

Reconciling Global Data-Centre Electricity Projections for 2030: A Boundary-Aware Evidence Map

Aditya Singh

Independent Researcher

Abstract: ***Background:** Published projections of global data-centre electricity use are frequently compared despite differences in geography, facility and workload coverage, units, and scenario design. **Methods:** We constructed a secondary-source-seeded, boundary-aware evidence map of projections for annual global operational electricity use by total data centres in 2030. Each publication was coded across eight comparability dimensions and assigned a transparent primary-source verification status. The primary synthesis treats publications reporting only range endpoints as interval-censored rather than interpreting their endpoints as complete scenarios. Publication-level bootstrap intervals and verification, peer-review, scenario-completeness, aggregation, and influence sensitivities were examined. **Results:** The corpus contains 28 values from 15 publications; 7 were directly checked in primary or author full text by one reviewer and 8 remain verification incomplete. The interval-censored model estimates a median latent publication centre of 912 TWh yr⁻¹ (95% publication-bootstrap interval 712–1217) and a geometric standard deviation of 1.69. A conventional geometric-scenario summary gives a median of 945 TWh yr⁻¹ (interquartile range 648–1167). Method, publisher, and vintage contrasts are imprecise and are retained only as exploratory sensitivities. **Conclusions:** Boundary harmonisation makes the literature more interpretable but does not identify a probability distribution for realised 2030 demand. The principal contribution is an auditable comparability taxonomy and minimum reporting standard, not a consensus forecast.*

Keywords: data centres; electricity demand; forecasts; evidence mapping; interval censoring; system boundaries; reporting standards

1. Introduction

Forecasts of data-centre electricity demand influence generation planning, network investment, technology policy, and assessments of artificial intelligence. Published 2030 values appear to span orders of magnitude, but that apparent disagreement combines differences in system definition, model architecture, and scenario design (Kamiya and Coroama, 2025; Koomey and Masanet, 2021). A national grid-connection estimate, an AI-only increment, a global whole-facility electricity forecast, and an estimate for the wider information and communication technology system do not measure the same quantity.

Converting such quantities to a common unit without first aligning their boundaries can create false precision. Capacity in gigawatts is not annual electricity consumption unless utilisation and time assumptions are known. IT equipment load is not whole-facility demand unless power-delivery and cooling overheads are included. Similarly, a low–high range does not reveal whether intermediate scenarios exist or whether either endpoint was considered probable by the original authors.

This study therefore examines a narrow estimand: annual global operational electricity used by total data centres in 2030. Its empirical objective is to describe the distribution of eligible projections without silently redefining them. Its methodological objective is to specify the minimum boundary information needed to determine whether forecasts are comparable. The study is deliberately presented as a secondary-source-seeded evidence map, not as a completed systematic review: the search did not cover all subscription databases, screening was conducted by one reviewer, and some retained sources remain verification incomplete.

1.1 Literature context and contribution

Global data-centre and wider ICT energy estimation has evolved from inventory studies to stock, traffic, workload, and facility models (Koomey, 2008; Malmudin et al., 2010; Van Heddeghem et al., 2014). Later work examined efficiency trends, sector boundaries, and the extent to which service growth has decoupled from electricity demand (Avgerinou et al., 2017; Malmudin and Lunden, 2018; Shehabi et al., 2018; Masanet et al., 2020). Updated inventories continue to revise both activity data and accounting scope (Jones, 2018; Malmudin et al., 2024).

Prospective estimates likewise use heterogeneous architectures. They include high-level traffic or demand extrapolation (Andrae and Edler, 2015; Andrae, 2019a, 2020b,a), bottom-up equipment models (Liu et al., 2020; Deloitte Global, 2024; International Energy Agency, 2025), and system-dynamics or scenario frameworks (Koot and Wijnhoven, 2021; Castro et al., 2024; Fan et al., 2026). Adjacent research on networks, streaming, devices, AI, water, and cryptocurrency is important but often uses different functional units and boundaries (Aslan et al., 2018; Shehabi et al., 2014; Preist et al., 2016; Henderson et al., 2020; de Vries, 2023; Mytton, 2021; Obringer et al., 2021; Stoll et al., 2019).

Previous critiques have warned against comparing incompatible ICT and Internet estimates (Freitag et al., 2021; Koomey and Masanet, 2021). This study extends that argument in three ways. First, it applies explicit record-level eligibility and verification rules to a defined 2030 data-centre estimand. Second, it treats endpoint-only publications as interval observations in the primary synthesis. Third, it translates the observed comparability problems into an

eight-dimensional reporting framework that can be used by forecast authors, reviewers, and planners.

1.2 Research Questions

The analysis asks four linked questions: which 2030 projections satisfy the common estimand; which boundaries determine comparability; how sensitive the distribution is to censoring, verification, source type, aggregation, and influence; and what the harmonised evidence can legitimately support for energy planning. The first two questions define the admissible evidence set, the third evaluates robustness, and the fourth translates rather than extends the statistical results.

2. Data and Methods

2.1 Design, discovery, and eligibility

The evidence map was seeded by Table 3.1 of the IEA-4E critical review (Kamiya and Coroama, 2025). Its bibliography was used to locate publisher, institutional, and author versions of candidate sources. A public-web update was conducted on August 5–6, 2026 with four archived strings combining “global data center” or “global data centre” with electricity, 2030, forecast, projection, and scenario. Crossref and OpenAlex requests were attempted but blocked by the execution environment. Strategies for Scopus, Web of Science Core Collection, IEEE Xplore, and Google Scholar are archived but were not executed because authenticated access was unavailable. The search archive therefore supports reproduction of the reported evidence map but not a claim of exhaustive or PRISMA-compliant retrieval. A record was eligible only when it provided:

- 1) Annual operational electricity, or an average load convertible without an undisclosed utilisation assumption;
- 2) Global geographic coverage;
- 3) Total data-centre coverage rather than an AI-only increment, external network, cryptocurrency-only load, or wider ICT total;
- 4) A 2030 target; and
- 5) A traceable source, quantity, unit, and scenario description.

Capacity-only estimates were excluded when utilisation or annualisation assumptions were unavailable. Multiple reports from the same author or organisation were retained when the forecast vintage, model, or scenario set differed. Full exclusion reasons and report linkages are supplied in the search archive.

2.2 Formal comparability and unit harmonisation

Each publication i is represented by an eight-component boundary vector

$$\mathbf{b}_i = (G_i, F_i, W_i, O_i, Q_i, S_i, M_i, C_i), \quad A_i = \prod_{d \in \mathcal{D}_*} \mathbf{1}(b_{id} \in \mathcal{A}_d), \quad (1)$$

where the components denote geography, facility, workload, ownership, quantity, scenario, model, and conversion information. $\mathcal{D}_*$ contains the dimensions required by the estimand, $\mathcal{A}_d$ is the admissible set for dimension d , and $A_i = 1$ is required for quantitative inclusion. Thus, a numerically convertible value is still excluded when its geography or workload boundary is incompatible.

$$E_i = 8.76 P_i u_i \kappa_i, \quad (2)$$

When a source reports power rather than annual energy, the harmonisation rule is

where E_i is TWh yr⁻¹, P_i is GW, u_i is annual utilisation, and κ_i is the disclosed facility-boundary multiplier (equal to PUE for IT load and one for whole-facility average load). A power estimate is not converted when u_i or κ_i is required but unavailable.

2.3 Primary-source verification and coding

The coding rubric records geography; facility, workload, and ownership boundaries; treatment of AI, cryptocurrency, networks, and on-premises facilities; original quantity and unit; conversion assumptions; scenario labels; base year; growth and PUE assumptions; and model architecture. Unknown is distinct from excluded or absent. A publication is labelled “directly checked” only when its value, unit, scope, and scenario can be located in primary or author full text. Public summaries of inaccessible reports remain explicitly verification incomplete.

Direct checks cover Andrae and Edler, two 2019 Andrae reports, Andrae (2020), Deloitte Global, Koot and Wijnhoven, and the IEA 2025 Base Case (Andrae and Edler, 2015; Andrae, 2019a,b, 2020a; Deloitte Global, 2024; Koot and Wijnhoven, 2021; International Energy Agency, 2025). The source audit corrected the rounded Andrae–Edler values and excluded five inherited candidates whose primary documents did not reproduce the required year, scope, outcome, or conversion. One 2026 candidate remains outside the synthesis pending exact scenario extraction.


2.4 Boundary taxonomy and source appraisal

Equation (1) operationalises the eight dimensions in Figure 1. Unknown ancillary assumptions are retained as missing data rather than imputed.

Source reporting was appraised using seven non-compensatory domains informed by CHARMS-style prediction-model extraction and AACODS grey-literature assessment (Moons et al., 2014; Tyndall, 2010): provenance, boundary applicability, outcome and unit integrity, scenario recoverability, model transparency, validation and uncertainty, and independence or update risk. Domains were coded as low concern, some concern, high concern, or insufficient information. No composite quality score or statistical quality weight was calculated.

Table 1: Characteristics and verification status of the analytical corpus

Characteristic	Count
Publications	15
Entered 2030 scenario values	28
Endpoint-only multi-scenario publications	5
Extrapolation method	5
Equipment-led, hybrid, or system-dynamics method	8
Method undocumented	2
Commercial publisher group	10
Independent publisher group	5
Primary full text directly checked	7
Primary source located; verification incomplete	8

**Figure 1:** Boundary taxonomy for determining whether electricity projections are comparable. Each dimension must be coded or explicitly marked unknown before pooling

2.5 Primary interval-censored synthesis

The unit of synthesis is the publication. For publications with adequately recovered scenario sets, the point summary is the publication-level geometric scenario mean,

$$\tilde{y}_i = \exp\left(\frac{1}{n_i} \sum_{j=1}^{n_i} \log y_{ij}\right). \quad (3)$$

This constructed summary is not described as an author-preferred forecast unless the source labels it accordingly. Publications represented only by a reported low–high range contribute the interval $[L_i, U_i]$ instead of two observed scenarios. The primary model assumes latent publication centres are normally distributed on the log scale. Let $z_i = \log \tilde{y}_i$, $a_i = \log L_i$, and $b_i = \log U_i$. For point publications $i \in \mathcal{P}$

and endpoint-only publications $i \in \mathcal{R}$, the log-likelihood is

$$\ell(\mu, \sigma) = \sum_{i \in \mathcal{P}} \left[-\log \sigma + \log \phi\left(\frac{z_i - \mu}{\sigma}\right) \right] + \sum_{i \in \mathcal{R}} \log \left[\Phi\left(\frac{b_i - \mu}{\sigma}\right) - \Phi\left(\frac{a_i - \mu}{\sigma}\right) \right], \quad (4)$$

where ϕ and Φ are the standard-normal density and distribution functions. Maximum likelihood gives median(Y) = $\exp(\hat{\mu})$ and GSD(Y) = $\exp(\hat{\sigma})$. Uncertainty is obtained from 999 publication-bootstrap resamples.

2.6 Secondary descriptions and sensitivity analyses

For K publications, the empirical distribution, quantiles, and equal-publication geometric mean are

$$\hat{F}_K(x) = \frac{1}{K} \sum_{i=1}^K \mathbf{1}(\tilde{y}_i \leq x), \quad \hat{Q}_p = \inf\{x: \hat{F}_K(x) \geq p\}, \quad \bar{y}_g = \exp\left(\frac{1}{K} \sum_{i=1}^K \log \tilde{y}_i\right). \quad (5)$$

Sensitivities restrict the corpus by verification, peer review, central-case availability, scenario completeness, and endpoint status; change the aggregation rule; and delete each publication in turn.

$$\log y_{ij} = \alpha + \mathbf{x}_i^\top \boldsymbol{\beta} + u_i + \varepsilon_{ij}, \quad u_i \sim N(0, \tau^2), \quad \varepsilon_{ij} \sim N(0, \sigma_w^2), \quad \text{ICC} = \frac{\tau^2}{\tau^2 + \sigma_w^2}. \quad (6)$$

Endpoint-only ranges and only four complete multi-scenario sets limit identification, so this decomposition and associations with method, publisher, and year are secondary. For association models, each publication receives total weight one. If publication i contributes n_i entered scenarios, $w_{ij} = 1/n_i$ and

$$p_{\text{perm}} = \frac{1}{|\Pi|} \sum_{\pi \in \Pi} \mathbf{1}(|T_\pi| \geq |T_{\text{obs}}|), \quad \text{CI}_{0.95}(\theta) = [Q_{0.025}\{\hat{\theta}^{*(b)}\}, Q_{0.975}\{\hat{\theta}^{*(b)}\}]. \quad (8)$$

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}^\top \mathbf{W} \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{W} \mathbf{z}, \quad R_k = \exp(\hat{\beta}_k), \quad (7)$$

where R_k is the multiplicative contrast for covariate k . Exact permutation and publication bootstrap calculations are

The extraction-and-scope sensitivity additionally propagates analyst-specified inclusion probabilities π_i and relative value uncertainties r_i :

$$I_i^{(b)} \sim \text{Bernoulli}(\pi_i), \quad y_i^{(b)} = y_i \exp(\epsilon_i^{(b)}), \quad \epsilon_i^{(b)} \sim N(0, [\log(1 + r_i)]^2). \quad (9)$$

CR2/Satterthwaite inference, multiplicity adjustment, and deletion checks provide additional diagnostics. These analyses describe the assembled corpus rather than causal or population relationships; the propagated uncertainty ranges are not empirically estimated confidence intervals.

2.7 Common-anchor growth comparison

Because the 2030 target has not expired, forecast accuracy cannot yet be evaluated. We instead compare each publication-level centre with the common IEA estimate of 485 TWh in 2025 (International Energy Agency, 2025, 2026) and calculate the annual growth required to reach each 2030 value:

$$g_i = \left(\frac{\tilde{y}_i}{485} \right)^{1/5} - 1. \quad (10)$$

The calculation is a scale and trajectory comparison, not independent validation. The IEA forecast is part of the analytical corpus, and differences in historical boundaries may remain.

2.8 Reproducibility

The accompanying replication package contains the coded records, source audit, search decisions, controlled rubric, generated tables, analysis scripts, fixed random seeds, software-environment capture, and automated integrity tests. It excludes copyrighted third-party full texts. Detailed sensitivity analyses, diagnostics, appraisal domains, and the

full auditable source table are provided in the Supplementary Material.

3. Results

3.1 Study selection and verification

The auditable flow begins with 19 candidate entries from the IEA-4E table, one additional IEA report, and six public-search records. One duplicate was removed from 26 identified records. Nine of 25 screened reports were excluded for incompatible geography, scope, target year, outcome, conversion, or source version; one remained pending exact extraction. The final synthesis contains 15 publications and 28 entered values (Figure 2). This flow describes the study's seeded candidate set, not the full universe of more than 50 publications considered by IEA-4E.

Seven included publications were directly checked by one reviewer. Eight were traceable to a primary report or publisher page but lacked complete page-level verification. Five publications were represented by endpoints without a recovered complete scenario set. All publications lacked sufficient information for a consistent independent validation and uncertainty appraisal; scenario recoverability and model transparency were the most frequent high-concern domains.




Figure 2: Study-selection flow for the secondary-source-seeded evidence map. The flow does not represent a rescreening of the complete IEA-4E inventory.




Figure 3: Aggregate reporting-appraisal profile across the 15 included publications. Counts are domain judgements rather than a composite quality score; the complete publication-level matrix is supplied in the Supplementary Material.

3.2 Primary distribution of publication centres

Treating endpoint-only reports as intervals gives a median latent publication centre of 912 TWh yr⁻¹, with a 95% publication-bootstrap interval of 712–1217 TWh (Table 2). The estimated geometric standard deviation is 1.69, indicating substantial multiplicative dispersion even after restricting the estimand.

Table 2: Primary interval-censored synthesis of publication centres.

Interval-censored estimand	Estimate	95% publication-bootstrap interval
Median publication centre (TWh)	912	712–1217
Geometric standard deviation	1.69	1.36–1.95

The conventional publication-level geometric summaries range from 411 to 2991 TWh. Their median is 945 TWh and their empirical interquartile range is 648–1167 TWh. The directly verified subset has the same median, with an interquartile range of 711–1452 TWh. These empirical distributions are not probability distributions for realised demand or for all forecasts that could have been published.




Figure 4: Entered 2030 projections. Lines show recovered low–high values and points show publication-level geometric scenario means. A point is a constructed descriptive summary unless the source explicitly identifies it as a central case.

3.3 Secondary variance and association results

The exploratory random-intercept model estimates between-publication and within-publication log variances of 0.18 and 0.28, respectively, but the ICC interval is 0.00–0.51 and the ordering reverses in the complete-scenario sensitivity. The data therefore do not establish whether scenario variation is systematically larger than publication variation.

Extrapolation publications have a point multiplier of 1.58 relative to the equipment-led or hybrid group, but the exact permutation result is $p = 0.151$ and CR2/Satterthwaite-style inference gives 0.67–3.70 with $p = 0.221$. The estimate falls to 1.28 after removing Andrae and Edler and to 1.16 under median regression. Publisher and publication-year contrasts are similarly imprecise after multiplicity adjustment. Accordingly, no forecast method, publisher type, or vintage is ranked as systematically high or low.




Figure 5: Secondary variance decomposition and empirical distribution of publication-level geometric scenario means. The variance comparison is descriptive because complete within-publication scenario sets are sparse.

3.4 Common-anchor growth comparison

From the 2025 IEA anchor, the median publication centre requires annual growth of 14.3% through 2030, with an empirical interquartile range of 5.9–19.2%. One publication-

level centre is already below the 2025 anchor. The comparison shows that forecasts with similar 2030 levels can imply very different trajectories depending on their historical base. It does not measure forecast error or provide independent confirmation of the IEA path.




Figure 6: Compound annual growth required to reach each publication-level 2030 centre from a common 485 TWh anchor in 2025. The dashed line shows the growth implied by the updated IEA path; it is a trajectory reference rather than a validation threshold.

4. Discussion

4.1 Principal findings

After applying a narrow common estimand, the literature still supports a broad range of 2030 values. The most

defensible central description is not a single consensus forecast but an interval-aware distribution of publication centres. The interval-censored median of approximately 912 TWh is close to the conventional 945 TWh publication median, yet its uncertainty and the wide empirical dispersion

caution against treating either number as a precise expected outcome.

The source audit is as important as the numerical synthesis. More than half of the included publications remain incompletely verified, five publications expose only range endpoints in the current extraction, and no publication permits a consistent independent validation judgment. Sophisticated modelling cannot recover assumptions that were not reported or source values that could not be checked. The sensitivity analyses therefore serve mainly to delimit what the corpus can support.

4.2 Relation to previous reviews

IEA-4E reports an envelope of approximately 210–7900 TWh across displayed 2030 scenarios (Kamiya and Coroama, 2025). The present evidence map asks a narrower question and excludes AI-only, national, network, capacity-only, and otherwise non-equivalent quantities. Its publication-level centres span 411–2991 TWh. The reduced descriptive spread does not invalidate the broader IEA-4E inventory; it demonstrates that part of the headline variation reflects different estimands and scenario architectures.

The corpus median is close to the included IEA 2025 Base Case (International Energy Agency, 2025). This proximity is not independent validation because the IEA observation contributes to the corpus. Earlier studies emphasizing efficiency gains and the imperfect link between digital service growth and electricity demand remain consistent with the present finding that assumptions and boundaries must stay visible (Masanet et al., 2020; Shehabi et al., 2018).

4.3 Boundary-reporting standard

The evidence map supports a minimum disclosure standard for future projections. Authors should report:

- 1) Geography and ownership classes;

- 2) IT-load, whole-facility, and grid-supply boundaries;
- 3) Treatment of AI, cryptocurrency, networks, and on-premises facilities;
- 4) Original quantities, units, and every annualisation or capacity conversion;
- 5) Base-year value and data source;
- 6) Workload, adoption, utilisation, hardware-efficiency, and PUE paths;
- 7) Verbatim scenario labels and the author's central case, if one exists; and
- 8) Machine-readable scenario outputs and versioned assumptions.

This standard is a reporting proposal, not yet a validated measurement instrument. Prospective evaluation should test whether independent coders using it obtain higher agreement, fewer unresolved conversions, and more reproducible numerical extractions than coders using unrestricted narrative descriptions.

4.4 Implications for planning and disclosure

For a demand case E_i , 2025 anchor E_0 , network-loss fraction λ_m , capacity factor f_m , and emission factor γ_m , the transparent planning translations are

$$\Delta E_i = \max(0, E_i - E_0), \Delta P_i = \frac{\Delta E_i}{8.76}, \quad (11)$$

$$G_{im} = \frac{\Delta E_i}{1 - \lambda_m}, K_{im} = \frac{G_{im}}{8.76 f_m}, M_{im} = G_{im} \gamma_m.$$

Here ΔP_i and K_{im} are GW, G_{im} is TWh yr⁻¹, and M_{im} is MtCO₂ yr⁻¹ when γ_m is tCO₂ MWh⁻¹. The empirical interquartile range implies approximately 19–78 GW of additional continuous average load. These are arithmetic translations, not probability-weighted planning cases; nameplate generation, emissions, transmission, and flexibility remain conditional on the assumed mix and regional constraints.




Figure 7: Illustrative translation of lower-quartile, median, and upper-quartile demand cases. Vertical ranges span the electricity-mix assumptions documented in the Supplementary Material and are not statistical confidence intervals

The more robust policy implication concerns disclosure. Regulators and system planners should require project developers and forecast publishers to distinguish energy from capacity, report workload and facility boundaries, preserve historical forecast versions, and release machine-readable assumptions. Global totals are useful for system-level scenario analysis, but local grid decisions require

region-specific pipelines, connection status, utilisation, and supply constraints.

4.5 Limitations

This is not an exhaustive or independently dual-screened systematic review. The candidate set was seeded by a

previous critical review, subscription-database searches were not executed, and one potentially eligible 2026 source remains unresolved. Only 7 of 15 included publications have single-reviewer direct verification; the full-corpus analysis therefore partly reflects secondary extraction.

The publication-level geometric mean is a constructed summary whose interpretation depends on scenario availability. The interval-censored model avoids treating range endpoints as observed scenarios but cannot identify within-publication scenario dispersion. Bootstrap intervals are wide, subset analyses are small and selected, and the association models involve few clusters. Analyst-specified extraction and inclusion probabilities are sensitivity assumptions rather than empirically estimated error rates.

The corpus is not a probability sample, so its median and quartiles do not estimate the distribution of every possible forecast. The common-anchor analysis does not establish accuracy because 2030 has not occurred and the anchor is not fully independent of the corpus. Finally, the proposed reporting standard has not yet undergone independent coder validation.

4.6 Research agenda

The next evidence update should execute the archived database strategies, directly verify all retained reports, and use two independent human reviewers for screening and extraction. Complete scenario sets would permit a cleaner comparison of within- and between-publication variation. Expired forecast vintages should be compared with scope-matched realised consumption using preregistered accuracy and calibration measures. The reporting standard should then be tested prospectively on newly published forecasts before it is presented as a validated instrument.

5. Conclusion

Global data-centre electricity projections cannot be reconciled by converting numbers to the same unit alone. Geography, facility scope, workloads, ownership, scenario status, model design, and conversion assumptions must also be aligned. In this evidence map, an interval-aware synthesis places the median publication centre near 912 TWh in 2030 but leaves substantial uncertainty and does not identify a probability distribution for realised demand. The durable contribution is therefore the auditable corpus and boundary-reporting framework, not a preferred point forecast or a ranking of model classes.

Data and code availability

The accompanying Replication Package contains the record-level data, primary-source audit, coding rubric, search protocol and logs, generated tables, study-selection counts, domain-level appraisal, common-anchor calculations, machine-readable results, software-environment capture, automated checks, and analysis code. Detailed audits, diagnostics, and sensitivity analyses are provided in the Supplementary Material. The archive excludes third-party full texts and copyrighted source documents. No persistent repository DOI or open reuse licence is claimed; these

require author and repository action before public deposition.

Declarations

- **Funding:** The author reports that this work received no external funding.
- **Competing interests:** The author declares no known financial or non-financial competing interests.
- **CRedit authorship statement:** Aditya Singh: Conceptualization, Methodology, Investigation, Data curation, Formal analysis, Software, Visualization, Writing- original draft, and Writing-review and editing.
- **Ethics approval:** Not applicable; the study uses published documents and contains no human-participant or personal data.

References

- [1] Anders Andrae. Comparison of several simplistic high-level approaches for estimating the global energy and electricity use of ict networks and data centers. *International Journal of Green Technology*, 5(1):50–63, 2019a. doi: 10.30634/2414-2077.2019.05.06.
- [2] Anders S. G. Andrae. Projecting the chiaroscuro of the electricity use of communication and computing from 2018 to 2030. Author manuscript, 2019b. URL https://www.researchgate.net/publication/331047520_Projecting_the_chiaroscuro_of_the_electricity_use_of_communication_and_computing_from_2018_to_2030.
- [3] Anders S. G. Andrae. New perspectives on internet electricity use in 2030. *Engineering and Applied Science Letters*, 3(2):19–31, 2020a. doi: 10.30538/psrp-easl2020.0038.
- [4] Anders S. G. Andrae. Hypotheses for primary energy use, electricity use and co2 emissions of global computing and its shares of the total between 2020 and 2030. *WSEAS Transactions on Power Systems*, 15:50–59, 2020b. doi: 10.37394/232016.2020.15.6.
- [5] Anders S. G. Andrae and Tomas Edler. On global electricity usage of communication technology: Trends to 2030. *Challenges*, 6(1):117–157, 2015. doi: 10.3390/challe6010117.
- [6] Joshua Aslan, Kieren Mayers, Jonathan G. Koomey, and Chris France. Electricity intensity of internet data transmission: Untangling the estimates. *Journal of Industrial Ecology*, 22(4):785–798, 2018. doi: 10.1111/jiec.12630.
- [7] Maria Avgerinou, Paolo Bertoldi, and Luca Castellazzi. Trends in data centre energy consumption under the european code of conduct for data centre energy efficiency. *Energies*, 10(10):1470, 2017. doi: 10.3390/en10101470.
- [8] Vitor Castro, Marios Georgiou, Thomas Jackson, Ian R. Hodgkinson, Lisa Jackson, and Steve Lockwood. Digital data demand and renewable energy limits: Forecasting the impacts on global electricity supply and sustainability. *Energy Policy*, 195: 114404, 2024. doi: 10.1016/j.enpol.2024.114404.
- [9] Alex de Vries. The growing energy footprint of artificial intelligence. *Joule*, 7(10):2191–2194, 2023. doi: 10.1016/j.joule.2023.09.004.

- [10] Deloitte Global. Powering artificial intelligence: A study of ai's environmental footprint—today and tomorrow. Technical report, Deloitte Global, November 2024. URL <https://www.deloitte.com/content/dam/assets-shared/docs/about/2024/powering-artificial-intelligence.pdf>.
- [11] Yee Van Fan, Charlie Wilson, George Kamiya, and Alessio Mastrucci. Data centre energy demand projections within shared socioeconomic pathways. *Energy and Climate Change*, 7:100253, 2026. doi: 10.1016/j.egycc.2026.100253.
- [12] Charlotte Freitag, Mike Berners-Lee, Kelly Widdicks, Bran Knowles, Gordon S. Blair, and Adrian Friday. The real climate and transformative impact of ict: A critique of estimates, trends, and regulations. *Patterns*, 2(9):100340, 2021. doi: 10.1016/j.patter.2021.100340.
- [13] Peter Henderson, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. Towards the systematic reporting of the energy and carbon footprints of machine learning. *Journal of Machine Learning Research*, 21(248):1–43, 2020. URL <https://jmlr.org/papers/v21/20-312.html>.
- [14] International Energy Agency. Energy and ai. Technical report, International Energy Agency, Paris, 2025. URL <https://www.iea.org/reports/energy-and-ai>.
- [15] International Energy Agency. Key questions on energy and ai. Technical report, International Energy Agency, Paris, 2026. URL <https://www.iea.org/reports/key-questions-on-energy-and-ai>.
- [16] Nicola Jones. How to stop data centres from gobbling up the world's electricity. *Nature*, 561(7722):163–166, 2018. doi: 10.1038/d41586-018-06610-y.
- [17] George Kamiya and Vlad C. Coroama. Data centre energy use: Critical review of models and results. Technical report, EDNA, IEA 4E Technology Collaboration Programme, March 2025. URL <https://www.iea-4e.org/wp-content/uploads/2025/05/Data-Centre-Energy-Use-Critical-Review-of-Models-and-Results.pdf>.
- [18] Jonathan Koomey and Eric Masanet. Does not compute: Avoiding pitfalls in assessing the internet's energy and carbon impacts. *Joule*, 5(7):1625–1628, 2021. doi: 10.1016/j.joule.2021.05.007.
- [19] Jonathan G. Koomey. Worldwide electricity used in data centers. *Environmental Research Letters*, 3(3):034008, 2008. doi: 10.1088/1748-9326/3/3/034008.
- [20] Martijn Koot and Fons Wijnhoven. Usage impact on data center electricity needs: A system dynamic forecasting model. *Applied Energy*, 291:116798, 2021. doi: 10.1016/j.apenergy.2021.116798.
- [21] Yanan Liu, Xiaoxia Wei, Jinyu Xiao, Zhijie Liu, Yang Xu, and Yun Tian. Energy consumption and emission mitigation prediction based on data center traffic and pue for global data centers. *Global Energy Interconnection*, 3(3):272–282, 2020. doi: 10.1016/j.gloi.2020.07.008.
- [22] Jens Malmmodin and Dag Lunden. The energy and carbon footprint of the global ict and e&m sectors 2010–2015. *Sustainability*, 10(9):3027, 2018. doi: 10.3390/su10093027.
- [23] Jens Malmmodin, Asa Moberg, Dag Lunden, Goran Finnveden, and Nina Lovehagen. Greenhouse gas emissions and operational electricity use in the ict and entertainment & media sectors. *Journal of Industrial Ecology*, 14(5):770–790, 2010. doi: 10.1111/j.1530-9290.2010.00278.x.
- [24] Jens Malmmodin, Nina Lovehagen, Pernilla Bergmark, and Dag Lunden. Ict sector electricity consumption and greenhouse gas emissions: 2020 outcome. *Telecommunications Policy*, 48(3):102701, 2024. doi: 10.1016/j.telpol.2023.102701.
- [25] Eric Masanet, Arman Shehabi, Nuo Lei, Sarah Smith, and Jonathan Koomey. Recalibrating global data center energy-use estimates. *Science*, 367(6481):984–986, 2020. doi: 10.1126/science.aba3758.
- [26] Karel G. M. Moons, Joris A. H. de Groot, Walter Bouwmeester, Yvonne Vergouwe, Susan Mallett, Douglas G. Altman, Johannes B. Reitsma, and Gary S. Collins. Critical appraisal and data extraction for systematic reviews of prediction modelling studies: The CHARMS checklist. *PLoS Medicine*, 11(10):e1001744, 2014. doi: 10.1371/journal.pmed.1001744.
- [27] David Mytton. Data centre water consumption. *npj Clean Water*, 4, 2021. doi: 10.1038/s41545-021-00101-w.
- [28] Renee Obringer, Benjamin Rachunok, Debora Maia-Silva, Maryam Arbabzadeh, Roshanak Nateghi, and Kaveh Madani. The overlooked environmental footprint of increasing internet use. *Resources, Conservation and Recycling*, 167:105389, 2021. doi: 10.1016/j.resconrec.2020.105389.
- [29] Chris Preist, Daniel Schien, and Eli Blevis. Understanding and mitigating the effects of device and cloud service design decisions on the environmental footprint of digital infrastructure. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 1324–1337, 2016. doi: 10.1145/2858036.2858378.
- [30] Arman Shehabi, Ben Walker, and Eric Masanet. The energy and greenhouse-gas implications of internet video streaming in the united states. *Environmental Research Letters*, 9(5):054007, 2014. doi: 10.1088/1748-9326/9/5/054007.
- [31] Arman Shehabi, Sarah J. Smith, Eric Masanet, and Jonathan Koomey. Data center growth in the united states: Decoupling the demand for services from electricity use. *Environmental Research Letters*, 13(12):124030, 2018. doi: 10.1088/1748-9326/aaec9c.
- [32] Christian Stoll, Lena Klaußen, and Ulrich Gellersdorfer. The carbon footprint of bitcoin. *Joule*, 3(7):1647–1661, 2019. doi: 10.1016/j.joule.2019.05.012.
- [33] Jess Tyndall. The AACODS checklist. Technical report, Flinders University, Adelaide, 2010. URL <https://fac.flinders.edu.au/dspace/api/core/bitstreams/e94a96eb-0334-4300-8880-c836d4d9a676/content>. Authority, Accuracy, Coverage, Objectivity, Date and Significance checklist for appraising grey literature.
- [34] Ward Van Heddeghem, Sofie Lambert, Bart Lannoo, Didier Colle, Mario Pickavet, and Piet Demeester. Trends in worldwide ict electricity consumption from 2007 to 2012. *Computer Communications*, 50:64–76, 2014. doi: 10.1016/j.comcom.2014.02.008.