

Ambient Conditions and the Projection of Full-Load Power Output

Abstract

The question concerns full-load power output association with ambient conditions. The secondary ordinary least squares analysis contains 9,568 observations and uses HC3 covariance. The coefficient on AT is -1.9775, with a 95 percent interval from -2.0092 to -1.9458. The full-load restriction defines the support of this projection. Contextual interpolation can describe how output relates to ambient conditions inside that support, while applications at other loads remain extrapolation. Kriging and contextual selection suggest where additional evidence could be useful, but are not tested by the present OLS fit or its in-sample explanatory power. The estimates describe the released collection and its conditioning set. No new data collection, simulator experiment, or deployment-policy evaluation is claimed.

The outcome and its observed population

The full-load restriction defines the support of this projection. Contextual interpolation can describe how output relates to ambient conditions inside that support, while applications at other loads remain extrapolation. Kriging and contextual selection suggest where additional evidence could be useful, but are not tested by the present OLS fit or its in-sample explanatory power.

The power-plant file contains 9,568 full-load records with ambient temperature, exhaust vacuum, atmospheric pressure, humidity, and net hourly electrical output. The repository describes observations from one combined-cycle plant and offers shuffled versions for benchmarking. The normalized CSV is used once; repeated shuffles are not concatenated as new observations. The released file does not provide a verified chronological index for a prospective time split. The analysis therefore reports within-file associations and explicitly avoids describing row order as operating time. HC3 intervals address heteroskedastic working-model errors while unobserved temporal dependence remains possible. Temperature curvature can be examined without measuring the causal effect of a plant control intervention. Full-load support also does not establish behavior at other load regimes.

Specification and conditional uncertainty

The prespecified main specification is $PE \sim AT + V + AP + RH$. Estimation uses ordinary least squares with HC3 covariance. The table reports coefficients, robust standard errors, and two-sided 95 percent intervals on the outcome's declared scale. The same retained observations determine the fitted projection and its within-file explanatory statistic. These are not prospective policy scores. Full coefficient tables, including categorical controls, are saved separately so that the selected terms shown in the manuscript can be interpreted within the complete specification. HC3 covariance is a leverage-sensitive heteroskedasticity-consistent calculation. It retains the fitted coefficients and does not repair omitted variables, nonrepresentative collection, or unsupported extrapolation. Where the file contains designed simulation outputs, the interval is read as a working-model diagnostic rather than population sampling or physical reliability. Logged outcomes or predictors retain their stated transformations. A coefficient per log unit is not a constant effect per raw unit. A recorded physical unit and a unitless projection statistic remain different quantities, and no uncorrected exponentiation is used to claim an average raw outcome. Polynomial and interaction terms are interpreted jointly: a main coefficient is not a constant marginal slope when its predictor also enters a square or interaction. The reported table keeps those additional terms explicit. Scientific notation is used for small nonzero quantities so that rounding does not make a standard error or confidence interval look exactly zero.

Observed estimates and outcome checks

The fitted specification uses 9,568 observations. Its reported R-squared is 0.9287. This within-file statistic summarizes the declared projection and does not evaluate a resource policy, a selected extreme, or a new deployment population. The coefficient on AT is -1.9775, with a 95 percent interval from -2.0092 to -1.9458. The interval is retained on its original scale. If it contains zero, the specification does not precisely distinguish its sign at this confidence level; if it excludes zero, the conditional association still does not identify an intervention effect. The other selected terms below have their own intervals, so their uncertainty is not inferred from the first term's result.

The accompanying data and outcome checks record raw n: 9568; no repeated shuffle concatenation: True. These definitions identify the sample and comparison represented by the estimates; they are not additional causal or deployment results.

Table 1. Declared terms from the completed secondary specification; complete coefficient tables retain all controls.

Term	Estimate	Robust SE	95% interval
AT	-1.9775	0.0162	-2.0092 to -1.9458
V	-0.2339	0.0072	-0.2481 to -0.2197
AP	0.0621	0.0092	0.0440 to 0.0801
RH	-0.1581	0.0041	-0.1661 to -0.1500

On the Convergence Rates of Stochastic Kriging Assisted-Simulation with Covariates

Stochastic kriging with covariates studies prediction from simulation observations collected before the eventual application context becomes known. The role of random covariates makes the prediction problem more demanding than interpolation over a fixed list of training locations. Offline simulation evidence must support an application context that will be observed later. Prediction uncertainty and the distribution of those future contexts interact: an error in a common region affects many decisions, while an error in a rare region may have unusually serious consequences. Rate statements should be interpreted with their coverage assumptions intact. (C. Li et al., 2023).

The contextual implication is to ask where the preferred alternative changes, rather than merely where an expected response changes. A validation should retain the context available before action and distinguish direct local observations from model-based borrowing. The source motivates this distinction but does not identify the preferred action in every new population. A local decision boundary remains an object to measure.

For the present question of dependence and the unit of information, this study supplies the starting comparison. Its strongest contribution is the explicit decision object in “On the Convergence Rates of Stochastic Kriging Assisted-Simulation with Covariates.” The next step would be to identify which observations can change that object at the operating budget. Interpolation accuracy and coverage of future covariate values must be assessed together. This is an implication for evaluation design rather than a result obtained by applying the source method to a new system.

Contextual Ranking and Selection with Gaussian Processes

The role of “Contextual Ranking and Selection with Gaussian Processes” in this argument is to qualify how the initial comparison should be interpreted. Under dependence and the unit of information, its information model deserves separate scrutiny. Covariance assumptions deserve stress tests wherever sparse contexts borrow most heavily from abundant ones. A later comparison should retain that distinction and assess the resulting action on evidence that did not determine the preferred estimate. Otherwise methodological sophistication can hide an unresolved application assumption.

Gaussian-process contextual ranking and selection connects alternative comparison with correlations across application contexts. The conference treatment of contextual selection develops the use of Gaussian-process information in choosing alternatives by context. It is useful as a statement of the contextual decision structure. Borrowing across locations changes both prediction and sampling priorities. A comparison with a noncontextual rule should measure where choices differ and whether that difference matters, rather than assuming that a more flexible representation must improve every operating region. (S. Cakmak et al., 2021).

For the second part of the comparison, information availability is decisive. A conditional rule can look attractive when evaluated with retrospectively available variables. Its operational interpretation requires those variables to be observed in time. Missingness and measurement error can change both the apparent context and the chosen alternative. The representation should therefore be checked against the information process of the intended application.

From an estimator to a defensible action

For “On the Convergence Rates of Stochastic Kriging Assisted-Simulation with Covariates,” a useful methodological question follows from dependence and the unit of information: interpolation accuracy and coverage of future covariate values must be assessed together. The power-plant file contains 9,568 full-load records with ambient temperature, exhaust vacuum, atmospheric pressure, humidity, and net hourly electrical output. The repository describes observations from one combined-cycle plant and offers shuffled versions for benchmarking. The normalized CSV is used once; repeated shuffles are not concatenated as new observations. The released file does not provide a verified chronological index for a prospective time split. The analysis therefore reports within-file associations and explicitly avoids describing row order as operating time. HC3 intervals address heteroskedastic working-model errors while unobserved temporal dependence remains possible. Temperature curvature can be examined without measuring the causal effect of a plant control intervention. Full-load support also does not establish behavior at other load regimes. (Tfekci and Kaya, 2014).

The requirement in “Contextual Ranking and Selection with Gaussian Processes” is sharpened by this principle: covariance assumptions deserve stress tests wherever sparse contexts borrow most heavily from abundant ones. Heteroskedasticity-consistent covariance changes the estimate of sampling uncertainty while retaining the fitted regression coefficients. This is useful when error variability differs across observed conditions. It does not address omitted variables, reverse causation, unsupported extrapolation, or dependence that needs a different covariance treatment. A robust interval should consequently be read as conditional on the selected model and observation assumptions. The useful practice is to report the response scale, the conditioning variables, and the role of the estimate in a decision, rather than allowing the word robust to imply a causal or universally reliable result. (White, 1980).

One way to examine dependence and the unit of information is to preserve the boundary identified for “On the Convergence Rates of Stochastic Kriging Assisted-Simulation with Covariates”: interpolation

accuracy and coverage of future covariate values must be assessed together. Cluster-robust inference treats observations within a declared group as potentially dependent. The group should follow the collection process, such as hours within a date or segments within a response. An arbitrary cluster choice can give an interval an unjustified interpretation. The effective information for some comparisons comes from the number of groups rather than the number of rows. Cluster adjustment does not turn convenience sampling into population representativeness, and it does not establish a causal effect. It supplies one part of an uncertainty calculation under a specified dependence structure. (Cameron and Miller, 2015).

For “Contextual Ranking and Selection with Gaussian Processes,” a useful methodological question follows from dependence and the unit of information: covariance assumptions deserve stress tests wherever sparse contexts borrow most heavily from abundant ones. Gaussian-process prediction uses a covariance function to express relationships among input locations. In a selection application those relationships determine how measurements of one configuration influence the estimated attractiveness of another. A smooth response surface can support useful borrowing, but abrupt feasibility changes or unrepresented regimes can make that borrowing misleading. The model should be assessed at the locations relevant to choice, not only on a random collection of average errors. Posterior uncertainty is conditional on the representation and should not be described as protection against every form of model discrepancy. (Rasmussen and Williams, 2006).

The requirement in “On the Convergence Rates of Stochastic Kriging Assisted-Simulation with Covariates” is sharpened by this principle: interpolation accuracy and coverage of future covariate values must be assessed together. Stochastic kriging distinguishes randomness in simulation observations from uncertainty about the response surface. The distinction is valuable when a fitted surface guides selection: repeated runs can reduce the first uncertainty at an observed configuration without fully resolving the second between configurations. A decision-focused validation should therefore examine both replication and spatial coverage. One can ask whether the model is precise because it has direct observations nearby or because its covariance assumptions imply precision. This question does not require abandoning a surrogate; it requires evaluating the kind of evidence the surrogate supplies. (Ankenman et al., 2010).

One way to examine dependence and the unit of information is to preserve the boundary identified for “Contextual Ranking and Selection with Gaussian Processes”: covariance assumptions deserve stress tests wherever sparse contexts borrow most heavily from abundant ones. Indifference-zone selection specifies a performance separation under which a correct-selection requirement is meaningful. The zone represents the precision of the decision task, rather than a promise to resolve arbitrarily small differences with a finite amount of evidence. In an application the separation needs an operational interpretation. Evaluating a rule only on comfortably separated alternatives can hide its cost near the boundary, while demanding exact distinction among practical ties can overstate the need for measurement. The useful comparison reports the zone together with the candidate gaps and the intended error protection. (Kim and Nelson, 2001).

For “On the Convergence Rates of Stochastic Kriging Assisted-Simulation with Covariates,” a useful methodological question follows from dependence and the unit of information: interpolation accuracy and coverage of future covariate values must be assessed together. Knowledge-gradient reasoning assigns value to a measurement through the improvement it is expected to produce in a later decision. It connects Bayesian learning to a choice about which information to purchase next. This makes the prior and observation model consequential: they determine how an unobserved outcome would update beliefs and affect the action. A practical evaluation should record the realized sequence and test sensitivity to those assumptions. A posterior value assigned to information is a model-based decision quantity, not direct evidence that the resulting sequence outperforms every alternative allocation. (Frazier et al., 2008).

The requirement in “Contextual Ranking and Selection with Gaussian Processes” is sharpened by this principle: covariance assumptions deserve stress tests wherever sparse contexts borrow most heavily from abundant ones. Robust optimization makes the cost of protection visible by comparing decisions under nominal and adverse conditions. The uncertainty description determines which departures are defended against, so conservatism should not be assessed in the abstract. A design can be robust within one description and fragile to a departure excluded from it. Applications can examine how the preferred action changes as protection is increased and identify the constraints responsible for the change. This exposes the tradeoff between nominal performance and the chosen protection without claiming that one uncertainty set is universally correct. (Bertsimas and Sim, 2004).

The information boundary of the comparison

The comparison between “On the Convergence Rates of Stochastic Kriging Assisted-Simulation with Covariates” and “Contextual Ranking and Selection with Gaussian Processes” turns on the following distinction. Interpolation accuracy and coverage of future covariate values must be assessed together. Covariance assumptions deserve stress tests wherever sparse contexts borrow most heavily from abundant ones. These qualifications identify different reasons a recommendation can fail. An assessment of dependence and the unit of information should preserve each reason rather than combine them through a generic accuracy score. The useful common outcome is a declared action evaluated on independent information at a counted resource limit. The common outcome may require new evidence not observed in one of the original settings. That gap is a prospective requirement rather than information supplied by the literature comparison itself.

A formal way to retain this distinction is to write the expected outcome of alternative a in context x as $\mu(a, x)$, and to define the admissible alternatives by the requirements in that context. A prediction model estimates a relationship; a selection procedure uses estimated relationships to recommend an action. The selected object can instead be a retained group or an acceptable set. Its loss depends on the consequence of an exclusion, acceptance, or wrong commitment. These objects should be stated before a budget is optimized. An accurate coefficient or a precise posterior does not define the loss, the requirements, or the operating population on its own.

The completed secondary analysis supplies an outcome definition and a transparent conditional projection. It has not implemented either reviewed selection procedure. A future comparison could use the recorded relationships to motivate scenarios, while keeping a separate model of how an alternative changes the outcome. Independent evaluation of that change would then assess the action rather than merely the statistical fit. The costs of acquiring new information, fitting the model, applying the decision rule, and checking its selected alternative belong in the same accounting boundary. Without those action outcomes, the present result remains evidence about the released collection.

Interpretation of the completed result

The full-load restriction defines the support of this projection. Contextual interpolation can describe how output relates to ambient conditions inside that support, while applications at other loads remain extrapolation. Kriging and contextual selection suggest where additional evidence could be useful, but are not tested by the present OLS fit or its in-sample explanatory power. The useful interpretation keeps the recorded association separate from a chosen action. A model can be informative about context or a diagnostic event while leaving the consequence of changing a configuration unidentified. It can also fit the existing collection well while performing poorly in a sparse region selected by an optimizer. Independent action outcomes would be needed to evaluate that wider claim. The present coefficients are therefore carried into the methodological comparison on their own scale, without imputed costs, fabricated intervention gains, or transferred theoretical guarantees.

Testing the remaining decision uncertainty

A study should identify the dependence structure from the collection process before choosing a covariance calculation. It can compare row-level and group-level resampling as a sensitivity, while explaining the assumptions of each. The selected action should be evaluated over independent groups when that matches the operating task. More observations within a group can improve its characterization, but should not be described as the same increase in population coverage as acquiring new groups. Unknown batch structure remains a limitation rather than a reason to invent convenient clusters.

For a prospective study, contexts would be divided by operating meaning rather than by an arbitrary row split. The analyst would fix the alternatives, estimate the conditional comparisons using training evidence, and preserve an independent set of contexts for final evaluation. Important strata would include locations with direct observations and locations supported mostly by borrowing. Context-specific action disagreement, feasibility, and realized loss would accompany aggregate performance. If a rare stratum has too little information for a precise estimate, that limitation would remain visible instead of being hidden by a smooth fitted surface. The protocol would record when each covariate becomes available and would treat newly acquired application outcomes as part of a later update policy. That policy would need its own cost and outcome assessment.

Conclusion

The released evidence supports the declared secondary projection and its observed outcome definition. The full-load restriction defines the support of this projection. Contextual interpolation can describe how output relates to ambient conditions inside that support, while applications at other loads remain extrapolation. Kriging and contextual selection suggest where additional evidence could be useful, but are not tested by the present OLS fit or its in-sample explanatory power. The next empirical step is an independent check of the action and population that the reviewed computational methods would target. That step should count the information and verification it requires rather than treating within-file fit as the completed selection experiment.

References

1. C. Li, S. Gao, and J. Du, "On the Convergence Rates of Stochastic Kriging Assisted-Simulation with Covariates," *INFORMS Journal on Computing*, vol. 35, no. 2, pp. 386-402, 2023.
2. S. Cakmak, E. Zhou, and S. Gao, "Contextual Ranking and Selection with Gaussian Processes," *Proceedings of the 2021 Winter Simulation Conference*, pp. 1-12, 2021.
3. Pnar Tfekci, Heysem Kaya. *Combined Cycle Power Plant*. UCI Machine Learning Repository, 2014. Dataset. <https://doi.org/10.24432/C5002N>.
4. White, Halbert. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica*, vol. 48, no. 4, 1980, pp. 817-838. <https://doi.org/10.2307/1912934>.
5. Cameron, A. Colin, and Douglas L. Miller. "A Practitioner's Guide to Cluster-Robust Inference." *Journal of Human Resources*, vol. 50, no. 2, 2015, pp. 317-372. <https://doi.org/10.3368/jhr.50.2.317>.
6. Rasmussen, Carl Edward, and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006. <https://gaussianprocess.org/gpml/>.
7. Ankenman, Bruce, Barry L. Nelson, and Jeremy Staum. "Stochastic Kriging for Simulation Metamodeling." *Operations Research*, vol. 58, no. 2, 2010, pp. 371-382. <https://doi.org/10.1287/opre.1090.0754>.
8. Kim, Seong-Hee, and Barry L. Nelson. "A Fully Sequential Procedure for Indifference-Zone Selection in Simulation." *ACM Transactions on Modeling and Computer Simulation*, vol. 11, no. 3, 2001, pp. 251-273. <https://doi.org/10.1145/502109.502111>.
9. Frazier, Peter I., Warren B. Powell, and Savas Dayanik. "A Knowledge-Gradient Policy for Sequential Information Collection." *SIAM Journal on Control and Optimization*, vol. 47, no. 5, 2008, pp. 2410-2439. <https://doi.org/10.1137/070693424>.
10. Bertsimas, Dimitris, and Melvyn Sim. "The Price of Robustness." *Operations Research*, vol. 52, no. 1, 2004, pp. 35-53. <https://doi.org/10.1287/opre.1030.0065>.