BASELINE-CALIBRATED ADAPTIVE GAUSSIAN PROCESSES FOR CONDITIONAL DOSE–RESPONSE ESTIMATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Estimating conditional dose-response functions from observational data requires capturing heterogeneous treatment responses while accounting for covariate-driven outcome variation. We propose Baseline-Calibrated Adaptive Gaussian Processes (BCA-GP), a structured Gaussian process estimator that combines an unpenalized covariate baseline with an adaptive dose kernel. The method learns dose directions from baseline-residualized outcomes and retains a common baseline space throughout direction learning, hyperparameter selection, and prediction. This construction prevents additive signals within the baseline space from altering dose adaptation. We establish exact baseline reproduction and equivariance of the complete fitting procedure: adding a baseline-space function to the outcomes produces the corresponding shift in predictions while preserving response-surface estimation error. We further characterize prediction as constrained kernel approximation and derive a deterministic error bound that separates function approximation, representation error, and adaptive noise propagation. Under standard causal identification assumptions, BCA-GP directly estimates conditional intervention responses. Experiments on semi-synthetic benchmarks evaluate conditional and average response accuracy, isolate the contributions of dose adaptation and baseline adjustment, and verify the predicted stability under controlled baseline shifts.

1 INTRODUCTION

Continuous-treatment decisions require understanding how outcomes vary with dose and individual characteristics. The conditional dose-response function (CDRF) captures this heterogeneity, whereas population averages can obscure differences between individuals. Estimation from observational data is challenging because each subject contributes an outcome at only one dose, and treatment assignment may depend on covariates. Generalized propensity-score methods provide adjustment strategies for nonbinary treatments (Hirano & Imbens, 2004; Imai & van Dyk, 2004), while doubly robust methods combine treatment and outcome models to estimate average dose-response curves (Kennedy et al., 2017). Our target is the conditional intervention mean, identified through outcome regression under consistency, conditional exchangeability, and adequate support.

Learning individual response curves requires sharing information across subjects and doses. Neural approaches use dose-specific architectures (Schwab et al., 2020; Nie et al., 2021), generative counterfactual models (Bica et al., 2020), and attention-based treatment-covariate interactions (Zhang et al., 2023). Complementary strategies address observational-data limitations through end-to-end balancing (Bahadori et al., 2022) and counterfactual supervision based on gradient interpolation and kernel smoothing (Nagalapatti et al., 2024). Gaussian processes (GPs) express information sharing through covariance kernels and accommodate explicit mean functions (Rasmussen & Williams, 2006). Kernel-based causal response estimators (Singh et al., 2024), observational dose-response GP models (Dailey, 2024), and learned kernel representations (Wilson et al., 2016) provide foundations for adapting this structure to heterogeneous responses.

A remaining design question is how covariate-driven baseline variation influences the learned dose structure. Adding a covariate-dependent offset at every dose changes a response curve's level but not its shape. An adaptive estimator may nevertheless respond by selecting a different dose

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

representation or regularization strength. Work on regularization-induced confounding motivates separating prognostic and treatment-related components (Hahn et al., 2020). For an outcome-adaptive GP, however, this separation must extend beyond the final predictor to the construction and selection of its kernel.

We introduce Baseline-Calibrated Adaptive Gaussian Processes (BCA-GP), which combine an unpenalized covariate baseline with a learned dose kernel. Nyström features (Williams & Seeger, 2001) define a finite nonlinear baseline space. A preliminary regression adjusts both outcomes and covariate-dose interactions for this space, then extracts dose directions from the fitted spline coefficients. These directions define candidate kernels, which are compared using a baseline-adjusted profile score. The selected GP estimates the complete response while retaining the same baseline throughout fitting.

Our contributions are:

- We propose BCA-GP for conditional dose-response estimation, enabling flexible dose similarity while separating baseline variation from dose adaptation.
- We prove that adding a baseline-space function leaves the learned dose directions and selected hyperparameters unchanged. Predictions shift by exactly that function, preserving error relative to the correspondingly shifted response surface.
- We characterize prediction as constrained kernel approximation and derive a deterministic CDRF error bound separating approximation, representation error, and adaptive noise propagation.

Experiments on five semi-synthetic benchmarks compare response accuracy against six competing methods, while component ablations and controlled baseline shifts assess the model's design choices and stability.

2 Related work

Continuous-treatment estimation encompasses propensity-based adjustment (Hirano & Imbens, 2004), doubly robust estimation (Kennedy et al., 2017), and kernel estimators of causal response functions (Singh et al., 2024; Kim & Wang, 2026). Neural approaches provide flexible conditional surfaces and counterfactual supervision (Schwab et al., 2020; Nie et al., 2021; Zhang et al., 2023; Nagalapatti et al., 2024). Beyond these modeling choices, an important question is how auxiliary components influence response learning. Double/debiased machine learning uses Neyman-orthogonal scores and cross-fitting to reduce sensitivity to nuisance estimation (Chernozhukov et al., 2018). BCA-GP addresses a different form of stability: residualizing outcomes and interaction features against a fixed baseline space makes dose adaptation invariant to shifts within that space. This finite-sample algebraic property does not imply Neyman orthogonality or double robustness.

The separation of baseline and residual structure also connects BCA-GP to semiparametric GP modeling. Explicit mean functions yield generalized least-squares prediction (Rasmussen & Williams, 2006, Section 2.7), while orthogonal GP models constrain the stochastic component to be orthogonal to the mean space to address identifiability (Plumlee & Joseph, 2018). Related causal models separate prognostic and treatment-related components (Hahn et al., 2020), use multitask GPs for individualized effects (Alaa & van der Schaar, 2017), or incorporate propensity uncertainty into observational dose–response inference (Dailey, 2024). BCA-GP does not impose an orthogonal residual GP prior. Instead, it preserves a shared, unpenalized baseline through direction learning, candidate scoring, and prediction, extending baseline reproduction to the complete adaptive fitting procedure.

Our prediction and adaptation steps connect to two further kernel perspectives. Functional balancing controls imbalance over a reproducing-kernel Hilbert space (Wong & Chan, 2018), and generalized optimal matching relates weights to approximation error and variance (Kallus, 2020). BCA-GP's constraint $F^\top w_z = f(x)$ instead reproduces baseline functions at each prediction location; it is not a global treatment-distribution balancing condition. Separately, centered kernel alignment (Cortes et al., 2012) and learned representations (Wilson et al., 2016) adapt similarity using supervised information. BCA-GP learns dose directions from baseline-adjusted regression coefficients rather than optimizing alignment or a deep feature map. The common baseline ensures that these outcome-dependent choices remain unchanged under the specified baseline shifts.

3 CAUSAL TARGET AND NOTATION

Let $Z_i = (X_i, A_i, Y_i)$ be i.i.d. observations, with pre-treatment covariates $X_i \in \mathbb{R}^d$ and dose $A_i \in [0, 1]$. Let $Y(a)$ denote the potential outcome under dose a. We target the conditional dose–response function (CDRF) and its population average:

$$\mu_0(x, a) = \mathbb{E}[Y(a) \mid X = x], \qquad \bar{\mu}_0(a) = \mathbb{E}_X[\mu_0(X, a)], \tag{1}$$

where $\bar{\mu}_0$ is the average dose–response function (ADRF).

Let $p_0(a \mid x)$ denote the observational treatment density and $q(a \mid x)$ a prespecified dose density used for evaluation. We use the standard causal identification conditions, with positivity restricted to the evaluation support (Hernán & Robins, 2020; Kennedy et al., 2017).

Assumption 1 (Causal identification)**.** *Consistency $Y = Y(A)$ and conditional exchangeability $Y(a) \perp A \mid X$ hold on the target support. The conditional treatment density $p_0(a \mid x)$ is positive wherever $q(a \mid x)$ is positive, for $P_X(dx)\,da$-almost every (x, a). Relevant conditional means are finite.*

Under Assumption 1, $\mu_0(x, a) = \mathbb{E}[Y \mid X = x, A = a]$ almost everywhere on the target support. Pointwise interpretation requires a specified continuous version of the conditional mean. We assess the target-distribution risk

$$\mathcal{E}_Q(\widehat{\mu}) = \int (\widehat{\mu}(x, a) - \mu_0(x, a))^2 \, dQ(x, a), \quad Q(dx, da) = P_X(dx) q(a \mid x)\, da. \tag{2}$$

Here P_X determines which covariate profiles are evaluated, while $q(a \mid x)$ determines how their doses are weighted. The risk therefore measures response-surface error over the evaluation support, rather than only at observed doses. Uniform-dose evaluation takes $q(a \mid x) = 1$ on $[0, 1]$ and assesses complete conditional response curves. Identification alone does not ensure accurate estimation in regions with little observational support.

Write $z = (x, a)$ for a prediction input, $y = (Y_1, \ldots, Y_n)^\top$, and $F \in \mathbb{R}^{n \times p}$ for a full-column-rank baseline design with rows $f(X_i)^\top$. The map f depends on training covariates and fixed auxiliary randomness, but not on outcomes or doses; its span contains constants. The columns of F represent covariate-only outcome patterns that we want the estimator to preserve without shrinkage. The residual projector

$$M_F = I - F(F^\top F)^{-1} F^\top$$

removes these patterns from a training vector: $M_F F = 0$, hence $M_F(y + Fc) = M_F y$. Appendix A.2 derives this projector from least squares. This identity motivates the construction below. It removes a specified baseline space, not all effects of confounding. Structural results condition on the observed input design and auxiliary randomness.

4 BASELINE-CALIBRATED ADAPTIVE GP

We model the response as $\mu(x, a) = f(x)^\top \beta + g(x, a)$. For each candidate kernel, $g \sim \mathcal{GP}(0, k_{\ell,C})$: the baseline specifies the mean, while the kernel describes the remaining covariate structure and heterogeneous dose responses. This explicit-mean formulation follows standard GP regression (Rasmussen & Williams, 2006). Our construction retains the same baseline while learning dose directions, selecting a kernel, and predicting the complete response.

Figure 1 shows the information flow, and Algorithm 1 summarizes training and prediction.

4.1 SEPARATING THE BASELINE FROM DOSE ADAPTATION

A covariate reference kernel is defined on standardized inputs. We select m_0 training covariate landmarks and construct Nyström features (Williams & Seeger, 2001), then append a constant. An SVD removes numerically dependent columns and yields a baseline matrix with $F^\top F = I$ and $p \leq m_0 + 1$. The resulting feature transformation is reused at prediction inputs.

To keep baseline amplitude from driving adaptation, compute $e = M_F y$ and use its root-mean-square magnitude, with a positive floor $\varepsilon_{\text{scale}}$, as a scale s. Normalize the outcomes as $\widetilde{y} = y/s$. Both e and

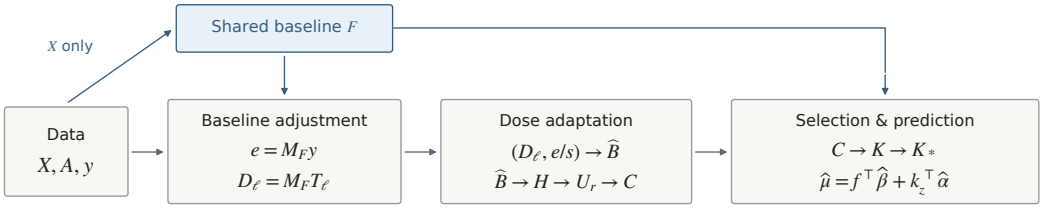

Figure 1: BCA-GP computation. Baseline-adjusted outcomes and interactions yield pilot coefficients $\widehat{B}$, whose second moment H determines dose directions and candidate kernels. Blue paths carry the same baseline F into adjustment, selection, and final prediction.

s are unchanged by adding Fc to y, so normalization preserves the baseline-shift property. The floor prevents division by zero.

To represent variation across doses, we use a spline feature map $b(a) \in \mathbb{R}^{d_a}$ whose linear combinations describe smooth dose-response patterns. Centering removes the constant component of each feature, while whitening puts the resulting directions on a common scale. For covariate lengthscale multiplier ℓ and positive semidefinite dose matrix C, define

$$k_{\ell,C}(z, z') = k_{X,\ell}(x, x')\{1 + b(a)^\top C b(a')\}. \tag{3}$$

The covariate factor shares information between similar profiles. Within the braces, the constant term permits remaining covariate-only variation, while $b(a)^\top C b(a')$ determines which dose patterns can be shared. Learning C therefore adapts the dose structure of the residual GP.

4.2 Residual-based dose adaptation

We construct candidates for C by fitting a preliminary response surface and extracting dose shapes with large fitted coefficients across covariate profiles. This preliminary fit supplies directions; the final GP estimates their contributions to the complete response.

A preliminary model of heterogeneous dose responses. For each ℓ, construct a Nyström map $z_\ell(x) \in \mathbb{R}^m$ from m_1 training covariate landmarks, with $m \le m_1$. Unlike the fixed baseline features, these features describe how dose coefficients vary with x and may change with ℓ. Separate landmark sets allow baseline dimension and direction-learning capacity to be chosen independently. Consider

$$h_B(x, a) = \theta_B(x)^\top b(a), \qquad \theta_B(x) = B^\top z_\ell(x), \qquad B \in \mathbb{R}^{m \times d_a}.$$

The vector $\theta_B(x)$ assigns coefficients to the shared spline functions for profile x. Thus one matrix B represents a family of individual dose curves. We estimate it from outcomes at the observed doses, adjusting for the common baseline.

Fitting the preliminary surface after baseline adjustment. Form the interaction design $T_\ell \in \mathbb{R}^{n \times (m d_a)}$ by multiplying each covariate feature by each dose feature at the observed inputs. If $w = \mathrm{vec}_{\mathrm{row}}(B)$ stacks the rows of B, then $(T_\ell w)_i = h_B(X_i, A_i)$. Thus T_ℓ expresses the preliminary surface as a linear regression in the entries of B.

Including the unpenalized baseline F and eliminating its coefficients by least squares gives the adjusted response $M_F \widetilde{y} = e/s$ and design $D_\ell = M_F T_\ell$. Each column of D_ℓ is the residual from regressing an interaction feature on F. We then solve

$$\widehat{B}_{\ell,\lambda} = \arg\min_B \left\{ \frac{1}{n} \left\| e/s - D_\ell \, \mathrm{vec}_{\mathrm{row}}(B) \right\|_2^2 + \lambda \|B\|_F^2 \right\}. \tag{4}$$

Adjusting both sides prevents the regression from explaining a baseline-space pattern through the interactions. The penalty $\lambda > 0$ stabilizes the coefficients. Appendix A.3 gives the derivation and closed-form solution.

Turning fitted curves into dose-kernel directions. For profile X_i, the fitted dose coefficients are $\widehat{\theta}_i = \widehat{B}_{\ell,\lambda}^{\top} z_\ell(X_i)$. Their second-moment matrix is

$$H_{\ell,\lambda} = \frac{1}{n}\sum_{i=1}^{n} \widehat{\theta}_i \widehat{\theta}_i^{\top}.$$

A unit vector u defines a dose shape $\phi_u(a) = u^{\top} b(a)$. The quantity $u^{\top} H_{\ell,\lambda} u = n^{-1}\sum_i (u^{\top}\widehat{\theta}_i)^2$ measures its average squared coefficient in the preliminary curves. The leading eigenvectors therefore identify orthogonal coefficient directions with the largest fitted energy. Using an uncentered second moment retains both shared dose patterns and variation across individuals.

Let $U_r = [u_1, \ldots, u_r]$ contain the leading r eigenvectors. For $r < d_a$, set $C = U_r U_r^{\top}/r$. The resulting dose kernel averages products of the learned shapes $\phi_j(a) = u_j^{\top} b(a)$ at two doses. The covariate kernel allows their contributions to vary across profiles. Only the directions are transferred to the GP, not the preliminary individual coefficients.

Dividing by r fixes $\mathrm{tr}(C) = 1$, keeping total dose-kernel weight comparable across ranks. A full-rank candidate $C = I_{d_a}/d_a$ retains all dose directions. Both constructions are positive semidefinite. We compare lengthscale multipliers $\ell \in \mathcal{L}$, direction penalties $\lambda \in \Lambda$, ranks $r \in \mathcal{R}$, and noise parameters $\eta \in \mathcal{H}$. These finite sets are fixed before fitting. The full-rank candidate is evaluated once per lengthscale and noise value because it does not depend on λ.

4.3 Common-baseline selection and prediction

We now choose among the candidate kernels and fit the complete response. Selection must retain the baseline adjustment: otherwise a baseline shift could leave the proposed dose directions unchanged but alter which kernel is selected.

Comparing candidates under a common baseline. For each candidate, let $K_{ij} = k_{\ell,C}(z_i, z_j)$ and $V = K + \eta I$, where $\eta > 0$. Estimate $\widetilde{\beta}_V$ by generalized least squares, minimizing the squared residual weighted by V^{-1}, and write $r_V = \widetilde{y} - F\widetilde{\beta}_V$. Define

$$G = F^{\top} V^{-1} F, \qquad \Pi = V^{-1} - V^{-1} F G^{-1} F^{\top} V^{-1}. \tag{5}$$

Then $\widetilde{\beta}_V = G^{-1} F^{\top} V^{-1} \widetilde{y}$ and $r_V^{\top} V^{-1} r_V = \widetilde{y}^{\top} \Pi \widetilde{y}$. Thus Π combines baseline fitting and weighted residual evaluation, and satisfies $\Pi F = 0$.

Scoring baseline-adjusted fit. We select the candidate minimizing

$$S(K, \eta) = \log\det V + \log\det G + (n-p)\log\left(\frac{\widetilde{y}^{\top}\Pi\widetilde{y}}{n-p}\right). \tag{6}$$

The final term measures the remaining weighted residual with $n - p$ residual degrees of freedom. The determinant terms account jointly for covariance volume and baseline estimation, so selection does not rely on residual fit alone.

For a fixed covariance family, this form follows from restricted likelihood based on error contrasts (Patterson & Thompson, 1971), after profiling out an overall variance scale. Appendix A.4 gives the derivation. Because our dose directions are learned from the same outcomes, we use this expression as a selection criterion without asserting fixed-kernel likelihood inference for the adaptive fit.

All candidates are compared on the full training data. Under $y \mapsto y + Fc$, the candidate kernels and scale s are unchanged, and $\Pi F = 0$ preserves every weighted residual term. The scores, and hence selection under fixed tie-breaking, are therefore unchanged.

Predicting the complete response. Using the selected covariance, we re-estimate the baseline and predict the residual function in original outcome units:

$$\begin{aligned} \widehat{\beta} &= G^{-1} F^{\top} V^{-1} y, \qquad \widehat{\alpha} = V^{-1}(y - F\widehat{\beta}) = \Pi y, \\ \widehat{\mu}(x, a) &= f(x)^{\top}\widehat{\beta} + k_z^{\top}\widehat{\alpha}. \end{aligned} \tag{7}$$

Algorithm 1 BCA-GP fitting and prediction

Require: Training data $\{(X_i, A_i, Y_i)\}_{i=1}^n$; candidate sets $\mathcal{L}, \Lambda, \mathcal{R}, \mathcal{H}$.
Construct the baseline map f, design F, and dose map b.
Compute $e = M_F y$, residual scale s, and $\widetilde{y} = y/s$.
Initialize an empty candidate collection $\mathcal{C}$.
for each $\ell \in \mathcal{L}$ **do**
  Construct z_ℓ, interaction design T_ℓ, and $D_\ell = M_F T_\ell$.
  for each $\lambda \in \Lambda$ **do**
    Fit $\widehat{B}_{\ell,\lambda}$ using Eq. (4) and form $H_{\ell,\lambda}$.
    for each $r \in \mathcal{R}$ with $r < d_a$ **do**
      Take the leading r eigenvectors U_r and add $(\ell, U_r U_r^\top / r)$ to $\mathcal{C}$.
    end for
  end for
  if $d_a \in \mathcal{R}$ **then**
    Add $(\ell, I_{d_a}/d_a)$ to $\mathcal{C}$.
  end if
end for
for each $(\ell, C) \in \mathcal{C}$ and $\eta \in \mathcal{H}$ **do**
  Form K, V, G, and Π; compute $S(K, \eta)$ using Eq. (6).
end for
Select the smallest score using deterministic tie-breaking.
Compute $\widehat{\beta}$ and $\widehat{\alpha}$ under the selected covariance using Eq. (7).
Ensure: Predictor $\widehat{\mu}(x, a) = f(x)^\top \widehat{\beta} + k_z^\top \widehat{\alpha}$.

Here $k_z = (k(z_1, z), \ldots, k(z_n, z))^\top$ uses the selected kernel. The baseline coefficients are unpenalized; the kernel and noise parameter regularize the remaining function. In particular, $\widehat{\beta}$ is re-estimated under the selected V, rather than carried over from the initial Euclidean projection. The implementation fits on y/s and rescales predictions by s, yielding the same expression. Varying a at fixed x gives the estimated CDRF.

5 Baseline invariance and CDRF error

We analyze how the shared baseline affects prediction, adaptation, and the remaining estimation error. Throughout, assume $n > p$, full column rank of F, and $\eta > 0$, so that V and G are positive definite. Structural statements condition on the observed inputs and auxiliary randomness. Results for the complete fitting procedure use exact arithmetic and deterministic tie-breaking. Proofs are collected in Appendix A.

We begin with prediction under a realized fitted kernel. Although the residual GP is regularized, the unpenalized baseline imposes an exact reproduction property on the prediction weights.

Proposition 1 (Response-level reproduction)**.** *For each realized fitted kernel, Eq.* (7) *can be written as* $\widehat{\mu}(z) = w_z^\top y$*, where* $F^\top w_z = f(x)$*. Consequently,* $w_z^\top F c = f(x)^\top c$ *for every* $c \in \mathbb{R}^p$*. Since the baseline space contains constants,* $\mathbf{1}^\top w_z = 1$*.*

Thus any function represented by the baseline features is reproduced at the prediction input. The weights need not be positive and depend on outcomes through kernel selection, so this representation does not make the complete learning algorithm linear in y. In particular, reproduction for a fixed kernel alone would not ensure stability if a baseline shift changed the learned dose directions or selected hyperparameters. The next result addresses this dependence by refitting the entire model after the shift.

Theorem 1 (Equivariance of the complete procedure)**.** *Hold the inputs and auxiliary randomness fixed. For any* $c \in \mathbb{R}^p$*, let* $y' = y + Fc$ *and rerun Algorithm 1. The residual scale, candidate dose matrices, and selected hyperparameters are unchanged, and*

$$\widehat{\mu}_{y'}(x, a) = \widehat{\mu}_y(x, a) + f(x)^\top c. \tag{8}$$

If the true surface is also shifted to $\mu_0'(x,a) = \mu_0(x,a) + f(x)^\top c$ with the same noise realization, then $\widehat{\mu}_{y'} - \mu_0' = \widehat{\mu}_y - \mu_0$ pointwise. The integrated CDRF risks are therefore equal under every common evaluation measure Q for which the integrals exist.

The two residualization identities connect this result to the training procedure: $M_F F = 0$ leaves the pilot regression and its normalization unchanged, while $\Pi F = 0$ preserves every candidate score. The selected residual predictor is consequently unchanged, and the fitted baseline absorbs the added function. Because that function depends only on x, the estimated curve shifts in level without changing shape. This property allows dependence between covariates and dose, but applies only to shifts within the chosen baseline space; it does not assert robustness to changes in treatment assignment.

To understand the error that remains, we examine how the prediction weights approximate functions outside the explicit baseline. This perspective is related to kernel matching, which connects weighting with function approximation and variance control (Kallus, 2020). For a kernel k, define the worst-case approximation error over its unit RKHS ball as $D_k(w,z) = \sup_{\|u\|_{\mathcal{H}_k} \le 1} |w^\top u_n - u(z)|$, where $u_n = (u(z_1), \ldots, u(z_n))^\top$. For each realized fit, the BCA-GP weights minimize $D_k(w,z)^2 + \eta\|w\|_2^2$ subject to $F^\top w = f(x)$, as shown in Appendix A.7. The constraint reproduces baseline functions at a particular prediction input; it is not a global treatment-distribution balancing condition.

A selected low-rank kernel may nevertheless omit genuine dose patterns. We therefore assess approximation in a common reference space containing all candidate kernels, rather than assume that the selected space contains the truth. Define

$$\bar{k}(z,z') = \sum_{\ell \in \mathcal{L}} k_{X,\ell}(x,x')\{1 + b(a)^\top b(a')\}. \tag{9}$$

Conditional on the inputs and auxiliary randomness, this reference kernel is independent of outcomes. Every candidate satisfies $k_{\ell,C} \preceq \bar{k}$ in the positive-semidefinite kernel order. The reference retains all d_a spline directions, although its finite dose basis can still leave some response structure unrepresented.

Theorem 2 (Deterministic CDRF error bound)**.** *Suppose that, for the realized design, $\mu_0(z) = f(x)^\top \beta_0 + u_0(z) + r(z)$, where $u_0 \in \mathcal{H}_{\bar{k}}$ and $\|u_0\|_{\mathcal{H}_{\bar{k}}} \le B$. Write $y_i = \mu_0(z_i) + \epsilon_i$ and $r_n = (r(z_i))_{i=1}^n$. The actual adaptive prediction weights satisfy*

$$|\widehat{\mu}(z) - \mu_0(z)| \le B D_{\bar{k}}(w_z, z) + |w_z^\top r_n - r(z)| + |w_z^\top \epsilon|. \tag{10}$$

The bound separates three sources of error: approximation within the reference space, propagation of the representation remainder, and noise propagated through the adaptive weights. The baseline contribution cancels exactly by Proposition 1, so its amplitude does not enter the bound. Squaring and integrating gives the CDRF risk bound in Appendix A.8.

This decomposition remains valid after outcome-dependent kernel selection because it retains the realized noise term $|w_z^\top \epsilon|$. Without further assumptions, that term cannot be replaced by a variance expression for weights fixed independently of the noise. Moreover, the training procedure minimizes the candidate score, not the error bound. The result therefore identifies the quantities that govern estimation error without claiming that selection automatically controls them.

Finally, evaluating complete response curves requires support across the target doses. If $q(a \mid x)/p_0(a \mid x) \le W$ almost everywhere on the target support, then $\mathcal{E}_Q(\widehat{\mu}) \le W \mathcal{E}_{P_{X,A}}(\widehat{\mu})$ for every realized predictor. This change-of-measure inequality relates target error to observational error, but does not establish a generalization guarantee for the training score. Together, the results establish exact finite-space stability and an explicit error decomposition; consistency for unrestricted CDRFs requires additional control of approximation, selection, and noise propagation.

6 Experiments

We evaluate response accuracy, examine the contributions of individual components, and test the baseline-shift stability established in Theorem 1.

We compare BCA-GP with GP (Rasmussen & Williams, 2006), DRNet (Schwab et al., 2020), VCNet and VCNet+TR (Nie et al., 2021), TransTEE (Zhang et al., 2023), and GIKS (Nagalapatti et al.,

Table 1: Mean $\pm$ sample SD of response-level MISE (lower is better).

CDRF MISE					
Model	IHDP	NEWS	TCGA-0	TCGA-1	TCGA-2
Base GP	0.9906 ± 0.0992	0.8359 ± 0.0361	0.0331 ± 0.0063	0.0263 ± 0.0034	0.0398 ± 0.0059
BCA-GP	$\mathbf{0.6088 \pm 0.0886}$	$\mathbf{0.6293 \pm 0.0360}$	$\mathbf{0.0084 \pm 0.0055}$	$\mathbf{0.0033 \pm 0.0004}$	$\mathbf{0.0149 \pm 0.0021}$
DRNet	1.7555 ± 0.3899	0.9889 ± 0.1026	0.0464 ± 0.0090	0.0294 ± 0.0053	0.2141 ± 0.1208
VCNet	1.7811 ± 0.3841	0.8685 ± 0.0387	0.0703 ± 0.0428	0.0108 ± 0.0030	0.0623 ± 0.0077
VCNet+TR	1.4348 ± 0.5095	0.8724 ± 0.0605	0.0633 ± 0.1133	0.0157 ± 0.0176	0.0427 ± 0.0143
TransTEE	1.5630 ± 0.2505	1.0735 ± 0.0541	0.0514 ± 0.0148	0.0360 ± 0.0065	0.0781 ± 0.0167
GIKS	1.3075 ± 0.5004	0.7224 ± 0.0410	0.0347 ± 0.0205	0.0093 ± 0.0027	0.0322 ± 0.0087
ADRF MISE					
Model	IHDP	NEWS	TCGA-0	TCGA-1	TCGA-2
Base GP	0.0839 ± 0.0276	0.0164 ± 0.0083	0.0105 ± 0.0060	0.0063 ± 0.0033	0.0088 ± 0.0056
BCA-GP	$\mathbf{0.0717 \pm 0.0226}$	$\mathbf{0.0085 \pm 0.0051}$	$\mathbf{0.0038 \pm 0.0048}$	$\mathbf{0.0003 \pm 0.0002}$	$\mathbf{0.0007 \pm 0.0006}$
DRNet	0.5149 ± 0.2610	0.0313 ± 0.0204	0.0191 ± 0.0097	0.0081 ± 0.0051	0.0647 ± 0.0224
VCNet	0.5221 ± 0.3320	0.0253 ± 0.0230	0.0601 ± 0.0413	0.0060 ± 0.0028	0.0367 ± 0.0090
VCNet+TR	0.2762 ± 0.1750	0.0299 ± 0.0276	0.0525 ± 0.1077	0.0067 ± 0.0096	0.0195 ± 0.0122
TransTEE	0.4984 ± 0.1055	0.0086 ± 0.0041	0.0152 ± 0.0102	0.0039 ± 0.0017	0.0332 ± 0.0140
GIKS	0.2303 ± 0.2197	0.0266 ± 0.0116	0.0275 ± 0.0209	0.0049 ± 0.0027	0.0212 ± 0.0086

Table 2: Component ablations: CDRF MISE, mean $\pm$ sample SD over identical original seeds. TCGA values are multiplied by 10^3. Full-rank fixes $C = I_{d_a}/d_a$; no adjustment omits baseline adjustment only in the pilot regression; intercept only removes covariate baseline features.

Dataset	BCA-GP	Full-rank	No adjustment	Intercept only
IHDP	0.6088 ± 0.0886	1.0497 ± 0.1292	1.7508 ± 0.5505	0.5902 ± 0.0874
NEWS	0.6293 ± 0.0360	0.8430 ± 0.0390	0.6677 ± 0.0450	0.6283 ± 0.0357
TCGA-0	8.438 ± 5.545	12.641 ± 8.981	12.641 ± 8.981	11.508 ± 6.135
TCGA-1	3.350 ± 0.403	3.326 ± 0.221	3.326 ± 0.221	4.276 ± 0.699
TCGA-2	14.912 ± 2.074	17.032 ± 0.889	17.032 ± 0.889	18.179 ± 4.851

2024) on IHDP, NEWS, and TCGA-0/1/2. Neural models use author-source implementations with documented compatibility and evaluation adapters. Data sources and training settings are documented in Appendix C.

For n_{te} test profiles and N_a equally spaced doses, let $e_{ij} = \widehat{\mu}(X_i, a_j) - \mu_0(X_i, a_j)$. We report

$$\mathrm{MISE}_{\mathrm{CDRF}} = \frac{1}{n_{\mathrm{te}} N_a} \sum_{i,j} e_{ij}^2, \qquad \mathrm{MISE}_{\mathrm{ADRF}} = \frac{1}{N_a} \sum_{j} \left(\frac{1}{n_{\mathrm{te}}} \sum_{i} e_{ij} \right)^2. \tag{11}$$

The first metric evaluates individual response curves, while the second evaluates their population average. Tables report descriptive means and sample SD; repeated splits are not independent population samples. The benchmark results are presented in Table 1.

The component ablations in Table 2 show that dose adaptation improves mean CDRF accuracy on most datasets, with little difference on TCGA-1. Removing pilot baseline adjustment is particularly detrimental on IHDP and selects the full-rank kernel on TCGA. An intercept-only baseline performs slightly better on IHDP and NEWS but worse on TCGA, indicating that a richer baseline does not uniformly improve accuracy. We therefore examine its stability role separately.

To test baseline equivariance, we add tFc to IHDP training outcomes and $tf(x)^\top c$ to test truth, choosing c from training covariates and refitting the complete model at each amplitude t. Figure 2 shows that BCA-GP preserves prediction error and hyperparameter selection up to numerical precision. The intercept-only and unadjusted-pilot variants are sensitive to these shifts, illustrating why retaining

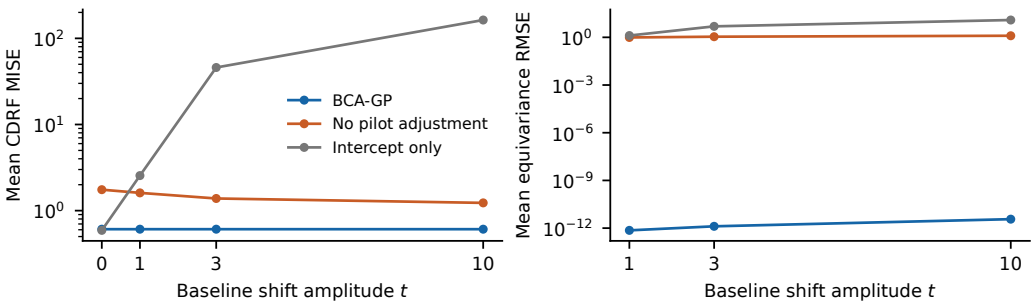

Figure 2: Controlled IHDP baseline shifts with full refitting. Left: mean CDRF error. Right: mean prediction-equivariance RMSE on a log scale, excluding the zero-shift case. Training outcomes and test truth receive the same baseline function.

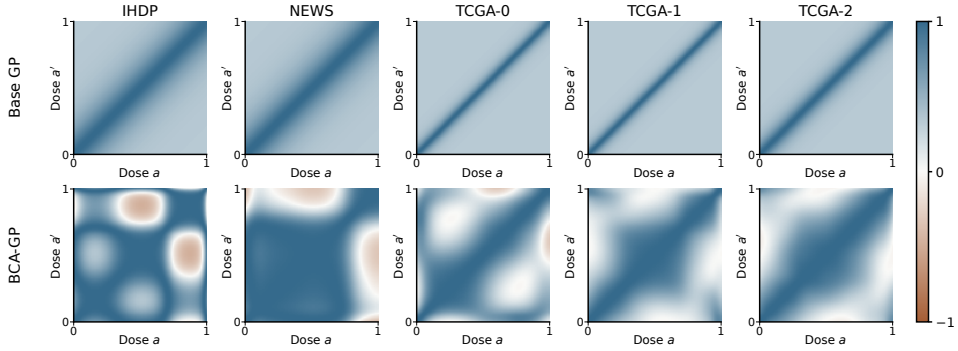

Figure 3: Fitted dose association at identical covariates: GP (top) and BCA-GP (bottom). Each kernel is normalized by its diagonal, then averaged across the original seeds. Both axes are dose; the shared color scale represents kernel correlation.

the same baseline throughout adaptation and prediction matters. This experiment tests shifts within the chosen baseline space.

Figure 3 compares the fitted dose associations at identical covariates. GP uses a stationary Matérn dose kernel, whose association depends on dose separation. BCA-GP's learned spline directions permit associations that also depend on dose location. The heatmaps illustrate these different information-sharing patterns; they do not represent causal correlations or independently establish prediction accuracy.

7 Conclusion

We presented BCA-GP, a Gaussian process framework for conditional dose-response estimation that combines an unpenalized covariate baseline with an adaptive dose kernel. Learning dose directions from baseline-adjusted outcomes allows the model to capture heterogeneous responses while preserving baseline structure throughout fitting. We established exact equivariance to baseline-space shifts and derived an error bound that accounts for adaptive kernel selection. Experiments across five benchmarks demonstrate competitive conditional and average response accuracy, while controlled baseline shifts confirm the stability of the complete procedure.

AI use statement

In this work, we used generative AI tools for clean and reformat dataset, provide feedback on experiments and manuscript polishing. We have not used generative AI tools for help develop theoretical

models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, assist in the writing of proofs, propose or refine hypotheses; and generate synthetic data sets and assist with translation are not applicable to this work.

Reproducibility statement

Algorithm 1 summarizes the BCA-GP training and prediction procedure. The supplementary material provides all proofs, implementation details, hyperparameter settings, and experimental protocols. It also includes a link to the BCA-GP source code and the implementation of the IHDP experiments. We will release the complete codebase for all experiments as open source.

References

Ahmed M. Alaa and Mihaela van der Schaar. Bayesian inference of individualized treatment effects using multi-task gaussian processes. In *Advances in Neural Information Processing Systems*, 2017.

Taha Bahadori, Eric Tchetgen Tchetgen, and David Heckerman. End-to-end balancing for causal continuous treatment-effect estimation. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 1313–1326. PMLR, 2022. URL https://proceedings.mlr.press/v162/bahadori22a.html.

Ioana Bica, James Jordon, and Mihaela van der Schaar. Estimating the effects of continuous-valued interventions using generative adversarial networks. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/bea5955b308361a1b07bc55042e25e54-Abstract.html.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018. doi: 10.1111/ectj.12097.

Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Algorithms for learning kernels based on centered alignment. *Journal of Machine Learning Research*, 13(28):795–828, 2012. URL https://jmlr.org/papers/v13/cortes12a.html.

Jake R. Dailey. Gaussian processes for observational dose-response inference. *arXiv preprint arXiv:2409.17043*, 2024. URL https://arxiv.org/abs/2409.17043.

P. Richard Hahn, Jared S. Murray, and Carlos M. Carvalho. Bayesian regression tree models for causal inference: Regularization, confounding, and heterogeneous effects (with discussion). *Bayesian Analysis*, 15(3):965–1056, 2020. doi: 10.1214/19-BA1195.

Miguel A. Hernán and James M. Robins. *Causal Inference: What If*. Chapman & Hall/CRC, Boca Raton, 2020. URL https://miguelhernan.org/whatifbook.

Keisuke Hirano and Guido W. Imbens. The propensity score with continuous treatments. In *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, pp. 73–84. Wiley, 2004.

Kosuke Imai and David A. van Dyk. Causal inference with general treatment regimes: Generalizing the propensity score. *Journal of the American Statistical Association*, 99(467):854–866, 2004. doi: 10.1198/016214504000001187.

Nathan Kallus. Generalized optimal matching methods for causal inference. *Journal of Machine Learning Research*, 21(62):1–54, 2020. URL https://jmlr.org/papers/v21/19-120.html.

Edward H. Kennedy, Zongming Ma, Matthew D. McHugh, and Dylan S. Small. Non-parametric methods for doubly robust estimation of continuous treatment effects. *Journal of the Royal Statistical Society: Series B*, 79(4):1229–1245, 2017.

Seok-Jin Kim and Kaizheng Wang. Estimating continuous treatment effects with two-stage kernel ridge regression. *arXiv preprint arXiv:2604.13410*, 2026. URL https://arxiv.org/abs/2604.13410.

Lokesh Nagalapatti, Akshay Iyer, Abir De, and Sunita Sarawagi. Continuous treatment effect estimation using gradient interpolation and kernel smoothing. In *AAAI Conference on Artificial Intelligence*, 2024.

Lizhen Nie, Mao Ye, Qiang Liu, and Dan Nicolae. VCNet and functional targeted regularization for learning causal effects of continuous treatments. In *International Conference on Learning Representations*, 2021.

H. D. Patterson and R. Thompson. Recovery of inter-block information when block sizes are unequal. *Biometrika*, 58(3):545–554, 1971. doi: 10.1093/biomet/58.3.545.

Matthew Plumlee and V. Roshan Joseph. Orthogonal gaussian process models. *Statistica Sinica*, 28 (2):601–619, 2018. doi: 10.5705/ss.202015.0404.

Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.

Patrick Schwab, Lorenz Linhardt, Stefan Bauer, Joachim M. Buhmann, and Walter Karlen. Learning counterfactual representations for estimating individual dose–response curves. In *AAAI Conference on Artificial Intelligence*, 2020.

Rahul Singh, Liyuan Xu, and Arthur Gretton. Kernel methods for causal functions: Dose, heterogeneous and incremental response curves. *Biometrika*, 111(2):497–516, 2024. doi: 10.1093/biomet/asad042.

Christopher K. I. Williams and Matthias Seeger. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems*, volume 13, pp. 682–688. MIT Press, 2001. URL https://proceedings.neurips.cc/paper/2000/file/19de10adbaa1b2ee13f77f679fa1483a-Paper.pdf.

Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P. Xing. Deep kernel learning. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pp. 370–378, 2016. URL https://proceedings.mlr.press/v51/wilson16.html.

Raymond K. W. Wong and Kwun Chuen Gary Chan. Kernel-based covariate functional balancing for observational studies. *Biometrika*, 105(1):199–213, 2018. URL https://raymondkww.github.io/publications/wong-chan18/.

Yi-Fan Zhang, Hanlin Zhang, Zachary C. Lipton, Li Erran Li, and Eric P. Xing. Exploring transformer backbones for heterogeneous treatment effect estimation. *Transactions on Machine Learning Research*, 2023. URL https://openreview.net/forum?id=1kl4YM2Q7P.

A PROOFS AND STATISTICAL SCOPE

A.1 IDENTIFICATION AND POSITIVE SEMIDEFINITENESS

Under consistency and conditional exchangeability, $\mathbb{E}[Y \mid X = x, A = a] = \mathbb{E}[Y(a) \mid X = x, A = a] = \mathbb{E}[Y(a) \mid X = x]$ on the supported region, with almost-everywhere interpretation for continuous doses. This identifies the CDRF in Eq. (1). The ADRF follows by integration over P_X.

For $C \succeq 0$, the factor $1 + b(a)^\top C b(a')$ is a positive semidefinite kernel with feature map $(1, C^{1/2} b(a))$. Its product with the positive semidefinite covariate kernel is positive semidefinite. Adding ηI makes the training covariance strictly positive definite. If F has full column rank, then $v^\top G v = (Fv)^\top V^{-1}(Fv) > 0$ for every nonzero v.

A.2 DERIVING THE BASELINE RESIDUAL PROJECTOR

The matrix M_F expresses ordinary least-squares residualization against the baseline features. Let $F \in \mathbb{R}^{n \times p}$ have full column rank. For any training vector $v \in \mathbb{R}^n$, consider

$$\widehat{\gamma}(v) = \arg\min_{\gamma} \|v - F\gamma\|_2^2, \qquad \gamma \in \mathbb{R}^p.$$

Differentiating the objective gives the normal equations

$$-2F^\top(v - F\widehat{\gamma}) = 0, \qquad F^\top F \widehat{\gamma} = F^\top v.$$

Since $F^\top F$ is positive definite, the solution is unique:

$$\widehat{\gamma}(v) = (F^\top F)^{-1} F^\top v.$$

The fitted baseline and its residual are therefore

$$\widehat{v}_F = F\widehat{\gamma}(v) = \underbrace{F(F^\top F)^{-1}F^\top}_{P_F} v,$$
$$v - \widehat{v}_F = \underbrace{\left[I_n - F(F^\top F)^{-1}F^\top\right]}_{M_F} v.$$

Thus P_F projects onto $\mathrm{col}(F)$ and $M_F = I_n - P_F$ retains the orthogonal residual. Indeed, direct multiplication gives

$$P_F^\top = P_F, \quad P_F^2 = P_F, \qquad M_F^\top = M_F, \quad M_F^2 = M_F, \qquad F^\top M_F v = 0.$$

The last identity means that the residual is orthogonal to every baseline feature on the training design. It does not imply statistical independence from the covariates.

Why baseline shifts disappear. Because $P_F F = F$, we have

$$M_F F = 0, \qquad M_F(v + Fc) = M_F v \quad \text{for every } c \in \mathbb{R}^p.$$

Adding a function in the chosen baseline space changes its fitted component but leaves the residual unchanged. In BCA-GP, setting $v = y$ gives $e = M_F y$; applying the same operation to each column of T_ℓ gives $D_\ell = M_F T_\ell$. This explains why both the response and the direction-learning design are residualized against the shared baseline.

Special cases and scope. If the baseline contains only a constant, $F = \mathbf{1}$, then $P_F v = \bar{v}\mathbf{1}$ and $M_F v = v - \bar{v}\mathbf{1}$, where $\bar{v} = n^{-1}\mathbf{1}^\top v$. Residualization is then ordinary mean centering. For the orthonormal features used in our implementation, $F^\top F = I_p$, so the general expression simplifies to $M_F = I_n - FF^\top$. These identities remove only components in the specified finite baseline space; they do not establish removal of all confounding. The coefficients $\widehat{\gamma}$ above define the preliminary Euclidean projection, whereas the final predictor estimates its baseline by covariance-weighted generalized least squares in Eq. (7).

A.3 RESIDUALIZED REGRESSION FOR DOSE DIRECTIONS

The interaction design has entries

$$[T_\ell]_{i,(j-1)d_a+k} = z_{\ell,j}(X_i)b_k(A_i).$$

Fix a covariate map z_ℓ and a penalty $\lambda > 0$. Write $w = \mathrm{vec}_{\mathrm{row}}(B)$, so that $w_{(j-1)d_a+k} = B_{jk}$ and $\|w\|_2^2 = \|B\|_F^2$. The interaction design in Section 4.2 satisfies

$$(T_\ell w)_i = \sum_{j=1}^{m}\sum_{k=1}^{d_a} z_{\ell,j}(X_i)B_{jk}b_k(A_i) = h_B(X_i, A_i).$$

Thus T_ℓ is computed directly from the covariate and dose features; B contains the coefficients to be estimated.

To adjust for the shared baseline, begin with the joint objective

$$\min_{\gamma,w}\left\{\frac{1}{n}\|\widetilde{y} - F\gamma - T_\ell w\|_2^2 + \lambda\|w\|_2^2\right\}.$$

The baseline coefficients γ are unpenalized. For fixed w, their least-squares solution is

$$\widehat{\gamma}(w) = (F^\top F)^{-1}F^\top(\widetilde{y} - T_\ell w).$$

Substituting this solution leaves the residual

$$M_F(\widetilde{y} - T_\ell w) = e/s - D_\ell w, \qquad D_\ell = M_F T_\ell.$$

This gives Eq. (4) and explains why the interaction design is residualized along with the response. Each column of D_ℓ is orthogonal to $\mathrm{col}(F)$; the regression therefore adjusts for the baseline rather than assigning its fitted component to the interaction coefficients.

Differentiating the reduced objective gives

$$(D_\ell^\top D_\ell/n + \lambda I)\widehat{w} = D_\ell^\top(e/s)/n.$$

Since $\lambda > 0$, the coefficient matrix is positive definite and the solution is unique:

$$\widehat{w} = (D_\ell^\top D_\ell/n + \lambda I)^{-1}D_\ell^\top(e/s)/n, \qquad [\widehat{B}_{\ell,\lambda}]_{jk} = \widehat{w}_{(j-1)d_a+k}.$$

Rearranging $\widehat{w}$ into $\widehat{B}_{\ell,\lambda}$ only restores the matrix indexing; it is not an additional fitting step. This preliminary matrix supplies candidate dose directions, while the final response prediction uses Eq. (7).

Equivalent forms of the dose structure. Collect the covariate feature rows in Z_ℓ. The second moment can be computed as

$$H_{\ell,\lambda} = \widehat{B}_{\ell,\lambda}^\top(Z_\ell^\top Z_\ell/n)\widehat{B}_{\ell,\lambda}.$$

For $C = U_r U_r^\top/r$, its contribution to the kernel expands as

$$b(a)^\top C b(a') = \frac{1}{r}\sum_{j=1}^{r}\phi_j(a)\phi_j(a'), \qquad \phi_j(a) = u_j^\top b(a).$$

This shows explicitly how the learned spline shapes determine dose similarity while the final GP refits their contributions.

A.4 GENERALIZED LEAST SQUARES, GP COEFFICIENTS, AND THE SELECTION SCORE

Throughout this subsection, fix a candidate kernel and noise parameter, write $V = K + \eta I$, and assume that $F \in \mathbb{R}^{n\times p}$ has full column rank. Thus $V \succ 0$. We first work in original outcome units; the selection score below uses $\widetilde{y} = y/s$.

The weighted Gram matrix G. A Gram matrix collects pairwise inner products of a set of vectors. Writing $F = [v_1, \dots, v_p]$, where $v_j \in \mathbb{R}^n$ contains the values of baseline feature j across training observations, the ordinary Gram matrix has entries $(F^\top F)_{jk} = v_j^\top v_k$. The positive-definite matrix V^{-1} defines the inner product

$$\langle u, v \rangle_{V^{-1}} = u^\top V^{-1} v.$$

The corresponding weighted Gram matrix is therefore

$$G = F^\top V^{-1} F, \qquad G_{jk} = \langle v_j, v_k \rangle_{V^{-1}}.$$

Equivalently, with $\overline{F} = V^{-1/2} F$, we have $G = \overline{F}^\top \overline{F}$. Thus G compares baseline feature columns after whitening the observation covariance; it is a $p \times p$ feature Gram matrix, whereas K is an $n \times n$ kernel matrix over observations. A Gram matrix is not generally a correlation matrix. Since F has full column rank, $c^\top G c = \|V^{-1/2} F c\|_2^2 > 0$ for nonzero c, so G is invertible.

Estimating the baseline. Under the working model $y \mid \beta \sim \mathcal{N}(F\beta, \tau^2 V)$, the generalized least-squares estimate minimizes

$$L(\beta) = (y - F\beta)^\top V^{-1} (y - F\beta) = y^\top V^{-1} y - 2\beta^\top F^\top V^{-1} y + \beta^\top G \beta.$$

Setting the derivative to zero gives

$$-2F^\top V^{-1} y + 2G\widehat{\beta} = 0, \qquad \widehat{\beta} = G^{-1} F^\top V^{-1} y.$$

The scalar variance factor $\tau^2 > 0$ does not change this minimizer. Unlike the preliminary Euclidean baseline fit used to form $e = M_F y$, this estimate depends on the selected covariance V.

Eliminating the baseline and obtaining Π. Substitution leaves the ordinary residual vector

$$r = y - F\widehat{\beta} = \left[I - F G^{-1} F^\top V^{-1} \right] y.$$

Applying the precision matrix V^{-1} to this residual yields

$$\begin{aligned} V^{-1} r &= V^{-1} y - V^{-1} F G^{-1} F^\top V^{-1} y \\ &= \underbrace{\left[V^{-1} - V^{-1} F G^{-1} F^\top V^{-1} \right]}_{\Pi} y. \end{aligned}$$

In particular, r and Πy are different quantities: $\Pi y = V^{-1} r$ is the precision-weighted residual. The minimized objective satisfies

$$\begin{aligned} L(\widehat{\beta}) &= y^\top V^{-1} y - (F^\top V^{-1} y)^\top G^{-1} (F^\top V^{-1} y) \\ &= y^\top \Pi y. \end{aligned}$$

The matrix Π is symmetric, positive semidefinite, and annihilates F. To see its geometry, let

$$P_{\overline{F}} = \overline{F} (\overline{F}^\top \overline{F})^{-1} \overline{F}^\top.$$

Then

$$\Pi = V^{-1/2} (I - P_{\overline{F}}) V^{-1/2}, \qquad \Pi F = V^{-1} F - V^{-1} F G^{-1} G = 0.$$

Thus whitening, removing the whitened baseline, and mapping back produces Π. It is not generally an idempotent Euclidean projector. In the special case $V = I$, $G = F^\top F$ and $\Pi = M_F$.

Why the GP coefficient vector is $\widehat{\alpha} = V^{-1} r$. For fixed β and covariance parameters, let $g_n = (g(z_1), \dots, g(z_n))^\top$ have covariance $\tau^2 K$, with independent observation noise of covariance $\tau^2 \eta I$. For a prediction location z, write $k_z = (k(z_1, z), \dots, k(z_n, z))^\top$. The joint working Gaussian model is

$$\begin{bmatrix} y - F\beta \\ g(z) \end{bmatrix} \sim \mathcal{N}\left(0, \tau^2 \begin{bmatrix} V & k_z \\ k_z^\top & k(z,z) \end{bmatrix} \right).$$

Gaussian conditioning gives

$$\mathbb{E}[g(z) \mid y, \beta] = (\tau^2 k_z^\top)(\tau^2 V)^{-1} (y - F\beta) = k_z^\top V^{-1} (y - F\beta).$$

Plugging in the generalized least-squares estimate defines the predictor

$$\widehat{g}(z) = k_z^\top \widehat{\alpha}, \qquad \widehat{\alpha} = V^{-1}(y - F\widehat{\beta}) = V^{-1}r = \Pi y,$$

and hence $\widehat{\mu}(z) = f(x)^\top \widehat{\beta} + k_z^\top \widehat{\alpha}$. The vector $\widehat{\alpha}$ is shared across prediction locations; only k_z changes. In computation, $(K + \eta I)\widehat{\alpha} = r$ is solved by factorization rather than by explicitly forming V^{-1}. This derivation conditions on fixed working parameters and then plugs in their fitted values; it does not establish calibrated posterior uncertainty after outcome-dependent kernel selection. Fitting $\widetilde{y} = y/s$ gives coefficients $\widetilde{\beta} = \widehat{\beta}/s$ and $\widetilde{\alpha} = \widehat{\alpha}/s$, so multiplying predictions by s recovers the same original-scale estimator.

Deriving the baseline-adjusted selection score. For the normalized working model $\widetilde{y} \sim \mathcal{N}(F\beta, \tau^2 V)$, the preceding calculation gives

$$\widetilde{\beta}_V = G^{-1}F^\top V^{-1}\widetilde{y}, \qquad (\widetilde{y} - F\widetilde{\beta}_V)^\top V^{-1}(\widetilde{y} - F\widetilde{\beta}_V) = \widetilde{y}^\top \Pi \widetilde{y}.$$

To derive the score, let L have $n - p$ orthonormal columns spanning the orthogonal complement of $\mathrm{col}(F)$. Then $L^\top F = 0$, and the baseline-free contrasts have distribution

$$L^\top \widetilde{y} \sim \mathcal{N}(0, \tau^2 L^\top V L).$$

The identities

$$L(L^\top V L)^{-1} L^\top = \Pi, \qquad \det(L^\top V L) = \frac{\det V \det(F^\top V^{-1} F)}{\det(F^\top F)}$$

follow by expressing V in an orthonormal basis for $\mathrm{col}(F)$ and its complement and using a Schur complement. Consequently, up to a constant common to all candidates, twice the negative log density of the contrasts is

$$(n - p)\log \tau^2 + \log\det V + \log\det G + \frac{\widetilde{y}^\top \Pi \widetilde{y}}{\tau^2}.$$

The omitted term $-\log\det(F^\top F)$ is common because every candidate uses the same F. For a positive residual quadratic form, minimizing over τ^2 gives

$$\widehat{\tau}^2 = \frac{\widetilde{y}^\top \Pi \widetilde{y}}{n - p}.$$

Substitution gives Eq. (6), after removing the additional common constant $n - p$. The implementation floors a zero or very small quadratic form as specified in Appendix B; at zero, the unmodified variance optimum is a boundary limit.

This calculation motivates the score for a fixed covariance. In BCA-GP, dose directions and hence the candidate covariances depend on the outcomes. We evaluate the derived expression to compare those candidates, without treating it as a fixed-covariance likelihood for post-selection inference. Its baseline-shift invariance follows algebraically from $\Pi F = 0$, the common residual scale, and unchanged candidate construction.

A.5 Proof of Proposition 1

Substituting the fitted coefficients into Eq. (7) yields

$$\widehat{\mu}(z) = w_z^\top y, \qquad w_z = \Pi k_z + V^{-1}FG^{-1}f(x). \tag{12}$$

The matrix Π is symmetric and

$$\Pi F = V^{-1}F - V^{-1}FG^{-1}F^\top V^{-1}F = 0.$$

Equation (7) therefore gives Eq. (12), and

$$F^\top w_z = F^\top \Pi k_z + F^\top V^{-1}FG^{-1}f(x) = f(x).$$

If $f(x)^\top c_1 = 1$ for all relevant x, then $Fc_1 = \mathbf{1}$ and $\mathbf{1}^\top w_z = c_1^\top f(x) = 1$. Multiplying the reproduction identity by arbitrary c proves the final assertion. The assumption that the constant function is retained refers to the mathematical feature space; finite-precision rank truncation realizes this property only to numerical tolerance.

A.6 Proof of Theorem 1

The covariate kernel, baseline map, dose basis, and landmark choices are unchanged because they depend only on the fixed design and auxiliary randomness. Since $M_F F = 0$,

$$M_F(y + Fc) = M_F y,$$

so the residual scale s, tensor-regression response, and learned dose matrices are unchanged. For every resulting candidate, V, G, and Π are the same and $\Pi F = 0$. Hence

$$(y + Fc)^\top \Pi (y + Fc) = y^\top \Pi y.$$

After division by the unchanged s^2, every candidate score in Eq. (6) is identical. A fixed numerical floor on the quadratic form also preserves this identity. Deterministic tie-breaking yields the same selected candidate.

In original outcome units, the final coefficients satisfy

$$\widehat{\alpha}' = \Pi(y + Fc) = \widehat{\alpha}, \qquad \widehat{\beta}' = G^{-1} F^\top V^{-1}(y + Fc) = \widehat{\beta} + c.$$

Substitution proves Eq. (8). Subtracting the shifted true response yields pointwise equality of errors; integration gives equality of risks. This is a paired, fixed-design transformation. The baseline map is design-dependent, and the theorem does not assert robustness to a different covariate or assignment distribution.

A.7 Constrained kernel prediction

Proposition 2 (Constrained kernel prediction)**.** *For a fixed selected kernel k, define*

$$D_k(w, z) = \sup_{\|u\|_{\mathcal{H}_k} \le 1} |w^\top u_n - u(z)|, \quad u_n = (u(z_1), \ldots, u(z_n))^\top.$$

Then $D_k(w, z)^2 = k(z, z) - 2w^\top k_z + w^\top K w$, and the weights in Eq. (12) *uniquely solve*

$$\min_{F^\top w = f(x)} \; D_k(w, z)^2 + \eta \, \|w\|_2^2 . \tag{13}$$

Proof. By the reproducing property, the linear functional $u \mapsto w^\top u_n - u(z)$ has representer $\sum_i w_i k(z_i, \cdot) - k(z, \cdot)$. Its squared Hilbert norm is

$$D_k(w, z)^2 = k(z, z) - 2w^\top k_z + w^\top K w.$$

After removing the constant $k(z, z)$, the objective in Eq. (13) is $w^\top V w - 2w^\top k_z$. Strict positive definiteness gives a unique constrained minimizer. Its stationarity and constraint equations are

$$Vw - k_z + F\lambda = 0, \qquad F^\top w = f(x).$$

Solving gives

$$w = V^{-1} k_z + V^{-1} F G^{-1} \{ f(x) - F^\top V^{-1} k_z \},$$

which equals Eq. (12). For any feasible perturbation v with $F^\top v = 0$, the objective increases by $v^\top V v$ from this solution. This also verifies uniqueness directly.

A.8 Reference kernel and proof of Theorem 2

Squaring and integrating the pointwise bound yields

$$\mathcal{E}_Q(\widehat{\mu}) \le 3B^2 \int D_{\bar{k}}(w_z, z)^2 dQ + 3 \int (w_z^\top r_n - r(z))^2 dQ + 3 \int (w_z^\top \epsilon)^2 dQ. \tag{14}$$

For each retained rank $r < d_a$, the eigenvalues of $C = U_r U_r^\top / r$ are $1/r$ on the selected subspace and zero elsewhere; for full rank they are all $1/d_a$. Thus $0 \preceq C \preceq I_{d_a}$. At a fixed ℓ, the difference between the full-direction kernel and $k_{\ell, C}$ is $k_{X,\ell}(x, x') b(a)^\top (I - C) b(a')$, which is positive semidefinite. Adding the other lengthscale kernels proves $k_{\ell, C} \preceq \bar{k}$.

Baseline reproduction gives the exact samplewise identity

$$\widehat{\mu}(z) - \mu_0(z) = w_z^\top u_{0,n} - u_0(z) + w_z^\top r_n - r(z) + w_z^\top \epsilon. \tag{15}$$

For the realized weights, the reproducing property in $\mathcal{H}_{\bar{k}}$ yields $|w_z^\top u_{0,n} - u_0(z)| \le B D_{\bar{k}}(w_z, z)$. The triangle inequality proves Eq. (10). Squaring with $(a+b+c)^2 \le 3(a^2+b^2+c^2)$ and integrating proves Eq. (14). No independence between weights and noise is used. The representation assumption and integrability must hold for the realized design; the result does not supply a uniform bound on B or the remainder over all designs.

A.9 Risk for a construction independent of fitting noise

For comparison, suppose the kernel and weights are fixed independently of the current fitting noise, conditional on the design, and $\mathbb{E}[\epsilon \mid z_{1:n}] = 0$, $\mathrm{Cov}(\epsilon \mid z_{1:n}) \preceq \sigma^2 I$. If $\mu_0 = f^\top \beta_0 + u_0 + r$ with $\|u_0\|_{\mathcal{H}_k} \le B$ for this fixed kernel, Eq. (15) implies

$$\mathbb{E}_\epsilon[(\widehat{\mu}(z) - \mu_0(z))^2 \mid z_{1:n}] \le \{B D_k(w_z, z) + |w_z^\top r_n - r(z)|\}^2 + \sigma^2 \|w_z\|_2^2 . \tag{16}$$

The bias–noise cross term vanishes because the weights are fixed and the noise is centered. When $r = 0$ and the noise covariance is $\sigma^2 I$, the worst-case expression over the RKHS ball is $B^2 D_k(w_z, z)^2 + \sigma^2 \|w_z\|_2^2$. It is attained in the direction of the functional representer, unless that representer is zero, in which case its bias term vanishes. With $\eta = \sigma^2/B^2$ and $B > 0$, Proposition 2 minimizes this expression among baseline-reproducing linear predictors.

BCA-GP selects both its kernel and η from outcomes, so Eq. (16) is not asserted for its complete fitting procedure. The deterministic bound in Theorem 2 retains the adaptive noise term. Establishing a stochastic rate would require control of selection, approximation, and weight concentration.

A.10 Change of measure and ADRF averaging

If $q/p_0 \le W$, then, holding any fitted function fixed,

$$\int (\widehat{\mu} - \mu_0)^2 q \, dP_X \, da \le W \int (\widehat{\mu} - \mu_0)^2 p_0 \, dP_X \, da.$$

The same inequality holds after expectation over the training sample. This proves the target-support statement in the main text.

For uniform-dose risk, Jensen's inequality gives

$$\int_0^1 \{\mathbb{E}_X[\widehat{\mu}(X, a) - \mu_0(X, a)]\}^2 da \le \mathcal{E}_Q(\widehat{\mu}).$$

The empirical analogue holds for the benchmark metrics because both use the same test profiles and dose grid. ADRF error can therefore be much smaller than CDRF error even when individual predictions are poor; it does not replace the individual-level metric.

B Implementation details

B.1 Notation and experimental settings

The main text uses m_0 for baseline landmarks, m_1 for direction-learning landmarks, m for their retained feature dimension, and d_a for the dose-basis dimension. Table A1 gives their benchmark values and the candidate sets $\mathcal{L}$, Λ, $\mathcal{R}$, and $\mathcal{H}$. We use a response-scale floor $\varepsilon_{\mathrm{scale}} = 10^{-10}$ and evaluate on $N_a = 64$ dose points. These settings are fixed across the reported BCA-GP benchmarks.

B.2 Covariate and dose features

We use a unit-amplitude Matérn-3/2 covariate kernel with $d(x, x') = \|\mathrm{scale}(x) - \mathrm{scale}(x')\|_2 / \sqrt{d}$ and $k_{X,\ell}(x, x') = (1 + \sqrt{3}\, d(x, x')/L_\ell) \exp\{-\sqrt{3}\, d(x, x')/L_\ell\}$. Here $L_\ell = \ell L_0$, and L_0 is the

median positive pairwise distance among at most 300 training profiles sampled using the model seed. All covariates are retained, without PCA.

For the baseline, sample $m_0 = \min(16, \lfloor n/3 \rfloor)$ landmarks without replacement. If the landmark Gram matrix has eigendecomposition $U\Lambda U^\top$, retain eigenvalues larger than 10^{-9} and form $Z_0 = K_{X,L}U\Lambda^{-1/2}$. Let $R = [\mathbf{1}, Z_0] = U_R S_R V_R^\top$. Retain singular values greater than 10^{-10} times the largest and set $F = RV_R S_R^{-1}$. The corresponding out-of-sample feature map uses the same landmarks and right transformation.

The dose map starts from cubic splines with seven knots at $0, 1/6, \ldots, 1$, omitting the redundant bias column. Center the resulting eight columns using 100-point Gauss–Legendre weights and whiten their weighted covariance, flooring eigenvalues at 10^{-12}. Constant extrapolation is configured, although benchmark doses remain in $[0, 1]$. These numerical operations approximate centering and whitening under uniform integration on $[0, 1]$: ideally, $\int_0^1 b(a)\, da = 0$ and $\int_0^1 b(a)b(a)^\top\, da = I_{d_a}$. Finite quadrature and eigenvalue flooring need not yield these identities exactly under Lebesgue integration. The positive-semidefinite kernel construction and baseline-equivariance results do not require exact Lebesgue orthogonality; the error bounds use the resulting feature map and its induced kernel.

Outcome normalization. The residual-based scale and normalized response are

$$e = M_F y, \qquad s = \max\{\|e\|_2/\sqrt{n}, \varepsilon_{\text{scale}}\}, \qquad \widetilde{y} = y/s.$$

The floor prevents division by zero, and the scale depends only on the baseline residual. Fitting on $\widetilde{y}$ and multiplying predictions by s gives the original-outcome-scale predictor in the main text.

B.3 Adaptive selection and numerical solves

The candidate grid is shown in Table A1. Dose-direction landmarks use the model seed plus 19. Ridge fitting adds a numerical 10^{-10} diagonal term to the specified direction penalty. Rank eight uses $I_8/8$ and is evaluated only once per lengthscale/noise pair, since it is independent of the direction penalty. There are $3(2 \cdot 3 + 1)3 = 63$ candidates. Selection and final fitting both use all training observations; no selection subset is used in the reported experiments. The quadratic form in Eq. (6) is floored at 10^{-15}, and all covariance and information solves use Cholesky factorization.

Table A1: BCA-GP configuration used on all five benchmarks.

Component	Values
Baseline landmarks m_0	16 in the benchmarks; $p \le m_0 + 1$
Direction landmarks m_1	32 in the benchmarks
Dose dimension d_a	8 centered, whitened cubic-spline features
Length multipliers $\mathcal{L}$	$1, 3, 10$
Direction penalties Λ	$10^{-2}, 10^{-4}$
Dose ranks $\mathcal{R}$	$1, 2, 4, 8$
Noise parameters $\mathcal{H}$	$10^{-2}, 10^{-1}, 1$
Selection criterion	Full-training profile score, Eq. (6)

The core implementation is `BCAGP` in `bca_gp.py`, with shared feature construction in `basis.py`. Prediction uses the factorization of Eq. (3) to evaluate all requested doses efficiently. Distance caching changes computation only; it does not sparsify the final GP. Each candidate requires a dense factorization, giving $O(Hn^3)$ factorization time for $H = 63$ candidates and $O(n^2)$ covariance memory, apart from cached distances and feature work. Low dose rank does not remove the dense sample-size cost.

C BENCHMARK SOURCES AND TRAINING PROTOCOLS

C.1 DATA AND REPEATED SPLITS

IHDP uses the VCNet author's generator with PyTorch and NumPy seed zero, producing one matrix with 747 subjects and 25 covariates. The author's ten random-permutation splits each contain 471 training and 276 test subjects. NEWS uses the distributed GIKS continuous-treatment matrix with 2,993 subjects and 498 features, with the author's ten 2,000/993 splits. These repeats share one generated response matrix within each dataset.

TCGA uses the publicly linked `tcga.p` asset with 9,659 subjects and 4,000 features and the GIKS `TCGA_Data` generator. Each treatment type defines a separate benchmark. For repeat s, the generator seed is $100s$, both selection-bias parameters are 2, and the test fraction is 0.2, giving 7,728/1,931 training/test subjects. These repeats change structural weights, treatment and noise draws, and splits. Raw generated outcome units are retained. Shared arrays retain generator precision; author neural models use their adapter's native dtype.

All methods receive the same outer split in a case. Their internal validation fractions and epoch budgets differ. Each training process saves its full predicted test surface before loading evaluation truth; truth comes from the corresponding author response function. Evaluation uses 64 equally spaced doses including both endpoints and arithmetic grid averages. This common metric need not equal a paper's original quadrature-based or ADRF-only metric.

C.2 MODEL CONFIGURATIONS

Base GP. The reference model uses the dense kernel $k_A + k_X + k_A k_X$ with unit-amplitude Matérn-$3/2$ components, standardized covariates, outcome normalization, and diagonal regularization 10^{-5}. Two lengthscales are optimized by L-BFGS-B over $[0.05, 20]$, with at most 60 iterations. It uses all training rows and all covariates. It differs from BCA-GP in kernel structure and selection as well as baseline treatment, so the comparison does not isolate a single component.

DRNet. The author model builder uses one treatment type, five dose strata, repeated dose inputs, two layers with eight units, batch size eight, and at most 150 epochs. The author learning-rate schedule is 0.05 for ten epochs and 0.005 thereafter. Validation uses 27% of outer-training data with patience 12. We use these general defaults rather than the author's hyperparameter search.

VCNet and VCNet+TR. IHDP uses the author IHDP loop for 800 epochs with a full batch of 471. NEWS uses the author NEWS loop with architecture 498–50–50, 800 epochs, batch size 500, SGD learning rate 5×10^{-4}, momentum 0.9, and weight decay 0.005. The TR coefficient is 0.5, with TR learning rate 5×10^{-4} and weight decay 10^{-5} on NEWS and TCGA. TCGA uses widths 4000–100–64 adapted from GIKS and retains the NEWS training settings. This is an explicit configuration transfer, not an official VCNet TCGA optimum. Prediction for VCNet+TR uses $Q + \mathrm{TR}/(g + 10^{-6})$ as in the author evaluation code.

TransTEE. IHDP and NEWS use the author's base attention branch, with embedding dimension ten, two heads, one attention layer, and 1,500 epochs. IHDP uses batch size 128; NEWS uses batch size 500 and learning rate 0.01 with a cosine schedule. TCGA uses the author dosage model with a single-treatment adapter, 100 projected covariates, hidden dimension 48, batch size 128, and 200 epochs of Adam with learning rate 0.001 and a cosine schedule. These are the branches without adversarial suffixes.

GIKS. The author training function runs 201 factual epochs followed by 201 GIKS epochs, using batch size 128 and validation-selected checkpoints. The same internal 30% validation split is retained across phases. Dataset-specific learning rates and GP/GI coefficients follow the archived launch configurations. The near-dose gradient-interpolation path remains active inside the uniform counterfactual loss. All model configuration records and adapter hashes accompany the experiment archive; no equal-budget hyperparameter search was performed.

C.3 ANONYMOUS CODE LINK

https://anonymous.4open.science/r/bcagp-F181/