International Journal of Scientific Research and Modern Education (IJSRME)
International Peer Reviewed - Refereed Research Journal, Website: www.crystalpen.in
Impact Factor: 7.137, ISSN (Online): 2455 - 5630, Volume 10, Issue 2, July - December, 2025

GEOGRAPHICAL ROUTING STRATEGIES AND THEIR IMPLICATIONS FOR WATER CONSUMPTION DURING DISTRIBUTED ARTIFICIAL INTELLIGENCE MODEL INFERENCE

Divya Vishwanath* & Paul Johnson Asamoah**
* Janardhan Rai Nagar Rajasthan Vidyapeeth, Udaipur, Rajasthan, India
** Power Life University College, Accra, Ghana, West Africa

Cite This Article: Divya Vishwanath & Paul Johnson Asamoah, "Geographical Routing Strategies and their Implications for Water Consumption During Distributed Artificial Intelligence Model Inference", International Journal of Scientific Research and Modern Education, Volume 10, Issue 2, July - December, Page Number 178-186, 2025.

Copy Right: © Crystal Pen Publication, 2025 (All Rights Reserved). This is an Open Access Article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

DOI: https://doi.org/10.5281/zenodo.24050541

Abstract:

This analysis quantifies how geographical routing priorities reshape water consumption in distributed artificial intelligence inference, addressing competing environmental and economic objectives. A deterministic comparison uses 4,570 destination and hour records across ten locations and 457 shared hours from 22 September to 11 October 2022, generating 1,371 routing outcomes under identical electricity allocations. Relative to water aware routing, carbon aware routing records 42.01 percent higher mean water consumption alongside 63.84 percent lower mean carbon emissions. Cost aware routing records 86.58 percent higher mean water consumption and 79.83 percent lower mean electricity expenditure. These differences arise because minimizing carbon intensity, water intensity or electricity price selects destinations with different resource coefficients. Hourly comparisons qualify the average ranking: cost aware routing consumes less water than carbon aware routing in 68 hours. The analysis refines resource allocation reasoning by demonstrating that geographical flexibility delivers objective dependent benefits. Its contribution is a reproducible matched comparison that separates resource intensity from electricity demand, supporting transparent assessment of water, carbon and expenditure consequences in international computing operations.

Key Words: Artificial Intelligence Inference; Geographical Routing; Water Consumption; Carbon Emissions; Matched Descriptive Analysis.

1. Introduction:

Where artificial intelligence inference is processed determines more than its electricity expenditure. Geographical routing connects computational demand to electricity systems and cooling conditions with different carbon and water intensities. A destination selected to minimize one burden may therefore increase another. This creates a practical decision problem for distributed computing: environmental performance depends on the allocation objective as well as the resources required to deliver the service. Carbon reduction, water conservation and electricity savings cannot be treated as interchangeable outcomes. Water consumption makes this distinction consequential. Data centres consume water through cooling and through the electricity generation supporting their operations. Mytton 2021 identifies these interconnected pathways and highlights measurement limitations that complicate comparisons between facilities. Operational water consumption must consequently be distinguished from electricity consumption, water withdrawal and broader ecological impact. A water coefficient expresses resource consumption per unit of electricity; it does not independently establish local scarcity or environmental damage. This analysis retains that distinction by examining modeled operational water consumption under explicitly standardized electricity allocations.

The underlying allocation problem arises because electricity price, carbon intensity and water intensity need not rank destinations identically. Carbon aware routing selects the destination with the lowest carbon coefficient, water aware routing selects the lowest water coefficient, and cost aware routing selects the lowest electricity price. Their water outcomes depend on the environmental characteristics of the selected destination. Macknick et al. 2012 show why alignment cannot be assumed: electricity generation and cooling technologies have different water requirements, and a transition towards lower carbon electricity can increase or decrease water consumption. The relevant theoretical expectation is therefore conditional convergence between objectives rather than a universal environmental advantage from optimizing any single criterion. Existing computational research establishes that location can materially alter operational emissions. Dodge et al. 2022 demonstrate the importance of geographical region and timing when assessing the carbon intensity of cloud computing. Spatial environmental accounting also shows why carbon considerations require a complementary water assessment. Siddik et al. 2021 quantify geographical differences in the carbon and water footprints of data centres and examine the tradeoffs associated with their location. Together, these contributions establish the importance of spatial resource variation, while motivating comparisons that preserve the distinct meanings of carbon emissions, water consumption and electricity expenditure.

The closest research directly addresses distributed inference. Yang et al. 2024 develop equity aware geographical load balancing to reduce environmental burdens in disadvantaged regions, using simulations involving ten geographically distributed data centres. Their work already establishes that allocation objectives can change the regional distribution of carbon and water footprints. Building on its public input traces, this analysis examines a narrower descriptive question: how do three transparent destination selection rules compare when each receives the same electricity allocation at the same hour? The contribution concerns the magnitude, distribution and geographical composition of those matched differences, rather than a new optimization algorithm or the first identification of environmental tradeoffs in inference routing. The distinction between electricity allocation and completed inference service is essential to interpreting this comparison. Luccioni et al. 2024 document substantial differences in inference energy requirements across model and task categories. A fixed electricity allocation therefore permits comparison of destination related resource intensity without establishing the footprint of an identical number of requests. This analysis holds allocated electricity at one kilowatt hour for each strategy and hour. Modeled water consumption then follows from the selected

destination's operational water coefficient, while the corresponding carbon coefficient and electricity price provide supporting resource and expenditure measures.

This study aims to quantify how geographical routing priorities shape modeled operational water consumption under matched hourly conditions. The research question asks how water consumption differs among carbon aware, water aware and cost aware routing across the available destinations. The analysis compares consumption levels and hourly differences, evaluates the accompanying carbon and expenditure outcomes, and identifies where destination selections coincide or diverge. These objectives use a balanced input panel containing 4,570 destination and hour records across ten locations and 457 shared hours from 22 September to 11 October 2022. Applying the three rules produces 1,371 calculated routing outcomes. The findings show that carbon aware routing achieves lower modeled carbon emissions with greater average water consumption than water aware routing, while cost aware routing achieves lower electricity expenditure with the highest average water consumption. However, the average ordering conceals hours in which cost aware routing uses less water than carbon aware routing. This combination of average differences and hourly reversals qualifies any fixed ranking of routing strategies. Its practical value is a reproducible comparison of the resource consequences accompanying each objective. The transferable insight concerns the need to evaluate routing choices across distinct resource measures; the numerical results remain conditional on the represented destinations, observation window and standardized allocation.

2. Theory, Conceptual Framework and Research Question Development:

Geographical routing connects computational demand to the environmental characteristics of the location supplying electricity and cooling. Its theoretical basis is resource allocation under competing objectives: a routing rule selects a destination according to one criterion, while the selected destination determines the associated resource burden. This analysis applies that logic to a fixed electricity allocation across alternative locations. Holding allocated electricity constant separates differences in environmental intensity from differences in computational scale. Masanet et al. 2020 demonstrate why this distinction matters: changes in computing demand do not translate proportionately into electricity consumption when infrastructure efficiency also changes. The conceptual framework therefore links geographical routing strategy to modeled operational water consumption through destination selection. Routing strategy is an observed categorical decision rule represented by three alternatives: carbon aware routing, water aware routing and cost aware routing. Each alternative selects one location from the same available destinations at the same hour. Modeled water consumption equals the selected location's water intensity multiplied by one allocated kilowatt hour. These alternatives represent mutually exclusive policies rather than independent continuous predictors. The framework describes a computational allocation sequence, with destination selection connecting the policy criterion to the calculated outcome.

Operational water consumption reflects the relationship between computing infrastructure, cooling requirements and electricity supply. A location can require water directly for cooling and indirectly through electricity generation. Consequently, electricity use alone does not identify the water burden associated with serving computational demand. Mytton 2021 explains these connected pathways and identifies measurement transparency as a constraint on evaluating data centre water consumption. This analysis uses the operational water coefficients supplied in the dataset rather than treating calculated consumption as direct metering of individual inference requests. Carbon aware routing selects the location with the lowest available carbon intensity during each hour. Its water implications depend on whether carbon intensity and water intensity identify the same destination. Where their rankings coincide, reducing modeled carbon emissions can also reduce water consumption. Where they diverge, the carbon criterion can select a destination with a higher water coefficient. Macknick et al. 2012 establish the technological basis for this ambiguity: electricity generation technologies and cooling configurations differ in their water requirements, and reducing carbon intensity can either increase or decrease water use. A universal directional relationship between carbon aware routing and water consumption is therefore unwarranted.

Water aware routing selects the destination with the lowest operational water coefficient during each hour. Given identical allocated electricity and the same destination set, this rule produces the lowest modeled water consumption by construction. The substantive descriptive comparison concerns the magnitude of the water difference relative to the other strategies and the carbon emissions associated with that selection. Siddik et al. 2021 provide spatial evidence that data centre location creates both tradeoffs and synergies between environmental footprints. Their findings justify examining carbon and water together without assuming that improvement in one establishes improvement in the other. Cost aware routing selects the destination with the lowest electricity price. Its mechanism is economic prioritization: the allocation rule responds to the recorded electricity price rather than the environmental coefficients. Water consumption consequently depends on the water intensity of the economically preferred destination. Price, carbon intensity and water intensity can identify the same location or different locations, so the water ranking of cost aware routing remains an empirical description of the available traces. This comparison concerns environmental intensity at fixed electricity demand. It does not measure changes in total inference demand, a distinction reinforced by de Vries 2023 in examining the energy requirements associated with expanding artificial intelligence deployment.

The available evidence comprises 4,570 location and hour input records and 1,371 calculated routing outcomes across ten locations and 457 shared hours in 2022. Each hour provides a matched comparison because all three strategies encounter the same destination set and receive the same electricity allocation. Carbon emissions and electricity expenditure provide supporting measures for interpreting the water comparison. The framework requires no latent constructs, composite indices or additional explanatory variables. The analysis addresses one descriptive research question: How does modeled operational water consumption differ among carbon aware, water aware and cost aware geographical routing strategies across matched hours? We resolve this question by comparing consumption levels and the size of the differences between strategies, while using their associated carbon emissions and electricity expenditure to interpret competing objectives. The interpretation remains conditional on the available locations, recorded hours and standardized electricity allocation.

3. Data:

This analysis establishes the provenance of the hourly environmental and electricity price inputs and documents their integration into a common analytical sample. Variable definitions and reproducible routing calculations connect geographical allocation decisions to modeled operational water consumption.

3.1 Data Sources, Population, and Sample Construction:

The dataset derives from the public Environmentally Equitable AI repository maintained by Li, Yang, Wierman and Ren. The repository contains location specific carbon intensity, operational water intensity and electricity price files. Its documentation attributes the underlying energy information to the United States Energy Information Administration, weather information to Iowa Environmental Mesonet and electricity prices to the International Energy Agency. These institutional datasets were accessed through the authors processed traces rather than downloaded independently. The repository carries an MIT licence; upstream data restrictions remain applicable. The attached workbook does not record the original extraction date or separate upstream release identifiers. Li et al. 2023 provide the source repository, while Yang et al. 2024 document the geographical allocation setting. The target population comprises geographical deployment opportunities for distributed artificial intelligence inference within the stated scope of 2010 to 2024. The accessible sampling frame comprises ten locations represented in the repository: Texas, Virginia, Georgia, Nevada, Germany, Belgium, the Netherlands, Denmark, Japan and Singapore. These locations combine subnational areas and countries and are treated as routing destinations rather than comparable national statistical units. All ten available destinations enter the analysis. Their selection reflects repository availability rather than probability sampling from the global data centre population.

The retained observation period extends from 22 September 2022 at 00:00 UTC to 11 October 2022 at 00:00 UTC, inclusive. The input dataset forms a balanced panel containing 457 hourly observations for each destination. Carbon and water variables are intensity coefficients attached to hourly timestamps, while electricity prices are hourly price observations. They are neither resource stocks nor cumulative consumption flows. The constructed dataset contains three policy outcomes for each hour, producing a matched descriptive structure in which alternative routing rules encounter identical contemporaneous inputs. Eligibility required a destination and timestamp to have carbon intensity, water intensity and electricity price information. The carbon and water files each contain 4,570 observations. The price files contain 4,714 observations, including 144 additional records without corresponding environmental coefficients. Table 1 distinguishes source records from constructed routing outcomes. The workbook records the inclusion rule but does not establish that it was registered before results were inspected.

Table 1: Sample construction

Stage	Units	Observations	Implemented Rule
Available source files	10 destinations across 30 files	13,854 component records	4,570 carbon, 4,570 water and 4,714 price records
Initial destination and hour frame	10 destinations	4,714 distinct keys	Union of timestamps available across the three input families
Excluded unmatched records	6 destinations affected	144 price records	Electricity prices without matching carbon and water coefficients
Integrated analytical inputs	10 destinations	4,570 complete records	Exact destination and UTC timestamp intersection
Constructed routing comparisons	3 strategies	1,371 outcomes	One destination selected for each strategy across 457 shared hours

3.2 Data Integration, Cleaning, and Quality Assurance:

Integration used destination and UTC timestamp as the joint key. Each source family contained one record per available key, supporting a one-to-one merge. The source spelling Belguim was harmonized to Belgium without changing numerical values. Timestamps were matched exactly, and no daily or annual aggregation was applied during integration. The resulting panel contains 457 observations for every destination, with no duplicate destination and hour keys. The audit identified no missing carbon, water or price values among the 4,570 retained records. Missingness arose at the integration boundary through the 144 additional price observations. These observations were excluded because corresponding environmental coefficients were unavailable. No zero substitution, interpolation, imputation or external matching was applied. Repeated timestamps across destinations represent the panel structure, while repeated timestamps across strategies represent alternative calculated decisions. Source carbon and water coefficients were retained without rescaling. Electricity prices remained in United States dollars per megawatt hour in the input dataset and were divided by 1,000 when calculating expenditure per allocated kilowatt hour. No additional currency conversion or inflation adjustment was introduced. Five negative electricity price observations were retained as recorded. Their sign does not create an undefined calculation because the analysis uses price levels rather than logarithms or percentage changes.

Quality checks covered key uniqueness, completeness, numerical units, positive environmental coefficients and agreement between selected destinations and routing criteria. Independent recomputation of all 1,371 routing outcomes reproduced the destination selections, water consumption, carbon emissions and electricity expenditure stored in the workbook. No minimum criterion ties occurred. The construction nevertheless specifies that a tie would select the first destination in the recorded source order. The analysis uses retrospective processed traces. The supplied documentation does not establish the historical publication timing or revision status of every upstream observation. Accordingly, contemporaneous allocation describes a comparison using recorded coefficients rather than an operational forecast based on information demonstrably available to a dispatcher. Distinguishing recorded water coefficients from direct consumption measurements is consistent with the measurement concerns identified by Mytton 2021.

3.3 Variable Measurement, Construction, and Validation:

Geographical routing strategy is a categorical allocation rule with three mutually exclusive alternatives. Carbon aware routing selects the lowest carbon coefficient at each hour, water aware routing selects the lowest water coefficient, and cost aware routing selects the lowest electricity price. Their binary indicators identify policy membership and do not represent three independently varying continuous characteristics. Modeled operational water consumption is the focal outcome. Equation 1 defines its implemented construction:

$$W_{st} = E_{st} \times w_{j_{st},t}$$

In Equation 1, W_{st} denotes modeled water consumption in litres for strategy s at hour t, E_{st} denotes allocated electricity in kilowatt hours, j_{st} denotes the destination selected by that strategy, and $w_{j_{st},t}$denotes the selected destination's operational water coefficient in litres per kilowatt hour. Allocated electricity equals one kilowatt hour for every strategy and hour. Consequently, numerical water consumption equals the selected coefficient, although the coefficient and calculated consumption retain different units. Carbon emissions equal allocated electricity multiplied by the selected carbon coefficient. Electricity expenditure equals allocated electricity multiplied by the selected price and divided by 1,000. These quantities support interpretation of the water comparison. Holding allocated electricity constant prevents differences in workload scale from determining the policy ranking, while leaving infrastructure efficiency and demand growth outside the calculated comparison. Masanet et al. 2020 establish the importance of distinguishing computational demand from electricity requirements.

Table 2: Variable dictionary

Variable or Indicator	Analytical Role	Definition and Construction	Source and Identifier	Unit and Frequency	Transformation and Direction	Missing Value Rule	Coverage and Limitation
Geographical routing strategy	Allocation variable	Carbon aware, water aware or cost aware destination selection	Routing dataset, strategy field	Category, hourly	Three mutually exclusive indicators	Constructed only for complete hours	1,371 calculated outcomes
Carbon intensity	Routing input	Recorded destination specific carbon coefficient	Repository carbon files, Carbon field	Kilograms per kilowatt hour, hourly	Unchanged; higher values indicate greater carbon intensity	Require complete matching inputs	4,570 processed trace observations
Operational water intensity	Routing input	Recorded destination specific operational water coefficient	Repository water files, WUE field	Litres per kilowatt hour, hourly	Unchanged; higher values indicate greater water intensity	Require complete matching inputs	4,570 modeled coefficients
Electricity price	Routing input	Recorded destination specific electricity price	Repository price files, Price field	United States dollars per megawatt hour, hourly	Divide by 1,000 for expenditure calculation	Exclude unmatched timestamps	4,570 retained values; negative prices preserved
Allocated electricity	Normalization input	One kilowatt hour assigned to each strategy and hour	Routing dataset, Allocated kWh field	Kilowatt hours, hourly	Fixed at one	No missing constructed values	Standardized allocation rather than measured inference demand
Selected destination	Construction output	Destination minimizing the applicable hourly criterion	Routing dataset, Selected location field	Category, hourly	Exact minimum selection	Require complete candidate set	Ten available destinations
Modeled water consumption	Focal outcome	Allocated electricity multiplied by selected water intensity	Routing dataset, Water consumption field	Litres, hourly	Higher values indicate greater modeled consumption	Complete constructed sample	1,371 outcomes; conditional on source coefficients
Modeled carbon emissions	Supporting measure	Allocated electricity multiplied by selected carbon intensity	Routing dataset, Carbon emissions field	Kilograms, hourly	Higher values indicate greater modeled emissions	Complete constructed sample	Same allocations as the water outcome
Electricity expenditure	Supporting measure	Allocated electricity multiplied by selected price divided by 1,000	Routing dataset, Electricity cost field	United States dollars, hourly	Higher values indicate greater electricity expenditure	Complete constructed sample	Excludes other service costs

Table 3 reports the distribution of retained inputs and the focal outcome. Standard deviations use the sample denominator. Carbon and water intensities are evaluated separately because electricity technologies can differ in their environmental rankings, as documented by Macknick et al. 2012.

Table 3: Descriptive statistics

Measure	Valid Observations	Mean	Standard Deviation	Median	Minimum	Maximum
Input carbon intensity, kilograms per kilowatt hour	4,570	0.3740	0.1106	0.3881	0.0947	0.5717
Input water intensity, litres per kilowatt hour	4,570	4.9685	1.2963	4.9599	1.3391	7.8924
Input electricity price, dollars per megawatt hour	4,570	161.7106	129.7570	142.0018	−65.5019	780.4130
Water consumption under carbon aware routing, litres	457	3.9028	0.2917	3.8982	3.0171	4.7707
Water consumption under water aware routing, litres	457	2.7482	0.6309	2.8436	1.3391	3.9225
Water consumption under cost aware routing, litres	457	5.1274	1.2626	5.3182	1.7756	7.6650

3.4 Analytical Framework, Identification, and Reproducibility:

The analysis is descriptive and compares modeled water consumption across alternative allocation rules at matched hours. Its principal quantities are policy specific consumption levels and differences between strategies under identical electricity allocations. The variation arises from destination specific coefficients and their changing hourly rankings. The three outcomes within an hour are related calculations, and consecutive hours form a time sequence rather than independent experimental replications. The implemented baseline consists of selecting destinations according to the three criteria and calculating resource outcomes from the corresponding source coefficients. It contains no fitted regression, additional controls, interaction terms or statistical inference procedure. Water aware routing produces the smallest water outcome within the available destination set by construction. The descriptive contribution therefore concerns the magnitude and distribution of differences and the associated carbon and expenditure outcomes. Interpretation remains conditional on unrestricted destination choice, the common hourly candidate set and one allocated kilowatt hour. The calculation does not incorporate transmission latency, capacity limits, network electricity, destination specific inference efficiency or differences in service quality. Spatial environmental comparisons are useful precisely because destination choices can produce competing resource outcomes, as demonstrated by Siddik et al. 2021; however, modeled consumption does not measure local ecological damage or water scarcity. Reproducibility rests on the pinned repository commit, source file paths, retained timestamps, documented location harmonization and editable workbook formulas. The independent recomputation verifies implementation consistency rather than validating the underlying environmental coefficients against external measurements. No alternative eligibility analysis or sensitivity specification is reported. The resulting evidence applies to the ten represented destinations and the observed hourly window.

4. Methods:

4.1 Research Design and Analytical Comparison:

This analysis uses a deterministic comparison of geographical routing strategies to quantify differences in modeled operational water consumption under matched hourly conditions. The unit of analysis is the routing strategy at a particular UTC hour. Each strategy selects one destination from the same ten available locations and receives one kilowatt hour of allocated electricity. Holding electricity allocation and observation time constant isolates the resource consequences of destination selection within the available traces. The comparison evaluates three decision rules: selecting the destination with the lowest carbon intensity, the lowest operational water intensity or the lowest electricity price. The analytical quantities are strategy specific consumption levels, matched hourly differences and the associated carbon emissions and electricity expenditure. The design answers how these quantities differ under alternative routing objectives. It does not estimate causal effects of observed deployment decisions. The approach builds on the public geographical load balancing setting documented by Yang et al. 2024. Its implementation uses unrestricted destination selection rather than reproducing their equity optimization model. No capacity, latency, network energy or service quality constraints enter the allocation rules. The calculated outcomes therefore describe standardized routing scenarios within the observed candidate set.

4.2 Routing Rules and Resource Calculations:

Equation 1 defines modeled operational water consumption as allocated electricity multiplied by the selected destination's operational water coefficient. Destination selection is specified in Equation 2:

$$j_{st} = \underset{j \in J}{\operatorname{argmin}}\ q_{sjt}$$

In Equation 2, j_{st} denotes the destination selected by strategy s at hour t, J denotes the common set of ten destinations, and q_{sjt} denotes the criterion evaluated for destination j. The criterion is carbon intensity for carbon aware routing, operational water intensity for water aware routing and electricity price for cost aware routing. All criteria use values recorded at the same hour. For each selected destination, carbon emissions equal allocated electricity multiplied by its carbon coefficient. Electricity expenditure equals allocated electricity multiplied by its electricity price divided by 1,000, converting dollars per megawatt hour to dollars per kilowatt hour. Water consumption, carbon emissions and electricity expenditure remain separate quantities. They are not normalized into a composite sustainability score. Allocated electricity equals one kilowatt hour in every scenario. This normalization permits comparison of geographical resource intensity but does not establish the footprint of an identical number of inference requests. Computational environmental accounting requires explicit separation of electricity requirements from the environmental coefficients applied to that electricity, consistent with Lannelongue et al. 2021. Selection uses the first minimum in the recorded destination order: Texas, Virginia, Georgia, Nevada, Germany, Belgium, the Netherlands, Denmark, Japan and Singapore. This provides a deterministic tie rule. No minimum criterion ties occur in the retained sample.

4.3 Sample Alignment and Descriptive Estimation:

The sample construction table and variable dictionary define the input coverage, eligibility rules and measurement units. The routing calculations use the balanced panel of 4,570 complete destination and hour records, containing 457 observations for each location. Each strategy is evaluated across every retained hour, producing 1,371 calculated outcomes. All policy comparisons therefore use an identical temporal sample and destination set. The sample represents the routing opportunities available in the source repository. It is not a probability sample of global data centres. Its combination of countries and subnational areas identifies candidate destinations rather than comparable national populations. No population weights or destination size weights are applied. Descriptive estimation uses arithmetic means, sample standard deviations, medians, minima and maxima. Every retained hour receives equal weight. Standard deviations describe observed dispersion and use the number of observations minus one as the denominator. Matched hourly water differences are defined in Equation 3:

$$D_{ab,t} = W_{at} - W_{bt}$$

In Equation 3, $D_{ab,t}$ denotes the water consumption difference between strategies a and b at hour t, while W_{at} and W_{bt} denote their respective modeled consumption values. Positive differences indicate greater consumption under strategy a. Comparisons cover carbon aware relative to water aware routing, cost aware relative to water aware routing, and cost aware relative to carbon aware routing. Relative differences between strategy means are calculated using Equation 4:

$$R_{ab} = 100\left[\frac{\overline{W}_a}{\overline{W}_b} - 1\right]$$

In Equation 4, R_{ab} denotes the percentage difference in mean water consumption relative to strategy b, and $\overline{W}_a$ and $\overline{W}_b$ denote the corresponding arithmetic means. The same calculation is applied separately to carbon emissions and electricity expenditure. These quantities are ratios of sample means rather than averages of hourly percentage differences. The analysis also counts hours with higher, equal or lower water consumption and tabulates selected destinations. Differences with an absolute magnitude no greater than 10^{-10} are classified as equal to avoid interpreting numerical rounding as substantive variation. Destination coincidence is assessed through exact agreement between the recorded location identifiers.

4.4 Computational Validation, Reproducibility and Interpretation:

The analytical sequence begins with the integrated destination and hour panel. For each hour, the procedure identifies the minimum criterion for each strategy, retrieves the corresponding destination coefficients, calculates resource outcomes and then constructs matched differences. Source coefficients remain unchanged. Negative electricity prices are retained, and no logarithmic transformation, trimming or winsorization is applied. Additional price observations without matching environmental coefficients are excluded before routing calculations; no imputation or interpolation enters the retained sample. The workbook implements selection and retrieval through MIN, MATCH and INDEX formulas. Summary calculations remain linked to the resulting routing records. Independent recomputation using Python's standard library checks each selected destination against the minimum source criterion and verifies water consumption, carbon emissions and electricity expenditure. Additional checks assess duplicate destination and hour keys, input completeness and equal temporal coverage. These checks evaluate construction consistency rather than independently validating the environmental coefficients. The workbook records source file paths, destination identifiers, retained timestamps, measurement definitions and editable calculations. The documented harmonization changes the source spelling Belguim to Belgium without altering numerical values.

Interpretation relies on descriptive magnitudes rather than sampling-based significance. Outcomes within an hour share the same inputs, and consecutive hours form a temporal sequence. Accordingly, the analysis does not treat the 1,371 calculated outcomes as independent experimental observations or report confidence intervals from that assumption. Water aware routing minimizes modeled water consumption by construction; the informative comparisons concern the size of the differences, their hourly variation and the accompanying resource outcomes. The validation establishes that the implemented rules reproduce the reported calculations. It does not demonstrate performance under alternative destination sets, operational constraints or environmental coefficient estimates. The methodological contribution is a transparent application of matched allocation comparisons, with conclusions confined to the represented routing opportunities and standardized electricity basis.

5. Findings:

This analysis compares modeled operational water consumption under three geographical routing strategies across matched hourly conditions. The findings quantify differences in water consumption, associated carbon emissions and electricity expenditure. Destination selections and hourly comparisons identify where the strategies converge and where their objectives produce competing resource outcomes. The analytical sample contains 4,570 complete destination and hour records covering ten locations over 457 shared hours from 22 September to 11 October 2022. Applying three routing strategies produces 1,371 calculated outcomes. Every strategy receives one kilowatt hour of allocated electricity at each hour, so differences reflect destination selection rather than differences in allocated electricity. The retained sample contains no missing routing outcomes. All comparisons use the same hours and available destinations. Table 4 presents the resource outcomes. Water aware routing records the lowest mean modeled water consumption at 2.7482 litres per allocated kilowatt hour. Carbon aware routing records 3.9028 litres, while cost aware routing records 5.1274 litres. The ordering differs for carbon emissions and electricity expenditure: carbon aware routing records the lowest mean carbon emissions, whereas cost aware routing records the lowest mean electricity expenditure. No strategy records the lowest sample mean across all three measures.

Table 4: Modeled resource outcomes by routing strategy

Measure	Carbon Aware Routing	Water Aware Routing	Cost Aware Routing
Shared hourly observations	457	457	457
Mean water consumption, litres	3.9028	2.7482	5.1274
Water standard deviation, litres	0.2917	0.6309	1.2626
Median water consumption, litres	3.8982	2.8436	5.3182
Minimum water consumption, litres	3.0171	1.3391	1.7756
Maximum water consumption, litres	4.7707	3.9225	7.6650
Mean carbon emissions, kilograms	0.1386	0.3834	0.3697
Mean electricity expenditure, United States dollars	0.2124	0.2198	0.0443

The comparison between carbon aware and water aware routing reveals a substantial separation between environmental objectives. Carbon aware routing consumes an additional 1.1546 litres per allocated kilowatt hour on average, equivalent to 42.01 percent above the water aware mean. Its mean carbon emissions are nevertheless 63.84 percent lower. Across the complete sequence of standardized allocations, carbon aware routing produces 1,783.57 litres of modeled water consumption and 63.35 kilograms of carbon emissions, compared with 1,255.90 litres and 175.20 kilograms under water aware routing. These totals describe the constructed allocation sequence rather than measured operational demand. The carbon and water comparison is therefore a resource tradeoff rather than a common improvement in environmental performance. This pattern is consistent with the technological distinction identified by Macknick et al. 2012, whose review establishes that lower carbon electricity can involve either higher or lower water requirements depending on generation and cooling technologies. The present calculations identify the direction and magnitude of that divergence within the available routing opportunities. Table 5 reports matched hourly water

differences. Carbon aware routing consumes more water than water aware routing in 443 hours and the same amount in 14 hours. Cost aware routing consumes more water than water aware routing in 434 hours and the same amount in 23 hours. These nonnegative differences follow from selecting the minimum water coefficient under water aware routing. Their informative content lies in their magnitude and frequency rather than in treating the minimum selection rule as an independently validated relationship.

Table 5: Matched hourly differences in modeled water consumption

Comparison	Mean Difference, Litres	Median Difference, Litres	Minimum Difference, Litres	Maximum Difference, Litres	Hours With Higher Consumption	Hours With Equal Consumption	Hours With Lower Consumption
Carbon aware relative to water aware	1.1546	1.1023	0.0000	2.7027	443	14	0
Cost aware relative to water aware	2.3793	2.5031	0.0000	5.0190	434	23	0
Cost aware relative to carbon aware	1.2246	1.4287	−2.5862	3.8303	339	50	68

Cost aware routing records mean water consumption 86.58 percent above water aware routing while reducing mean electricity expenditure by 79.83 percent. Its carbon emissions are only 3.56 percent lower than the water aware mean. Within these traces, prioritizing electricity price therefore produces a large expenditure difference alongside a substantial water difference and a comparatively small difference in average carbon emissions. This divergence supports evaluating water separately from electricity expenditure, consistent with Mytton 2021, who identifies distinct electricity supply and cooling pathways in data centre water consumption. The comparison between cost aware and carbon aware routing is less uniform at the hourly level. Cost aware routing consumes more water in 339 hours, the same amount in 50 hours and less water in 68 hours. Its average water consumption is 31.38 percent higher, but the negative minimum difference shows that this average ordering does not apply at every hour. The lowest price destination occasionally offers a lower water coefficient than the lowest carbon destination. Electricity expenditure and water consumption consequently exhibit conditional alignment rather than a fixed opposing relationship. Destination selections clarify the allocation pattern underlying these outcomes. Carbon aware routing selects Denmark in 444 hours and Nevada in 13 hours. Water aware routing selects the Netherlands in 269 hours, Japan in 174 hours and Denmark in 14 hours. Cost aware routing selects seven destinations, with Texas accounting for 244 selections. The strategies therefore differ in both their resource outcomes and the concentration of their selected destinations.

Table 6: Destination selections across the shared hourly sample

Destination	Carbon Aware Routing	Water Aware Routing	Cost Aware Routing
Belgium	0	0	7
Denmark	444	14	50
Germany	0	0	9
Japan	0	174	103
Netherlands	0	269	8
Nevada	13	0	0
Singapore	0	0	36
Texas	0	0	244
Georgia	0	0	0
Virginia	0	0	0
Total selections	457	457	457

Carbon aware and water aware routing select the same destination in 14 hours. Cost aware routing coincides with water aware routing in 23 hours and with carbon aware routing in 50 hours. These overlaps demonstrate that competing objectives sometimes identify a common destination, although most hourly selections differ. The geographical separation is consistent with the spatial tradeoffs documented by Siddik et al. 2021. Selection frequency alone, however, does not establish local water scarcity or ecological damage. Independent recomputation reproduces all 1,371 destination selections and their associated water, carbon and expenditure outcomes without discrepancies. No minimum criterion ties occur. This agreement establishes computational consistency. The standardized allocation also isolates environmental intensity from electricity demand, a distinction supported by Masanet et al. 2020 in their analysis of computing activity and infrastructure efficiency. The descriptive research question is resolved within the represented locations and hours: water aware routing minimizes modeled water consumption, carbon aware routing achieves lower carbon emissions with greater average water consumption, and cost aware routing achieves lower electricity expenditure with the highest average water consumption. The comparisons concern resource intensity at fixed electricity allocation. They do not establish total deployment footprints, which also depend on the scale and energy requirements of inference activity, as emphasized by de Vries 2023.

6. Discussion:

Reducing the carbon footprint of inference does not necessarily reduce its water footprint. Table 4 demonstrates this separation under identical electricity allocations: carbon aware routing achieves the lowest modeled carbon emissions while consuming more water than water aware routing. Equation 1 isolates the accounting relationship responsible for this result, since fixed allocated electricity leaves the selected destination's water coefficient to determine consumption. The contribution is a matched quantification of the resource consequences of competing routing objectives, rather than the discovery that environmental tradeoffs exist. Hadian and Madani 2015 established that evaluating energy systems through one environmental criterion can produce rankings that differ from broader sustainability assessments. This analysis makes that distinction operational for distributed inference by comparing the destinations selected under three explicit decision rules. Destination selection provides the computational link between routing priorities and modeled water consumption. Table 6 shows that carbon aware routing concentrates allocations in Denmark, whereas water aware routing predominantly selects the Netherlands and Japan. Combined with Equation 1, these selections explain how changing the decision criterion changes the water coefficient assigned to the same electricity allocation. This is a transparent allocation pathway rather than an estimated mediation effect. Ristic et al. 2015 locate data centre water consumption within interconnected electricity and infrastructure systems. The present comparison sharpens that perspective by showing that geographical flexibility acquires its environmental meaning through the criterion governing its use. Flexibility alone does not identify which resource burden will fall or whether several burdens will improve together.

The strongest distinction is between minimizing a resource directly and allowing its consumption to follow another objective. Table 5 shows that water aware routing never exceeds the alternatives in modeled water consumption, as required by its selection rule. More substantively, cost aware routing uses less water than carbon aware routing during some hours despite having higher average consumption. These reversals reveal that the relationship between economic and environmental rankings varies across available opportunities. Table 4 therefore describes an average ordering, while Table 5 identifies its limits. Lannelongue et al. 2021 emphasize explicit computational accounting when assessing environmental footprints. Here, the additional interpretive value comes from retaining matched hourly comparisons: average resource differences can coexist with individual hours in which objectives align, making a universal ranking of routing strategies inappropriate. The concentration of destination selections also changes how geographical optimization should be interpreted. Table 6 records a narrow set of preferred locations under each environmental objective, while cost aware routing draws on a broader destination set. These counts describe the distribution of modeled allocations; they do not establish infrastructure capacity, ecological pressure or the feasibility of concentrating actual traffic. Dodge et al. 2022 demonstrate that geographical region materially affects operational carbon intensity in cloud computing. The present results place a boundary around translating that finding into a general sustainability recommendation: a destination preferred for carbon can perform differently for water. The international relevance lies in this dependence on regional resource coefficients, rather than in treating the represented countries and subnational areas as interchangeable or representative of all deployment settings.

Routing decisions should therefore be assessed through their accompanying water, carbon and expenditure outcomes rather than through the optimized criterion alone. Tables 4 and 5 provide the relevant comparison: electricity savings under cost aware routing accompany greater average water consumption, while carbon savings under carbon aware routing also involve a water difference. Equation 1 keeps these conclusions conditional on fixed electricity allocation. Luccioni et al. 2024 show that inference energy requirements differ across model and task categories, so destination coefficients cannot by themselves establish the footprint of a completed service. The practical implication is to report the resource consequences of a routing choice alongside its intended benefit and the electricity basis used for comparison. This refines the allocation framework by treating routing performance as objective dependent, with the supported contribution confined to transparent comparisons of modeled operational resource use.

7. Conclusion and Implications:

Cleaner computation cannot be identified through a single resource objective. This analysis shows that prioritizing emissions, water conservation or electricity savings directs identical allocations towards different destinations, producing competing environmental and economic outcomes. Their alignment depends on the available geographical opportunities, with hourly reversals qualifying the average ranking of alternatives. The contribution is a transparent, matched comparison that makes these conditional differences visible without conflating resource intensity with computational demand. The findings refine resource allocation reasoning by showing that geographical flexibility has no uniform sustainability benefit independent of the criterion guiding its use. For managers, this means assessing the accompanying resource consequences of a preferred destination before treating its principal advantage as evidence of overall performance. Operational reporting should present water consumption, emissions and electricity expenditure on a common allocation basis while preserving their distinct meanings. Procurement and routing reviews can then identify where objectives coincide and where a preferred allocation requires an explicit tradeoff. For policymakers, the evidence supports comparable disclosure of operational resource coefficients and the assumptions underlying environmental claims. Such transparency can improve scrutiny of sustainability decisions across jurisdictions. Its societal relevance lies in making competing resource priorities accountable rather than allowing one favorable indicator to obscure another burden. The available evidence covers ten destinations over a limited hourly window and uses modeled coefficients with standardized electricity allocations. Longer observation periods and additional locations would establish whether the reported rankings and reversals persist across seasons and electricity systems. Integrating measuredReferences

References:

1. De Vries, A. 2023. The growing energy footprint of artificial intelligence. Joule, 7, issue 10, 2191 to 2194. https://doi.org/10.1016/j.joule.2023.09.004
2. Dodge, J., Prewitt, T., Tachet des Combes, R., Odmark, E., Schwartz, R., Strubell, E., Luccioni, A. S., Smith, N. A., DeCario, N., & Buchanan, W. 2022. Measuring the carbon intensity of AI in cloud instances. Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, 1877 to 1894. Association for Computing Machinery. https://doi.org/10.1145/3531146.3533234

3. Hadian, S., & Madani, K. 2015. A system of systems approach to energy sustainability assessment: Are all renewables really green? Ecological Indicators, 52, 194 to 206. https://doi.org/10.1016/j.ecolind.2014.11.029
4. Lannelongue, L., Grealey, J., & Inouye, M. 2021. Green algorithms: Quantifying the carbon footprint of computation. Advanced Science, 8, issue 12, Article 2100707. https://doi.org/10.1002/advs.202100707
5. Li, P., Yang, J., Wierman, A., & Ren, S. 2023. Environmentally Equitable AI. Source code and data, commit 182da36005763374f072c74791b18cd9fd11d724. GitHub. Pinned source repository
6. Luccioni, A. S., Jernite, Y., &Strubell, E. 2024. Power hungry processing: Watts driving the cost of AI deployment? Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, 85 to 99. Association for Computing Machinery. https://doi.org/10.1145/3630106.3658542
7. Macknick, J., Newmark, R., Heath, G., & Hallett, K. C. 2012. Operational water consumption and withdrawal factors for electricity generating technologies: A review of existing literature. Environmental Research Letters, 7, issue 4, Article 045802. https://doi.org/10.1088/1748-9326/7/4/045802
8. Masanet, E., Shehabi, A., Lei, N., Smith, S., &Koomey, J. 2020. Recalibrating global data center energy use estimates. Science, 367, issue 6481, 984 to 986. https://doi.org/10.1126/science.aba3758
9. Mytton, D. 2021. Data centre water consumption. npj Clean Water, 4, issue 1, Article 11. https://doi.org/10.1038/s41545-021-00101-w
10. Ristic, B., Madani, K., & Makuch, Z. 2015. The water footprint of data centers. Sustainability, 7, issue 8, 11260 to 11284. https://doi.org/10.3390/su70811260
11. Siddik, M. A. B., Shehabi, A., & Marston, L. 2021. The environmental footprint of data centers in the United States. Environmental Research Letters, 16, issue 6, Article 064017. https://doi.org/10.1088/1748-9326/abfba1
12. Yang, J., Li, P., Ren, S., & Wierman, A. 2024. Towards environmentally equitable AI via geographical load balancing. Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems, 291 to 307. Association for Computing Machinery. https://doi.org/10.1145/3632775.3661938