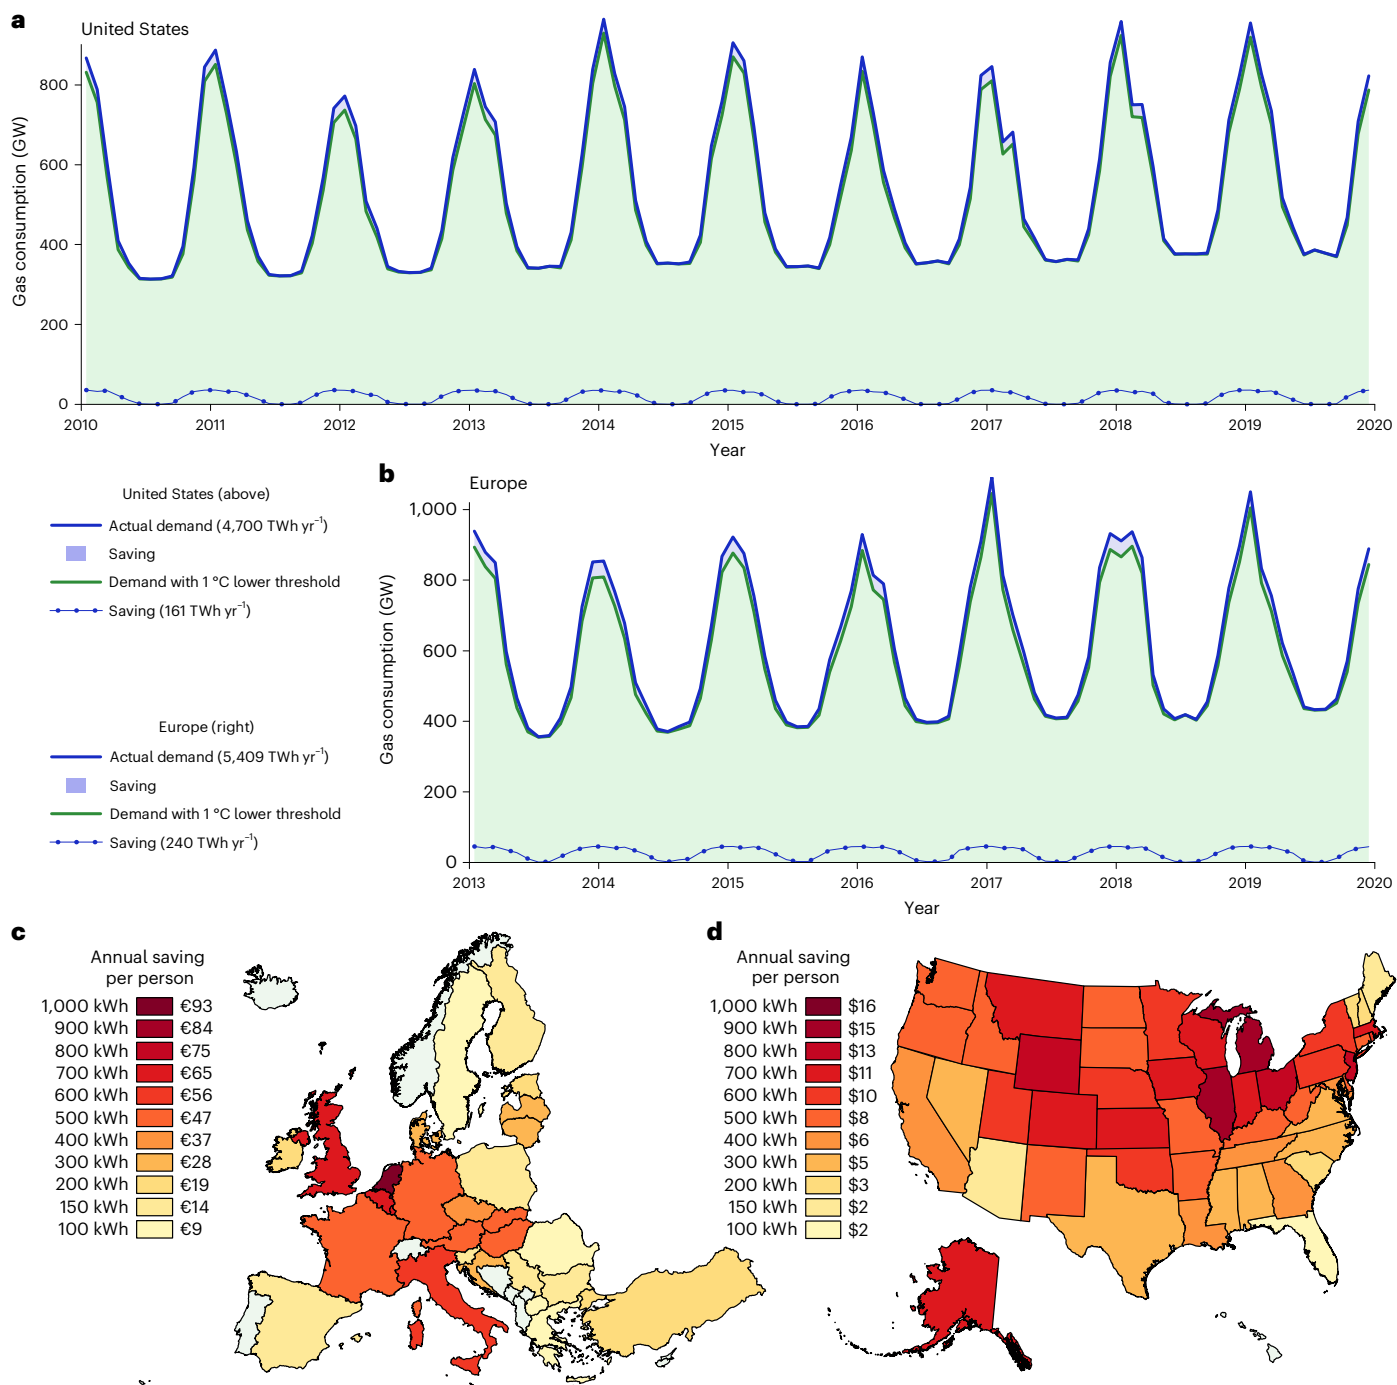


Fig. 7 | Reduction in natural gas demand across the United States and Europe from lowering all building temperatures by 1 °C. a, b, Total monthly natural gas demand and savings in the United States (a) and Europe (b). Measured demand is represented by the solid blue line and is taken from refs. 43,72. Modelled demand with a 1 °C lower heating threshold is represented by the green line and the shaded green area beneath it. The saving is shown as the shaded blue area and the dotted blue line. **c, d,** The regional distribution of per capita demand reductions in Europe (c) and the United States (d). Economic savings are given on

the secondary axis of each colour scale, calculated using the average wholesale gas price during winter 2021–2022. The European gas price is the average across Dutch TTF⁷⁴ and British NBP⁷⁵ and the US gas price is the Henry Hub price⁷⁶, both averaged across the period from 1 September 2021 to 31 March 2022. The United States is represented by the 48 contiguous states and Alaska; Europe is represented by the EU27, the United Kingdom and Turkey. Countries and states shown in light green had no data available for gas demand.

Increasing CDDs due to climate change

Extended Data Figure 3 shows maps of HDDs and CDDs across the world. Regions at high latitudes have the highest numbers of HDDs and the lowest numbers of CDDs, whereas equatorial regions have the opposite. Note that meteorological reanalysis data do not incorporate urban heat island effects, and thus modelled CDD values will be lower, and HDDs

higher, in highly urbanized regions. While cooling thresholds vary around the world (for example, they tend to be higher in low-income countries), this analysis used the model’s generic parameters to simulate CDDs in all regions (specified in Methods).

Figure 8 shows how CDDs have changed over the past 43 years. In almost all regions of the world (1,431 of the 1,518 countries, states and

provinces we consider) CDDs have increased between 1980 and 2022 due to rising temperatures. The absolute increase in CDDs is greatest in central Africa. For example, northern Angola now experiences an additional 686 CDDs per year on average when compared with 40 years ago (an increase of 172 per decade). This is largely because this equatorial region has high year-round temperatures, hence any increase in daily average temperature results in additional CDDs across more days. This rising need for cooling combined with limited capacity to meet existing—let alone new—cooling demand contributes towards this region being among the most vulnerable to impacts of future warming.

In contrast, the relative increase in CDDs is greatest in northern Europe because temperate regions experienced far fewer CDDs initially. For example, CDDs in London, United Kingdom, increased by an average of 5.0% per year (64% per decade) between 1980 and 2022, although the absolute increase was only five CDDs per decade. Nonetheless, small increases in CDDs across Europe, where only a small percentage of residential buildings are fitted with air conditioning (that is, existing demand for cooling is largely unmet), could have a disproportionate effect on electricity demand⁴⁰. Globally, we estimate that 5 billion people experience >25 additional annual CDDs each decade, and 1.3 billion people experience an increase of >50 each decade.

The impact of climate change on electricity demand for cooling can be assessed by combining the long-term increase in CDDs with cooling power coefficients for each region. By holding the cooling power coefficients constant, we isolate the effect of changing temperatures from sociotechnical drivers such as population, building stock, uptake of air conditioning, and system efficiency. The United States consumes 66 TWh more electricity per year for space cooling than it would have done with the weather of 40 years ago. This increase is greatest in Florida, where annual CDDs have risen from 1,320 to 1,590 since 1980, and the cooling power coefficient of 1.65 GW °C⁻¹ equates to 40 GWh per degree day, giving 10.7 TWh additional demand due to hotter summers. Across Europe, additional demand is 19 TWh per year (greatest in Italy, at 5.6 TWh). In Japan and Australia, additional demand is 9 TWh and 3 TWh respectively. Supplementary Table 9 summarizes the relevant parameters for all studied regions.

Discussion and conclusions

We simulate hourly energy demand in buildings using meteorological reanalysis data and validate simulated results against measured energy demand across 27 countries and ~5,000 buildings spanning four continents. The Demand.ninja can be applied globally and at specific locations, enabling users to simulate hourly energy demand in buildings worldwide. We show that it predicts energy demand more accurately (by a factor of 1.5) than similar but more geographically restricted models^{8–10}.

The model's main limitations centre around the availability of temporally resolved energy demand data for calibration. While data used here come from a wider range of countries than for similar models^{8–10}, there are still notable coverage gaps in South America and Africa. We require hourly or daily energy demand data at the national, regional or building level to expand calibration to new regions. As new data become available, simulations on the Demand.ninja platform can be updated. Energy demand data are also key for understanding how physical parameters of buildings (for example floor area, insulation and glazed area) translate to the high-level parameters the Demand.ninja model uses (that is BAIT and power coefficients) in individual buildings. For example, ref. 41 uses metered demand with rich metadata from 14,000 Danish households to correlate heating power coefficients with building area to infer average *U* values and other physical parameters. This would also be possible using the Demand.ninja if such data were available with global scope. However, users can employ the model's generic parameters, which simulate out-of-sample demand

accurately, or subjectively consider the range of parameters we find across large populations of diverse buildings to select those that best fit the building they seek to simulate.

We demonstrate two applications of the Demand.ninja by performing exploratory analyses. Prompted by the war in Ukraine, European nations sought to reduce their reliance on Russian gas imports^{12–15}. A substantial opportunity exists in lowering thermostat set-point temperatures. Unlike other measures for reducing gas demand (for example, accelerating the deployment of renewables and improving energy efficiency^{13,42}), thermostats can be adjusted instantly with no upfront financial costs, albeit at the expense of comfort levels and potentially health. We find that simply turning thermostats down by 1 °C would allow European countries to reduce gas consumption by 240 TWh yr⁻¹ (equivalent to one-sixth of historical imports from Russia) and the United States to reduce consumption by 161 TWh yr⁻¹ (equivalent to one-tenth of exports)^{13,43}. The emissions savings (49 and 33 MtCO₂ in Europe and the United States respectively) are comparable to emissions from the two largest coal power plants in Europe and the United States respectively, a combined 15 GW of capacity^{44,45}. With 2021–2022 gas prices, consumer energy bills could be reduced by €22 billion across Europe, and an average UK household would save around £200 annually by lowering building temperatures by 1 °C. Thermostat adjustments should therefore be considered by policy-makers as an important lever for addressing all three dimensions of the energy trilemma—security, affordability and sustainability.

Meanwhile, nearly half a million heat-related deaths occur annually, concentrated largely in Asia (46%) and Europe (37%)¹⁷, highlighting a substantial unmet demand for space cooling. We find that, in almost all regions of the world, CDDs are increasing as regional climates warm. This is not contingent on speculative future warming scenarios: it is evidence from historical weather data. In some regions, primarily in Europe, CDDs are increasing by up to 5% per annum. Globally, more than 5 billion people are experiencing >100 additional CDDs per year compared with just a generation ago (1980s versus 2020s). As living standards continue to rise, increased demand for cooling could combine with greater uptake of space cooling technologies, putting electricity systems under greater strain while increasing cooling-related emissions. These problems will only be exacerbated as heatwaves become more intense⁴⁶ and countries seek to reduce their energy consumption—factors that have already prompted curbs on air conditioning.

The broad applicability of the Demand.ninja complements tools for modelling spatial and temporal variability in energy supply such as www.renewables.ninja. There is high demand for spatially and temporally resolved energy demand data within the energy modelling community⁴⁷. However, commercial confidentiality and other factors often force researchers to produce their own simulations, requiring time and effort to duplicate work that has been undertaken elsewhere. The model and simulations developed here are therefore made available online for others to use freely via www.demand.ninja.

The Demand.ninja could increase research productivity in many different areas. By disaggregating historical demand into heating, cooling and non-thermal components, the impacts of heating electrification or uptake of air conditioning can be explored at the building, street, regional or national level. Hourly time series of demand from individual buildings can be synthesized using only annual demand and generic parameters (as in Fig. 5). Generating high-resolution time series of demand, including peak demand coincident across many buildings, could be used for system sizing to assess the need for infrastructure upgrades in local distribution networks, or capacity adequacy in national electricity systems. Drawing on several decades of historical data allows system stress to be explored during extreme winters or summers. When combined with contemporaneous simulations of renewable power production, this tool could allow 100% renewable energy

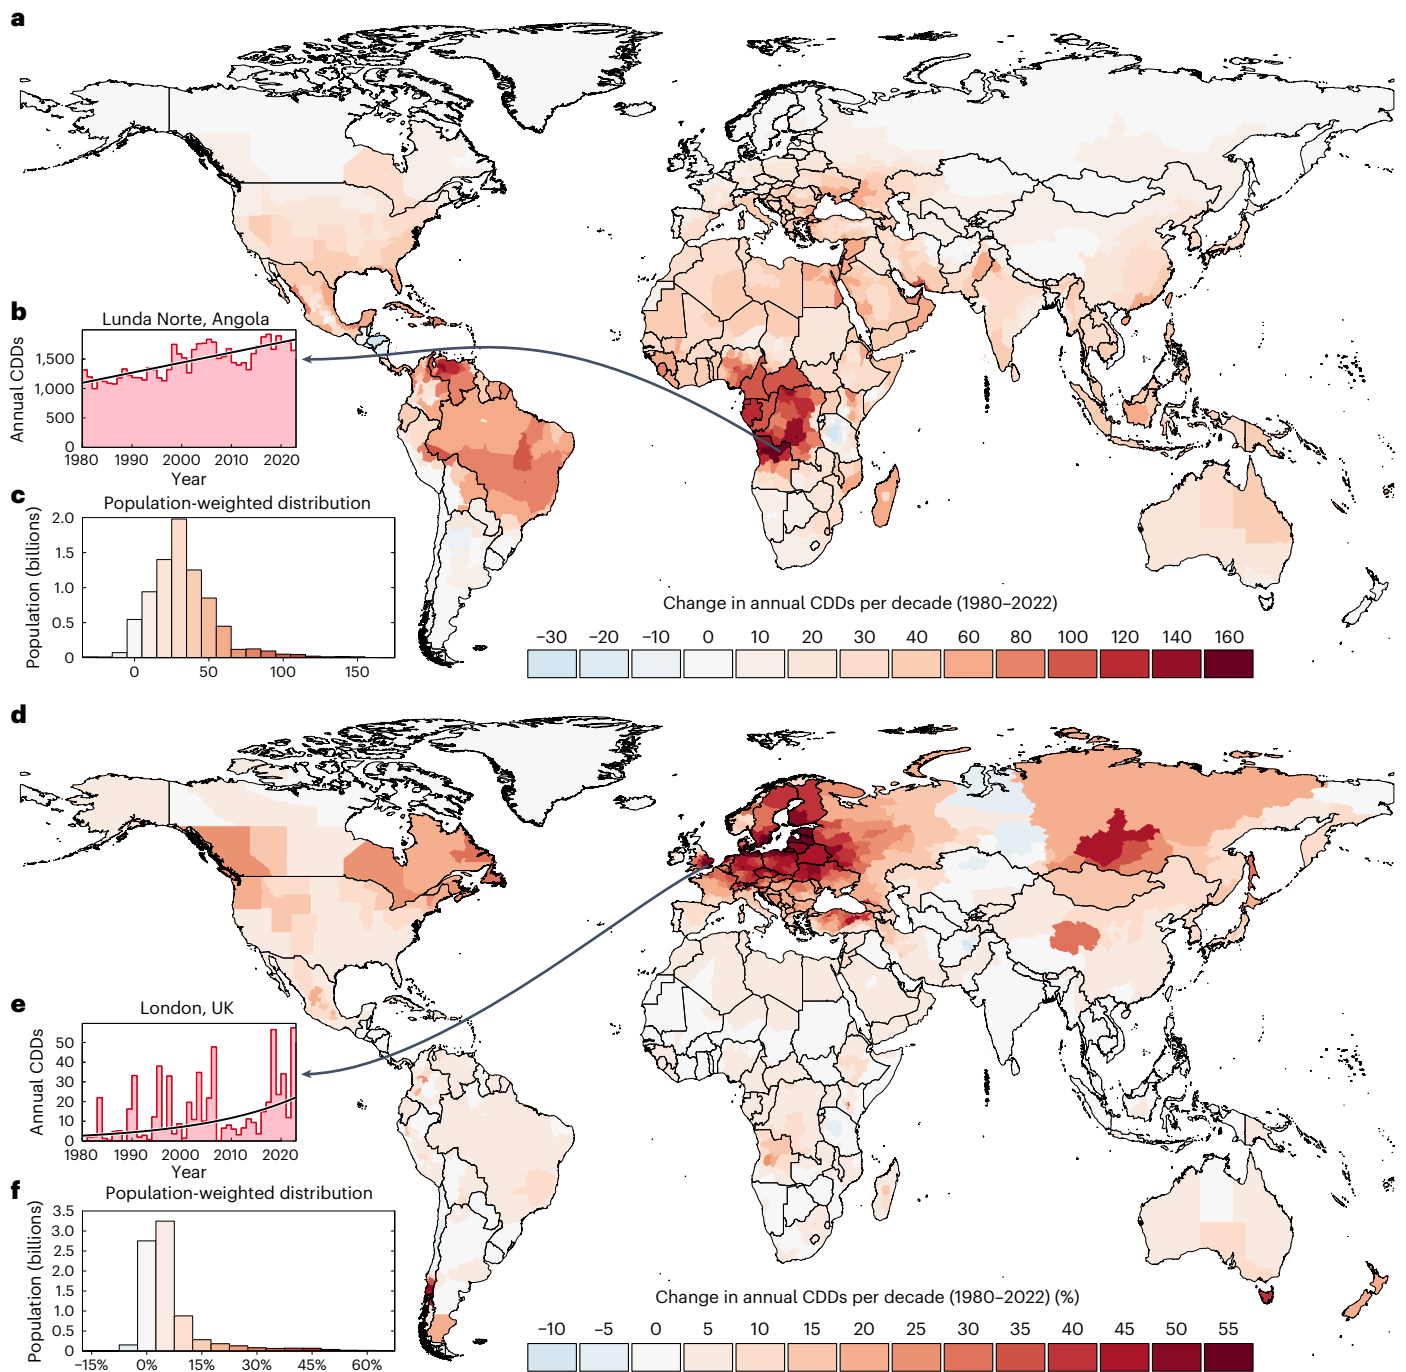


Fig. 8 | Global maps of the change in annual CDDs over the last four decades. All countries are modelled with a cooling threshold of 20 °C derived from our global calibration. **a**, The increase in absolute annual CDDs per decade, determined through linear regression of data from 1980 to 2022. **d**, The relative increase in CDDs per decade, determined through log-linear regression of data from 1980 to 2022. CDDs are evaluated at the level of administrative subregions (that is states, provinces) within larger countries (Methods). **b, e**, The 43 year time series of annual CDDs for two exemplary regions in Angola and the United Kingdom, which experience large absolute and relative increases respectively.

Black lines indicate the linear and log-linear fits to the data respectively. **c, f**, Population-weighted histograms of the increase in absolute and relative CDDs respectively, with shading that matches the colour scales used in the corresponding maps (**a, d**). These panels show how many people worldwide are exposed to different levels of CDD increase. Population exposure is evaluated at the level of administrative regions (for example, the 0.9 million citizens of Lunda Norte are modelled to all have experienced a 172 CDD increase per decade, or 738 additional CDDs per year in 2022 relative to 1980).

scenarios to be assessed with greater fidelity, help network operators plan infrastructure and capacity upgrades, or allow individuals to assess the level of self-consumption of building-integrated renewables. Adjusting weather inputs or temperature thresholds could allow users to explore the impacts of climate change on future demand, or the impacts of changing consumer preferences (for example, imposing

indoor temperature limits on municipal buildings). Equally, applications beyond energy could be explored, for example using cooling threshold temperatures to explore heatwave-induced excess mortality. We hope that by releasing the model open source we will help to overcome barriers to research in this area and enable new questions to be answered.

Methods

Data

Gridded meteorological data (box i in Fig. 1a) were sourced from MERRA-2⁴⁸, specifically temperature (T2M), solar radiation (SWGDN), wind speed (U2M and V2M) and specific humidity (QV2M). The newer ERA5 reanalysis offers higher spatial resolution but is not used because of an issue affecting Europe and the United States with mismatched wind speeds between model assimilation cycles (every 12 h)⁴⁹, which also appears to affect surface temperatures. Meteorological data are aggregated from hourly resolution to daily means for computational tractability and because intradiurnal variations in energy demand are strongly driven by user behaviour and occupancy patterns.

Gridded population data (box ii) were sourced from GPW⁵⁰ to calculate the population living within each MERRA-2 grid cell ($0.625^\circ \times 0.5^\circ$), and used to generate spatially aggregated HDDs, CDDs and demand. Other model inputs—weather coefficients (iv), threshold temperatures (v) and power coefficients (vi)—were derived by calibrating the model against real-world energy demand data (iii). These may also be customized to fit specific contexts or to perform exploratory analyses.

National and regional datasets of daily and monthly electricity and gas demand were sourced from as many regions as were readily available (Supplementary Table 2), and were used to calibrate and validate aggregated outputs of the model. These datasets cover all end-uses (for example appliances, lighting, industry, transport) as well as space heating and cooling. With electricity demand data, we ensured that ‘embedded’ renewable generation (for example residential solar photovoltaics) was not netted off total demand. For natural gas, consumption in power stations was specifically excluded, as this is transformation rather than end-use, and would introduce spurious noise based on seasonal variation in wind power output in some countries. Economic activity is typically lower on weekends and during national holidays, meaning that electricity and gas demand are lower, so we separate these from weekdays. Data were sourced for as many years as possible up to the end of 2019, so that energy-related impacts of the COVID-19 pandemic do not spuriously influence results⁵¹.

Electricity and gas demand data from individual buildings were used to calibrate and validate the model’s performance at the building level. These were sourced from the smart metering trials summarized in Supplementary Table 3, hence datasets were of varying size, quality and temporal resolution. As with national data, these were aggregated to daily resolution and weekends and holidays were separated. Buildings with less than 200 days of data were removed as they do not capture the full seasonal variation in demand. We then ran a change-point regression⁵² of demand in each building against temperature to remove those where energy demand was completely independent of the weather. This removed unconditioned buildings (for example, warehouses with only lighting and appliances), or buildings that were conditioned with a fuel not being metered (for example, electricity demand from a home with gas heating and no air conditioning). Buildings with heating but no cooling (or vice versa) are included, but report only one change point for the service that is used (see, for example, Fig. 5e,f).

As a compromise between obtaining a representative number of buildings and maintaining tractability, we typically discarded profiles with an R^2 coefficient below 0.6 in smaller datasets, or below 0.8 in datasets with over 10,000 buildings. These thresholds were chosen to exclude profiles that have a weak (or no) correlation with the weather due to ‘exogenous’ factors that cannot reasonably be expected to be captured. Such factors include buildings with non-standard occupancy patterns (for example, holiday homes or heating timers) or that suffered monitoring or data entry errors. Examples of discarded profiles are shown in Supplementary Fig. 167. Profiles that instead showed a weak correlation due to poor model skill were generally above these R^2 thresholds (and included in the analysis), but any such profiles that

had been excluded were manually identified and retained to provide a fair evaluation. Examples of such profiles can be seen in Supplementary Figs. 154, 156 and 157 (for example, $R^2 < 0.5$). Each of the remaining profiles was then screened, and anomalous points were removed (for example, sequences of zero values, which indicated holidays or failure of the data loggers), following the process in ref. 53.

The building-adjusted internal temperature

The BAIT is a proxy for the temperature it would ‘feel like’ inside a building with no active heating or cooling, accounting for solar gains, wind chill, humidity and thermal inertia. In other words, it is the temperature that an unconditioned building’s thermostat would register, which serves as the physical link between weather conditions and heating and cooling demand. The conception of this approach is discussed in Supplementary Note 5. BAIT should more adequately describe heating and cooling demand compared with temperature alone, as demonstrated in Supplementary Fig. 168.

Solar gains through glazed areas increase the internal temperature of a building, while higher wind speeds increase ventilation and the air change rate of a building, thereby decreasing the internal temperature of the building³³. The effect of humidity depends on air temperature. At high temperatures humid air prevents the body from sweating, which increases the feels-like temperature, while at low temperatures humid air deposits cold moisture on warm skin, increasing its thermal conductivity and thus decreasing the feels-like temperature⁵⁴.

We chose to incorporate these variables because they are used in existing human-related temperature indices and have been shown to improve the accuracy of physical and statistical models of building energy demand^{33,55,56}. Existing indices use a variety of mathematical forms to express deviations from outdoor temperatures (that is linear, quadratic and inverse relationships)^{54,57,58}. We use simple linear relationships to derive our index, as more complex forms were found to make little difference to BAIT but increase the risk of overfitting.

We derive BAIT in a way that is unbiased towards temperature so that measures derived from it (such as HDDs and CDDs, which rely on threshold temperatures) have magnitudes comparable to those derived from raw temperature. While a BAIT of zero will depend on how sunny, windy and humid it is, it should generally correspond to an air temperature of zero and not be systematically higher or lower (as seen in Fig. 1b). If BAIT were modified by absolute values of solar irradiance, wind speed and humidity it would yield lower temperatures during winter and higher temperatures during summer, hence the threshold temperatures for heating and cooling become skewed. To eliminate this bias, we calculate the global average solar radiation, wind speed and humidity as a function of temperature to be used as counterfactuals. We perform a combined linear regression across all locations considered in this study, correlating solar radiation, wind speed and humidity against temperature. Results from all regressions are shown in Supplementary Fig. 169. The equations derived from these regressions are simplified to give

$$\text{BAIT}_{d,l} = T + x(S - S^*) - y(W - W^*) + z(H - H^*)(T - T^*) \quad (1)$$

where for each day (d) and grid location (l) BAIT is in $^\circ\text{C}$, T is outdoor air temperature ($^\circ\text{C}$), S is solar radiation (W m^{-2}), W is wind speed (m s^{-1}), H is relative humidity (g water per kg air), x , y and z are coefficients for solar radiation ($^\circ\text{C W}^{-1}\text{m}^2$), wind speed ($^\circ\text{C m}^{-1}\text{s}$) and relative humidity ($^\circ\text{C g}^{-1}\text{kg}$) respectively and S^* , W^* , H^* and T^* are the counterfactuals for air temperature, solar radiation, wind speed and humidity respectively.

Counterfactuals are expressed as linear or exponential functions in relation to temperature (equation (2)). The counterfactual temperature is the point at which humidity has no influence on energy demand and is given the value of an assumed balance point temperature, where neither heating nor cooling is required.

$$S^* = 100 + 7T, W^* = 4.5 - 0.025T, H^* = e^{1.1+0.06T}, T^* = 16. \quad (2)$$

Buildings also store heat within their thermal envelopes, providing an inertia to fluctuations in the outdoor temperature. The thermal inertia of buildings is related to their level of insulation: well insulated buildings retain heat for longer than poorly insulated buildings, so have greater inertia. In much of the degree-day literature the effects of thermal inertia are ignored^{59,60}; however, recent studies^{8,61} and industry analyses^{62,63} have employed temperature ‘smoothing’ to account for the thermal inertia of buildings. Here, we smooth BAIT over the preceding two days using

$$BAIT_{d,l} = \frac{BAIT_{d,l} + \sigma BAIT_{d-1,l} + \sigma^2 BAIT_{d-2,l}}{1 + \sigma + \sigma^2} \quad (3)$$

where for each d and l , σ is the smoothing factor, the weight applied to the previous day’s temperature, in the range of [0–1].

BAIT and smoothing coefficients are bespoke to a given building, and can be defined by the user for exploratory analysis, or found through calibration to match a specific dataset of energy demand (Calibrating the model). Including wind speed, solar radiation, humidity and temporal smoothing improves the model’s ability to describe heating demand but worsens that for cooling demand. We hypothesize that this is because of human behaviour. When it is hot, some occupants open windows to provide a cooling breeze. However, when it is cold, occupants do not open windows as this cools the building further. Therefore, at low temperatures, buildings are continuously insulated from the outdoors, retaining a ‘memory’ of the previous day’s conditions, and are constantly affected by solar gains and wind chill^{33,64}. At high temperatures, the accumulated effects of the previous days are reset once any windows are opened, hence the raw instantaneous outdoor temperature becomes a better descriptor of cooling demand^{33,64}. We address this issue by creating the final index from a weighted blend of BAIT and raw temperature, with weighting factors that depend on the raw temperature. The blend is active between a lower and upper threshold (the point at which people start to open windows and the point at which everyone who is going to has opened their windows), and these are mapped to the range of 10 (from –5 to +5) of a sigmoid function (meaning that they correspond to a 1% and 99% blend):

$$B = \frac{B_{max}}{1 + e^{-B'}} \quad (4)$$

$$B' = (T - 0.5(B_U + B_L)) \frac{(10)}{B_U - B_L} \quad (5)$$

where B is the blending amount, B_{max} is the maximum amount of blending to apply, B' is the input to the sigmoid function and B_U and B_L are the upper and lower thresholds for blending. Default values of $B_{max} = 0.5$, $B_U = 15$ °C and $B_L = 23$ °C were applied as these were found to give the best fit with the datasets we use. The interpretation of this is that between daily-average temperatures of 15 and 23 °C people begin opening windows, and beyond 23 °C buildings lose half of the influence of temperature smoothing and other meteorological variables due to changes in occupant behaviour.

The final temperature index is then calculated using

$$BAIT_{d,l} = (BAIT_{d,l}(1 - B_{d,l})) + (T_{d,l}B_{d,l}). \quad (6)$$

Calculating degree days

HDDs and CDDs are measures of by how much (in °C) and for how long (in days) the outdoor air temperature was higher or lower than a given balance point temperature. The thresholds of the balance

point are the temperatures above which heating is required and below which cooling is required. The threshold temperature for heating is typically lower than that for cooling because the balance point covers a range of temperatures in which occupants require neither heating nor cooling (that is, people are willing to tolerate a range of internal temperatures throughout the year). Here, we use BAIT and thresholds based on BAIT in place of outdoor temperature to calculate HDDs and CDDs:

$$HDD_{d,l} = (T_{heat} - BAIT_{d,l})^+ \quad (7)$$

$$CDD_{d,l} = (BAIT_{d,l} - T_{cool})^+ \quad (8)$$

where for each d and l T_{heat} is the threshold BAIT for heating (°C), T_{cool} is the threshold BAIT for cooling (°C) and the + indicates that only positive values are counted.

The geographic extent of individual grid cells in reanalysis models (0.625° × 0.5° for MERRA-2) means that “climatic variables represent the background regional conditions without any influence from urban areas”⁶⁵. They neglect the urban heat island effect, which increases surface temperatures in urban areas by around 1–2 °C (ref. 65), and so threshold temperatures in major cities may have a negative bias.

As part of our calibration, we use national and regional gas and electricity demand data, and so HDDs and CDDs must be aggregated accordingly. We calculate HDDs and CDDs for each grid cell that lies within the country or region, and then sum these together weighting each location’s values by the population contained within that grid cell:

$$HDD_{d,c} = \sum_{l \in c} \frac{HDD_{d,l} p_l}{p_c} \quad (9)$$

$$CDD_{d,c} = \sum_{l \in c} \frac{CDD_{d,l} p_l}{p_c} \quad (10)$$

where for each day (d), country/region (c) and grid location (l) p is population density, taken from ref. 50.

This approach relies on four simplifications: it uses daily average temperatures, which ignore diurnal variation; calibration and validation are based on total energy consumption, which includes end-uses other than space heating and cooling; regional aggregates are weighted according to population density, which neglects the locational distribution of non-residential buildings, and the aggregation assumes a simple relationship between population density and building size. These points are expanded upon in Supplementary Note 6.

Our validation reveals heterogeneity in building weather response and occupant behaviour within and between the datasets of individual buildings. This would be neglected if a single set of model parameters were used to simulate national or regional demand, as this would imply that every building has the same response to the weather (for example, level of insulation, air-tightness, amount of glazing) and every person has the same preferences for internal space temperature. To reflect the observed heterogeneity, the degree days for each grid cell were calculated 36 times using different combinations of parameters, and then averaged across all samples. Each parameter was varied around its central value using maximin Latin hypercube sampling⁶⁶. A six-dimensional hypercube was constructed with each dimension representing a model parameter, and samples were optimized to maximize the minimum distance between design points so that the combinations of parameters were well spread across the input space. These six parameters and the optimization used to derive their central values are described in Calibrating the model. Each dimension of the hypercube was then

translated from a [0–1] uniform distribution to a normal distribution, with an s.d. that reflects the mean s.d. seen within all building datasets (Supplementary Table 3).

Translating degree days into energy demand

We employ a simple linear model to translate degree days (derived from BAIT) into energy demand. This is defined by the five parameters illustrated in Fig. 1c, which shows the relationship between BAIT and electricity demand in a single home in Austin, TX.

The vertical lines (T_{heat}) and (T_{cool}) show the optimized threshold temperatures for heating and cooling. These thresholds are influenced by both behaviour (for example, the required comfort levels of occupants) and technical characteristics of the building (for example, passive gains from occupants and appliances are more influential in well insulated buildings, lowering the heating threshold for a given level of comfort). The region between T_{heat} and T_{cool} is the balance point temperature where neither heating nor cooling is required, and P_{base} in this window is independent of BAIT (for example, energy demand for water heating, lighting and appliances). When considering only the energy consumed for space heating and cooling, this baseline is equal to zero. The slopes (P_{heat}) and (P_{cool}) are the power coefficients for heating and cooling, which determine the additional energy consumed per HDD and CDD. For the home in Fig. 1c, we estimate heating and cooling power coefficients of 0.64 and 0.79 kWh d⁻¹ °C⁻¹. During the exceptional winter storm of February 2021, when temperatures in Austin fell to –12 °C, we predict that this home would consume an additional 16.2 kWh d⁻¹ for heating, contributing to the power crisis that left 11 million people without power⁶⁷. While heating and cooling power coefficients are roughly symmetrical in Fig. 1c, they can be highly asymmetrical elsewhere.

The Demand.ninja can be used for two purposes: (1) ‘what if’ simulations, where a user supplies the above parameters to estimate the energy demand of a building in a given location, and (2) to predict the historical energy demand in a building or region. For the latter, we perform a linear regression with fixed effects between measured daily energy demand, HDD, CDD and two time indicators. The first of these, W , is a binary indicator that splits working days from weekends and national holidays. The second, D , is used when modelling national or regional demand to account for growth in demand due to population or economic factors. For example, electricity demand in Great Britain has fallen by 1.5% per year since 2012 due to improving efficiency⁶⁸. The regression for total power demand (P_{total}) is given by

$$P_{\text{total}} = P_{\text{base}} + P_{\text{heat}} \text{HDD} + P_{\text{cool}} \text{CDD} + \alpha W + \beta D + \varepsilon \quad (11)$$

where P_{base} is in GW, P_{heat} and P_{cool} are in GW °C⁻¹ and ε is the model error term. α (in GW) and β (in GW yr⁻¹) are the time coefficients, which represent the impacts of differences in societal behaviour and long-term trends in power demand.

The impact of α is visible in Fig. 2e as the gap between pink and green points (weekdays and weekends). As this is defined as a fixed offset (in the case of New York, –1.5 GW) it has no impact on thermal demand. It does not affect the quality metrics reported throughout the paper, as these focus on weekday demand only. It is incorporated so that the model can produce 8,760 h time series of demand including the weekend/weekday split. The impact of β is not visible in Fig. 2 as New York has near-zero annual growth in electricity demand. A more prominent example is given in Supplementary Fig. 17, the electricity demand in Ireland, which has increased in part due to the growth of the technology industry and data centres.

A similar approach was used to create estimates of total national demand from other models for comparison with the Demand.ninja, as described in Supplementary Note 7 and shown in Supplementary Figs. 100–153.

Calibrating the model

This section is concerned with deriving the optimal values for BAIT parameters (box iv in Fig. 1a) and balance point temperature thresholds (box v).

For each individual building and each national dataset, we used a two-stage process to optimize the model parameters. A differential evolution algorithm⁶⁹ implemented in R⁷⁰ was used to find the BAIT parameters and temperature thresholds ($x, y, z, \sigma, T_{\text{heat}}, T_{\text{cool}}$) that minimized the s.d. of residuals between modelled and measured demand. This first tests a randomly selected population of candidate parameters, and evolves the population towards better solutions (with lower error on modelled demand) by mimicking evolutionary processes. An initial population of 150 parents was used (25 times the number of parameters) and 250 generations were tested to ensure convergence.

Tests were repeated with different random seeds to ensure that the stochastic initial population did not materially influence the final outcome (that is, no deviation within the first four significant figures of any resulting parameter). For each set of parameters, BAIT was constructed (equations (1)–(6)), and HDDs and CDDs were calculated from it (equations (7)–(10)). Energy demand was then estimated (equation (11)) to find the optimal energy parameters ($P_{\text{base}}, P_{\text{heat}}, P_{\text{cool}}, \alpha, \beta$).

The datasets used for calibration are summarized in Supplementary Tables 2 and 3. The results from this process are summarized in Supplementary Table 4 and Supplementary Figs. 43 and 44. The spread of results both within and across datasets reflects the diversity of building types, living standards and occupant behaviours. Supplementary Note 8 discusses regional differences in parameters and, wherever possible, compares our estimated parameters with others in the literature.

Resolving profiles to hourly resolution

We use daily average meteorological data rather than hourly or sub-hourly data. Variations in demand from day to day are primarily driven by weather, while variations in hourly demand are often driven by behavioural choices and occupancy patterns. For example, many commercial buildings are only heated or cooled during working hours (for example, 9:00 to 17:00, Monday to Friday), while many residential buildings are heated only in the evenings and mornings. If demand were modelled using hourly temperature the greatest demand for heating would occur during the night when temperatures are at their lowest.

We upscale the time series of heating, cooling and total demand from daily to hourly resolution using representative 24 h diurnal profiles for each component of demand. We use globally representative profiles for heating and cooling demand as we find these are relatively consistent across regions. When calculating total electricity demand, we use regionally and seasonally disaggregated profiles for baseline demand (for non-thermal end-uses) to reflect the different cultural and economic contexts around the world.

Diurnal profiles for the three components of demand are derived by following seven steps.

- (1) Run the model with optimal parameters for each country/region to derive HDDs and CDDs for each day.
- (2) Group days according to weekday/weekend and summer/winter/shoulder (spring and autumn) days.
- (3) Classify days in each group according to the number of HDDs and CDDs. Days that are above the 90th percentile for HDDs or CDDs for that specific region and season are classified as cold or hot days, respectively. Days that fall below the 10th percentile for both HDDs and CDDs are classified as mild days.
- (4) Identify electricity demand within each hour of the day for all days within each of these groups (cold/mild/hot, weekday/weekend, summer/winter/shoulder) and summarize the diurnal profile for each group using its mean and s.d.
- (5) Subtract mean hourly electricity demand profiles on mild days from cold and hot days for each group to isolate the profile of

- temperature-dependent electricity demand, which is used as a proxy for heating and cooling demand.
- (6) For heating and cooling profiles, calculate the mean profile across all days (aggregating weekday/weekend and summer/winter/shoulder) to give a characteristic hourly diurnal profile for each country or region.
 - (7) Further aggregate the heating and cooling profiles by calculating the mean across all national and regional profiles, weighted according to the share of demand within each region, to give a single globally representative hourly diurnal profile for each.

The profiles from steps 1–4 are visualized in Supplementary Figs. 170–206 for all countries and regions that we study. Comparison between the panels within each subfigure shows that baseline profiles (shown by black lines) are relatively variable over time within any given country, but additional demands from heating and cooling (shown by the heights of the red and blue shaded areas) are relatively constant in comparison. This provides the rationale for step 6: the baseline profile (for non-thermal electricity demand) needs to be specific to the season and working versus non-working days, but the profiles for heating and cooling do not. The profiles for heating and cooling demand that result from step 5 are visualized in Supplementary Fig. 207 for all countries and regions we study. Comparison between the panels shows that these profiles are relatively consistent around the world, providing the rationale for step 7. These simplifications are discussed further in Supplementary Note 6, and the final output of step 7 is shown in Extended Data Fig. 1.

The final hourly time series of demand for each country/region is then assembled by first classifying days according to their type (weekday/weekend) and season (summer/shoulder/winter). The non-thermal demand, P_{base} , on each day is then convoluted with the respective baseline profile (black lines in Supplementary Figs. 170–206) for the specific country/region. Finally, the demands for heating and cooling are convoluted with the global average heating and cooling profiles (Extended Data Fig. 1) across all days, and added to this baseline.

Modelling the impact of lowering building temperatures

We model the impact of turning down the thermostat by adjusting the heating threshold temperature, which relies on the assumption that reducing the threshold temperature by 1 °C is equivalent to reducing the desired internal temperature of the building by 1 °C. This implies that heating a house to 20 °C when it is 10 °C outside should require the same energy as heating to 21 °C when it is 11 °C outside. This holds in general, as the rate of heat loss from a building is proportional to the difference between internal and external temperatures for the primary sources of heat loss (conduction and air changes)⁷¹.

We model gas savings in individual buildings and across regions using the same approach but with different data inputs. For buildings, we use daily gas demand data from 2,062 houses from the UK-based Energy Demand Research Project dataset³⁷. For regions, we use monthly gas demand data for 29 countries in Europe (EU27, United Kingdom and Turkey)⁷² and 49 US states (all except Hawaii)⁴³. We first calibrate the model using measured demand in each building or region to give optimal parameters (summarized in Supplementary Table 4). We then decrease the optimal heating threshold by 1 °C across all buildings and regions. The difference between gas demand in the optimal simulation and with the decreased threshold temperature is the gas saving. Gas savings are converted into emission savings and economic savings using the emission factor in ref. 73 and the average price of natural gas from 1 September 2021 to 31 March 2022. The Dutch TTF⁷⁴ and British NBP⁷⁵ prices are averaged for Europe, and the Henry Hub price⁷⁶ is used for the United States.

Modelling global heating and cooling demand

The Demand.ninja can be generalized to estimate heating and cooling demand in any region or building globally. We first calculate BAIT

at all locations globally and then use BAIT to calculate degree days. Degree days are calculated at the subnational level within larger countries, using GADM1 subdivisions (that is, at the level of US states or Chinese provinces), except in Europe, where we use NUTS level 2. We calculate HDDs and CDDs at the national level in 132 smaller countries, with a land area below 15,000 km² and a population below 50 million. Deriving bespoke BAIT parameters and heating and cooling thresholds is useful for maximizing the quality of the fit when predicting metered demand in known buildings; however, these parameters must be generalized when modelling energy demand at locations without temporally resolved data on energy demand. We therefore calculate BAIT using the following global average parameters (approximated from the average values in Supplementary Table 4): a smoothing coefficient of $0.50 \pm 0.23 \text{ d}^{-1}$, solar coefficient of $0.012 \pm 0.008 \text{ °C W}^{-1} \text{ m}^{-2}$, wind coefficient of $-0.20 \pm 0.17 \text{ °C (m s}^{-1})^{-1}$ and humidity coefficient of $0.050 \pm 0.065 \text{ °C (g kg}^{-1})^{-1}$. HDDs and CDDs are then calculated using thresholds of 14 ± 2 and $20 \pm 2 \text{ °C}$ respectively, approximately the average across all regions for which we have daily electricity and/or gas demand data for model calibration. These are the default parameters used in the Demand.ninja's web interface.

To produce the maps in Fig. 8, CDDs were calculated using the same generic parameters in all world regions for each year between 1980 and 2022, and the change in CDDs over time was determined through linear (for absolute change) or log-linear (for percentage change) regression.

Data availability

Results from the model that support the findings of this study are presented in the article and its Supplementary Information. Additional data are available from the Demand.ninja website (<https://demand.ninja>). The metered energy demand data used for validation are available from the sources given in Supplementary Information. The weather data used in this study are available from NASA (<https://gmao.gsfc.nasa.gov/reanalysis/MERRA-2/>). Source data are provided with this paper.

Code availability

The Demand.ninja model is available from the GitHub code repository in Python (<https://github.com/renewables-ninja/demand-ninja>) and R (https://github.com/iain-staffell/demand_ninja).

References

1. IEA *Renewables 2019* <https://www.iea.org/reports/renewables-2019> (2020).
2. IEA *Is Cooling the Future of Heating?* <https://www.iea.org/commentaries/is-cooling-the-future-of-heating> (2020).
3. Isaac, M. & van Vuuren, D. P. Modeling global residential sector energy demand for heating and air conditioning in the context of climate change. *Energy Policy* **37**, 507–521 (2009).
4. IEA *The Future of Cooling* <https://www.iea.org/reports/the-future-of-cooling> (2018).
5. Gi, K., Sano, F., Hayashi, A., Tomoda, T. & Akimoto, K. A global analysis of residential heating and cooling service demand and cost-effective energy consumption under different climate change scenarios up to 2050. *Mitig. Adapt. Strateg. Glob. Chang.* **23**, 51–79 (2018).
6. Pfenninger, S. & Staffell, I. Long-term patterns of European PV output using 30 years of validated hourly reanalysis and satellite data. *Energy* **114**, 1251–1265 (2016).
7. Staffell, I. & Pfenninger, S. Using bias-corrected reanalysis to simulate current and future wind power output. *Energy* **114**, 1224–1239 (2016).
8. Ruhnau, O., Hirth, L. & Praktiknjo, A. Time series of heat demand and heat pump efficiency for energy system modeling. *Sci. Data* **6**, 189 (2019).

9. Pezzutto, S. et al. *D2.3 WP2 Report—Open Data Set for the EU28* https://www.hotmaps-project.eu/wp-content/uploads/2018/03/D2.3-Hotmaps_for-upload_revised-final_.pdf (2019).
10. Wilson, E. J. H. et al. *End-Use Load Profiles for the U.S. Building Stock: Methodology and Results of Model Calibration, Validation, and Uncertainty Quantification*. Report No. NREL/TP-5500-80889 <https://www.osti.gov/biblio/1854582/> (US Department of Energy Office of Scientific and Technical Information, 2022).
11. Ruhnau, O. & Muessel, J. Update and extension of the When2Heat dataset. *EconStor* <http://hdl.handle.net/10419/249997> (ZBW—Leibniz Information Centre for Economics, 2022).
12. IEA *A 10-Point Plan to Cut Oil Use* <https://www.iea.org/reports/a-10-point-plan-to-cut-oil-use> (2022).
13. IEA *A 10-Point Plan to Reduce the European Union's Reliance on Russian Natural Gas* <https://www.iea.org/reports/a-10-point-plan-to-reduce-the-european-unions-reliance-on-russian-natural-gas> (2022).
14. Associated Press *Nations “United” in Seeking to Cut Russian Oil, Gas Imports* <https://apnews.com/article/russia-ukraine-climate-business-united-states-global-trade-fe40c84b36ed311ac60bf2ecdbdc20f5> (2022).
15. Creutzig, F. Fuel crisis: slash demand in three sectors to protect economies and climate. *Nature* **606**, 460–462 (2022).
16. Waite, M. & Modi, V. Electricity load implications of space heating decarbonization pathways. *Joule* **4**, 376–394 (2020).
17. Zhao, Q. et al. Global, regional, and national burden of mortality associated with non-optimal ambient temperatures from 2000 to 2019: a three-stage modelling study. *Lancet Planet. Health* **5**, e415–e425 (2021).
18. Kikstra, J. S., Mastrucci, A., Min, J., Riahi, K. & Rao, N. D. Decent living gaps and energy needs around the world. *Environ. Res. Lett.* **16**, 095006 (2021).
19. Turner, C. & Frankel, M. *Energy Performance of LEED for New Construction Buildings* <https://newbuildings.org/resource/energy-performance-leed-new-construction-buildings/> (New Buildings Institute, 2010).
20. Chung, W. Review of building energy-use performance benchmarking methodologies. *Appl. Energy* **88**, 1470–1479 (2011).
21. Pang, X., Wetter, M., Bhattacharya, P. & Haves, P. A framework for simulation-based real-time whole building performance assessment. *Build. Environ.* **54**, 100–108 (2012).
22. Azar, E. & Menassa, C. C. A comprehensive analysis of the impact of occupancy parameters in energy simulation of office buildings. *Energy Build.* **55**, 841–853 (2012).
23. Scofield, J. H. Do LEED-certified buildings save energy? Not really.... *Energy Build.* **41**, 1386–1390 (2009).
24. Karlsson, F., Rohdin, P. & Persson, M.-L. Measured and predicted energy demand of a low energy building: important aspects when using building energy simulation. *Build. Serv. Eng. Res. Technol.* **28**, 223–235 (2007).
25. Tsanas, A. & Xifara, A. Accurate quantitative estimation of energy performance of residential buildings using statistical machine learning tools. *Energy Build.* **49**, 560–567 (2012).
26. Robinson, C. et al. Machine learning approaches for estimating commercial building energy consumption. *Appl. Energy* **208**, 889–904 (2017).
27. Fumo, N. & Rafe Biswas, M. A. Regression analysis for prediction of residential energy consumption. *Renew. Sustain. Energy Rev.* **47**, 332–343 (2015).
28. Waite, M. et al. Global trends in urban electricity demands for cooling and heating. *Energy* **127**, 786–802 (2017).
29. Westermann, P. & Evins, R. Surrogate modelling for sustainable building design—a review. *Energy Build.* **198**, 170–186 (2019).
30. Van Gelder, L., Das, P., Janssen, H. & Roels, S. Comparative study of metamodelling techniques in building energy simulation: guidelines for practitioners. *Simul. Model. Pract. Theory* **49**, 245–257 (2014).
31. Westermann, P., Welzel, M. & Evins, R. Using a deep temporal convolutional network as a building energy surrogate model that spans multiple climate zones. *Appl. Energy* **278**, 115563 (2020).
32. Catalina, T., Iordache, V. & Caracaleanu, B. Multiple regression model for fast prediction of the heating energy demand. *Energy Build.* **57**, 302–312 (2013).
33. ASHRAE *ASHRAE Handbook 2021—Fundamentals* (2021).
34. Peacock, M., Fragaki, A. & Matuszewski, B. J. in *Energy and Sustainable Futures. Springer Proceedings in Energy* (eds Mporas, I. et al.) 53–60 (Springer, 2021).
35. US Census Bureau *American Community Survey* <https://data.census.gov/cedsci/table?t=HeatingandAirConditioning%28HVAC%29&g=0100000US%240400000&y=2020&tid=ACSDT5Y2020.B25040&tp=true> (2020).
36. Staffell, I. & Green, R. How does wind farm performance decline with age? *Renew. Energy* **66**, 775–786 (2014).
37. AECOM Building Engineering *Energy Demand Research Project: Early Smart Meter Trials, 2007–2010* <https://doi.org/10.5255/UKDA-SN-7591-1> (2014).
38. British Gas *“The Fixed One v19”—Gas & Electricity Quotes* <https://www.britishgas.co.uk/GetAQuote/new/quote-details> (2022).
39. Friedlingstein, P. et al. Global carbon budget 2021. *Earth Syst. Sci. Data* **14**, 1917–2005 (2022).
40. UK Department for Business Energy and Industrial Strategy *Cooling in the UK* <https://www.gov.uk/government/publications/cooling-in-the-uk> (2021).
41. Gianniou, P., Reinhart, C., Hsu, D., Heller, A. & Rode, C. Estimation of temperature setpoints and heat transfer coefficients among residential buildings in Denmark based on smart meter data. *Build. Environ.* **139**, 125–133 (2018).
42. Johnson, N. J., Gross, R. & Staffell, I. Stabilisation wedges: measuring progress towards transforming the global energy and land use systems. *Environ. Res. Lett.* **16**, 64011 (2021).
43. US Energy Information Administration *Natural Gas Monthly* <https://www.eia.gov/naturalgas/data.php> (2022).
44. Ember *Top 10 EU Emitters All Coal Power Plants in 2021* <https://ember-climate.org/insights/research/top-10-emitters-in-the-eu-ets-2021/> (2022).
45. US Environmental Protection Agency *Power Sector Facility Level Comparisons: 2021 vs. 2022* <https://www.epa.gov/power-sector/facility-level-comparisons#LongTerm> (2023).
46. Perkins-Kirkpatrick, S. E. & Lewis, S. C. Increasing trends in regional heatwaves. *Nat. Commun.* **11**, 3357 (2020).
47. Chatterjee, S. et al. Existing tools, user needs and required model adjustments for energy demand modelling of a carbon-neutral Europe. *Energy Res. Soc. Sci.* **90**, 102662 (2022).
48. Gelaro, R. et al. The modern-era retrospective analysis for research and applications, version 2 (MERRA-2). *J. Clim.* **30**, 5419–5454 (2017).
49. Hersbach, H. et al. The ERA5 global reanalysis. *Q. J. R. Meteorol. Soc.* **146**, 1999–2049 (2020).
50. Center for International Earth Science Information Network (CIESIN)—Columbia University *Gridded Population of the World, Version 4 (GPWv4): Population Count Adjusted to Match 2015 Revision of UN WPP Country Totals* <https://doi.org/10.7927/H4F47M65> (2018).
51. Mehlig, D., ApSimon, H. & Staffell, I. The impact of the UK's COVID-19 lockdowns on energy demand and emissions. *Environ. Res. Lett.* **16**, 054037 (2021).

52. Muggeo, V. M. R. Estimating regression models with unknown break-points. *Stat. Med.* **22**, 3055–3071 (2003).
53. Staffell, I. *Fuel Cells for Domestic Heat and Power: Are They Worth It?* PhD thesis, Univ. Birmingham (2010).
54. Masterson, J. & Richardson, F. A. *Humidex; a Method of Quantifying Human Discomfort Due to Excessive Heat and Humidity* (Environment Canada, 1979).
55. Crawley, D. B., Hand, J. W., Kummert, M. & Griffith, B. T. Contrasting the capabilities of building energy performance simulation programs. *Build. Environ.* **43**, 661–673 (2008).
56. Seyedzadeh, S., Rahimian, F. P., Glesk, I. & Roper, M. Machine learning for estimation of building energy consumption and performance: a review. *Vis. Eng.* **6**, 5 (2018).
57. Osczevski, R. & Bluestein, M. The new wind chill equivalent temperature chart. *Bull. Am. Meteorol. Soc.* **86**, 1453–1458 (2005).
58. Moran, D. et al. An environmental stress index (ESI) as a substitute for the wet bulb globe temperature (WBGT). *J. Therm. Biol.* **26**, 427–431 (2001).
59. Atalla, T., Gualdi, S. & Lanza, A. A global degree days database for energy-related applications. *Energy* **143**, 1048–1055 (2018).
60. Spinoni, J. et al. Changes of heating and cooling degree-days in Europe from 1981 to 2100. *Int. J. Climatol.* **38**, e191–e208 (2018).
61. Ramon, D., Allacker, K., De Troyer, F., Wouters, H. & van Lipzig, N. P. M. Future heating and cooling degree days for Belgium under a high-end climate change scenario. *Energy Build.* **216**, 109935 (2020).
62. National Grid Gas Demand Forecasting Methodology <https://www.nationalgas.com/document/132516/download> (2016).
63. Elexon Load Profiles and their Use in Electricity Settlement <https://bscdocs.elexon.co.uk/guidance-notes/load-profiles-and-their-use-in-electricity-settlement> (2018).
64. D'Oca, S., Fabi, V., Corgnati, S. P. & Andersen, R. K. Effect of thermostat and window opening occupant behavior models on energy use in homes. *Build. Simul.* **7**, 683–694 (2014).
65. Manoli, G. et al. Magnitude of urban heat islands largely explained by climate and population. *Nature* **573**, 55–60 (2019).
66. Stein, M. Large sample properties of simulations using Latin hypercube sampling. *Technometrics* **29**, 143–151 (1987).
67. Gold, R. The Texas electric grid failure was a warm-up. *Texas Monthly* <https://www.texasmonthly.com/news-politics/texas-electric-grid-failure-warm-up/> (2022).
68. Green, R. & Staffell, I. The contribution of taxes, subsidies, and regulations to British electricity decarbonization. *Joule* **5**, 2625–2645 (2021).
69. Price, K., Storn, R. & Lampinen, J. *Differential Evolution—A Practical Approach to Global Optimization* (Springer, 2006).
70. Mullen, K., Ardia, D., Gil, D. L., Windover, D. & Cline, J. DEoptim: an R package for global optimization by differential evolution. *J. Stat. Softw.* **40**, 1–26 (2011).
71. Davies, M. G. *Building Heat Transfer* (Wiley, 2004).
72. Eurostat Eurostat Energy Database (nrg_cb_gasm) <https://ec.europa.eu/eurostat/web/energy/data/database> (2022).
73. Staffell, I. *The Energy and Fuel Data Sheet* https://www.academia.edu/1073990/The_Energy_and_Fuel_Data_Sheet (2011).
74. Yahoo Finance Dutch TTF Natural Gas Calendar <https://finance.yahoo.com/quote/TTF%3DF/history?period1=1621848432&period2=1653384432&interval=1wk&filter=history&frequency=1wk&includeAdjustedClose=true> (2022).
75. ERCE UK Natural Gas NBP Spot Price—ERCE <https://www.erce.energy/graph/uk-natural-gas-nbp-spot-price/> (2022).
76. US EIA Henry Hub Natural Gas Spot Price (Dollars per Million Btu) <https://www.eia.gov/dnav/ng/hist/rngwhhdm.htm> (2022).
77. Halttunen, K. et al. Global assessment of the merit-order effect and revenue cannibalisation for variable renewable energy. Preprint at SSRN <https://doi.org/10.2139/ssrn.3741232> (2020).
78. National Grid Data Item Explorer <https://mip-prd-web.azurewebsites.net/DataItemExplorer> (2022).
79. Crozier, C. & Baker, K. The effect of renewable electricity generation on the value of cross-border interconnection. *Appl. Energy* **324**, 119717 (2022).

Acknowledgements

I.S. was funded by the EPSRC under EP/T023031/1. N.J. was funded by the EPSRC under EP/T023031/1 and EP/R045518/1.

Author contributions

I.S. conceived and led the study, developed the analysis and visualized the results. S.P. developed the online interactive version of the model. N.J. wrote the paper with support from I.S., and acquired and processed the data. All authors edited and discussed the paper.

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at <https://doi.org/10.1038/s41560-023-01341-5>.

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41560-023-01341-5>.

Correspondence and requests for materials should be addressed to Iain Staffell.

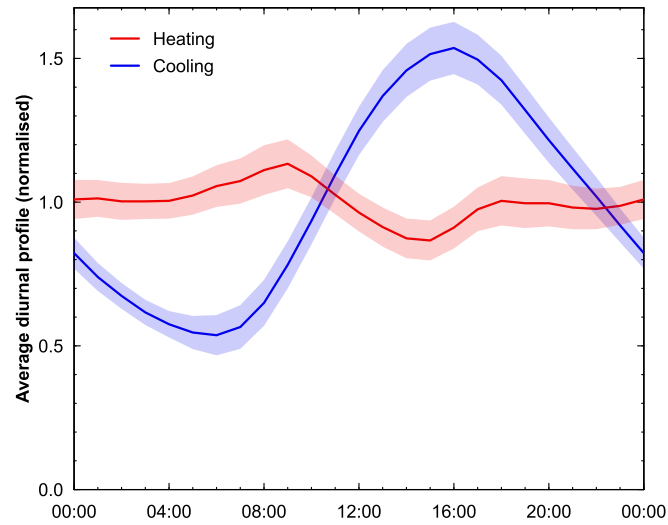
Peer review information *Nature Energy* thanks Lukas Kranzl, Jesus Lizana and Christoph Reinhart for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

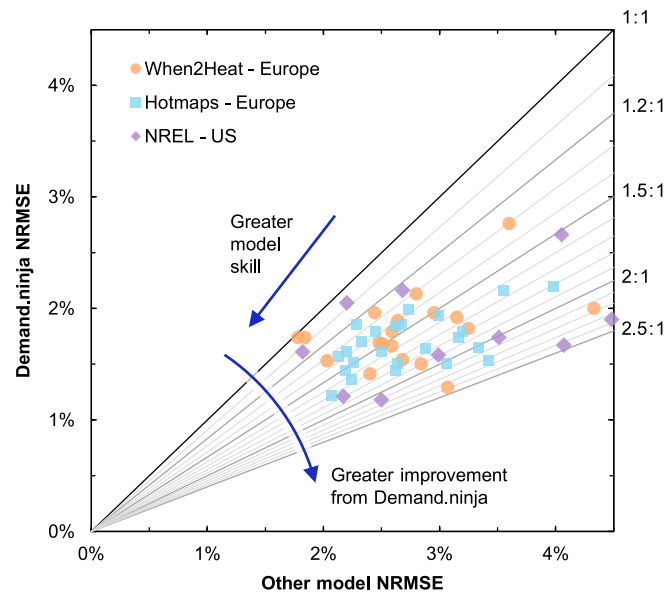
Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023, corrected publication 2023



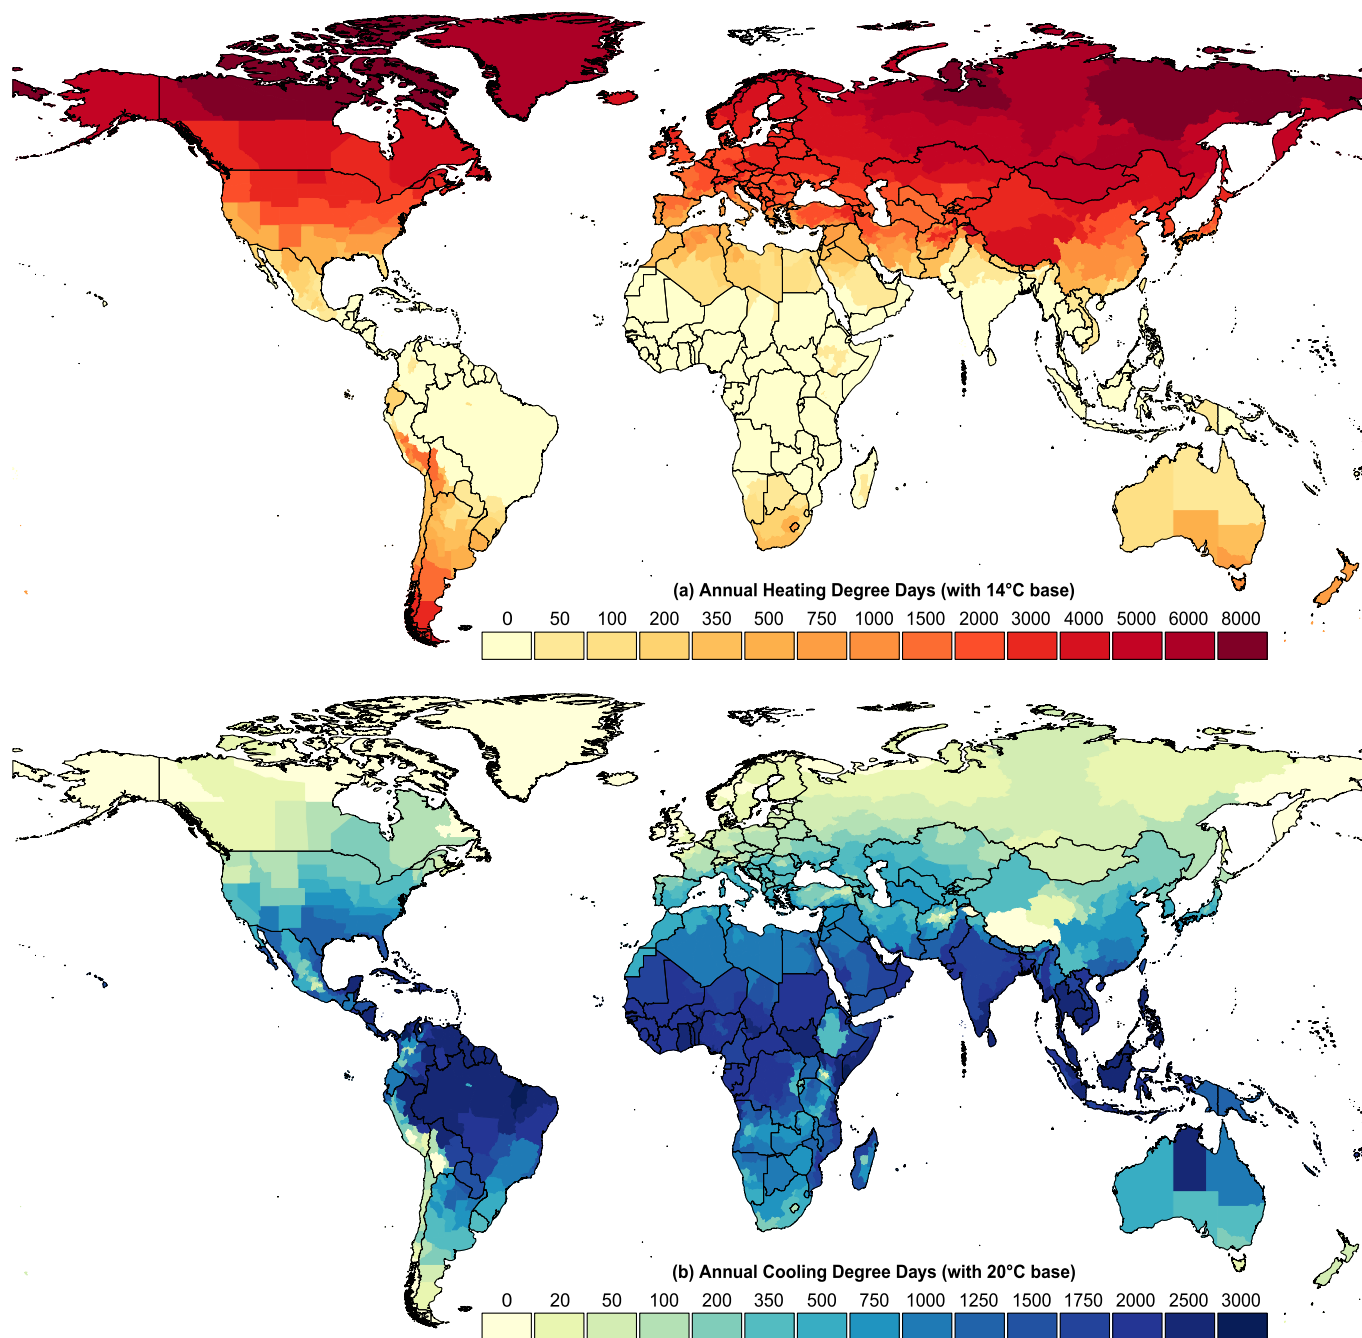
Extended Data Fig. 1 | Global average diurnal profiles for heating and cooling demand. Lines show the mean across all the countries and regions we consider (shown in Supplementary Fig. 207), weighted by annual demand in each region.

Shaded areas show the weighted standard deviation across the regional means, which captures the diversity between countries but excludes the diversity between days within countries.



Extended Data Fig. 2 | Improvement in accuracy of modelling daily electricity demand with Demand.ninja relative to other models. Data points show the normalised root mean squared error (NRMSE) between modelled and measured

demand from Demand.ninja and other models: 2016–2019 for When2Heat¹¹, 2010 for Hotmaps⁹ and 2018 for NREL¹⁰. Demand.ninja is shown with temperature thresholds and BAIT parameters calibrated for each region. See Fig. 4e.



Extended Data Fig. 3 | Global maps of average annual heating and cooling degree days. The colour of each region shows the mean number of degree days experienced across the years 2010-19. Panel (a) shows heating degree days (HDDs) calculated with a threshold temperature of 14 ± 2 °C. Panel (b) shows cooling degree days (CDDs) calculated with a threshold temperature of

20 ± 2 °C. Both panels use the global average BAIT parameters in all regions: a smoothing coefficient of $0.50 \pm 0.23 \text{ day}^{-1}$, solar coefficient of $0.012 \pm 0.008 \text{ °C}/(\text{Wm}^2)$, wind coefficient of $-0.20 \pm 0.17 \text{ °C}/(\text{ms}^{-1})$, and humidity coefficient of $0.050 \pm 0.065 \text{ °C}/(\text{gkg}^{-1})$.