

541 **A The GP-SHAP algorithm and discussion on computation techniques**

We present the complete algorithm for both GP-SHAP and BayesGP-SHAP in Algorithm 1.

---

**Algorithm 1** GP-SHAP / BayesGP-SHAP

---

**Input:** Posterior mean function  $\tilde{m}$ , posterior covariance function  $\tilde{k}$ , inducing locations  $\tilde{\mathbf{X}}$ , explanation instances  $\mathbf{X}$ , number of coalition samples  $n_Z$ , hyperparameter  $\lambda, n_0, \sigma_0^2$ , base kernel  $k$ , algorithm **algo**,

- 1: Compute  $n_I =$  number of inducing location,  $n =$  number of explanation instances,  $d =$  number of features.
- 2: Compute Cholesky decomposition on posterior covariance  $\mathbf{L}\mathbf{L}^\top = \tilde{\mathbf{K}}_{\tilde{\mathbf{X}}\tilde{\mathbf{X}}}$
- 3: Sample coalitions  $\mathcal{S} = \{S_1, \dots, S_{n_Z}\}$  from  $[d]$ , build binary matrix  $\mathbf{Z} = \{0, 1\}^{n_Z \times d}$  from  $\mathcal{S}$ , and compute weights  $\mathbf{W} = \text{diag}[w_1, \dots, w_{n_Z}]$  with  $w_i = \frac{d-1}{\binom{d}{|S_i|}|S_i|(d-|S_i|)}$ .
- 4: Compute  $\mathbf{A} = (\mathbf{Z}^\top \mathbf{W} \mathbf{Z})^{-1} \mathbf{Z}^\top \mathbf{W}$  ▷ Shape:  $d \times n_Z$
- 5: Compute  $\mathbf{B}(\mathbf{X}, \mathcal{S}) = [(\mathbf{K}_{\tilde{\mathbf{X}}_S \tilde{\mathbf{X}}_S} + \lambda I)^{-1} k_S(\tilde{\mathbf{X}}_S, \mathbf{X}_S) \text{ for } S \text{ in } \mathcal{S}]$  ▷ Shape:  $n_Z \times n_I \times n$
- 6: Compute  $\mathbf{Q}$  where  $\mathbf{Q}_{i,l,k} = \sum_j \mathbf{B}(\mathbf{X}, \mathcal{S})_{i,j,k} \mathbf{L}_{j,l}$  ▷ Shape:  $n_Z \times n \times n_I$
- 7: Compute  $\mathbf{R}$  where  $\mathbf{R}_{i,k,l} = \sum_j \mathbf{A}_{i,j} \mathbf{Q}_{j,k,l}$  ▷ Shape:  $d \times n \times n_I$
- 8: Compute  $\mathbf{V}$  where  $\mathbf{V}_{i,m,k,n} = \sum_{j,l} \mathbf{R}_{i,j,k} \mathbf{R}_{m,l,n}$  ▷ Shape:  $d \times d \times n \times n$
- 9: Compute  $\mathbf{E}$  where  $\mathbf{E}_{i,k} = \sum_j \mathbf{B}(\mathbf{X}, \mathcal{S})_{i,j,k} \tilde{m}(\tilde{\mathbf{X}}_j)$  ▷ Shape:  $n_Z \times n$
- 10: Compute  $\Phi = \mathbf{A} \mathbf{E}$  ▷ The mean stochastic Shapley values of shape  $d \times n$
- 11: **if algo = GP-SHAP then**
- 12:     **return** mean explanations  $\Phi$  and covariance  $\mathbf{V}$  between  $d$  features and  $n$  instances
- 13: **else if algo = BayesGP-SHAP then**
- 14:     Compute  $s^2 = \text{diag}((\mathbf{E} - \mathbf{Z}\Phi)^\top \mathbf{W}(\mathbf{E} - \mathbf{Z}\Phi)) + \text{diag}(\Phi^\top \Phi)$  ▷ Shape:  $n \times 1$
- 15:     Sample  $\sigma^2$  from Scaled-Inv- $\chi^2(n_0 + n_Z, \frac{n_0 \sigma_0^2 + n_Z s^2}{n_0 + n_Z})$  ▷ Shape:  $n \times 1$
- 16:     **return** mean explanations  $\Phi$  and covariance  $\mathbf{V} + (\mathbf{Z}^\top \mathbf{W} \mathbf{Z})^{-1} \sigma^2$
- 17: **end if**

---

542

543 **Computational considerations.** In terms of computational complexity, one of the most demanding  
544 operations in the algorithm is the computation of conditional mean embeddings in step 5. Instead of  
545 naively inverting an  $n \times n$  matrix, which would have a computational cost of  $\mathcal{O}(n^3)$ , we employ the  
546 conjugate gradient method to reduce the computation of the conditional mean embedding component  
547 to  $\mathcal{O}(n^2 a)$ , where  $a \ll n$  represents the number of conjugate gradient iterations. Additionally, to  
548 further reduce runtime, we utilize the variational sparse GP model [48]. This model learns a set of  
549 inducing locations  $\tilde{\mathbf{X}}$  with a size of  $n_I \ll n$ , which can be reused for the estimation of conditional  
550 mean embeddings in the algorithm. Consequently, the computation of the conditional expectation is  
551 reduced from  $\mathcal{O}(n^2 a)$  to  $\mathcal{O}(n_I^2 a)$ . Another computational burden arises from the computation of the  
552 full covariance matrix across  $d$  features and  $n$  instances, which requires storage of a  $n^2 d^2$  matrix.  
553 However, since the full covariance matrix can be factorized into the  $\mathbf{R}$  component from step 7 of the  
554 algorithm, we can store this low-rank component and compute covariances between specific instances  
555 when necessary. It is worth noting that this decomposition of the covariance matrix allows us to avoid  
556 redundant computations when computing the covariance component, as we no longer need to iterate  
557 over all possible coalitions twice. Finally, we can further speed up our computational by parallelising  
558 computation across the sub-sampled coalitions in step 5.

559 **B Proofs and derivations**

560 **B.1 Section 2 proofs: Stochastic Shapley values**

561 We include the full proof of the derivation of stochastic Shapley values for completeness. The proof  
 562 is analogous to the original work of Shapley's [1] but extended to random variable payoffs. Ma et al.  
 563 [16] has also proved the same theorem but used a different proving strategy. They started with the  
 564 solution and showed it satisfies the axioms and then prove uniqueness, whereas the following proof  
 565 starts from the characterisation of s-games and derive the solution from a bottom-up fashion.

566 To facilitate the proof, we first introduce the concept of stochastic symmetric game.

567 **Proposition 15** (s-symmetric games). *Let  $C$  be a real-valued random variables, then the symmetric*  
 568 *game  $\nu_{C,R}(S) := C\mathbf{1}[R \subseteq S]$  gets a stochastic shapley value as,*

$$\phi_i(\nu_{C,R}) = \frac{C}{r} \quad (16)$$

569 where  $r = |R|$ .

570 *Proof.* Take any  $i, j \in R$ , pick a permutation  $\pi \in \Pi(U)$  so that  $\pi R = R$  and  $\pi i = j$ , so the induced  
 571 game  $\pi\nu_{C,R} = \nu_{C,R}$ , and therefore by the s-symmetry axiom,

$$\phi_j(\nu_{C,R}) = \phi_i(\nu_{C,R}) \quad (17)$$

572 Now by the s-efficiency axiom,

$$C = \nu_{C,R}(R) = \sum_{j \in R} \phi_j(\nu_{C,R}) = r\phi_i(\nu_{C,R}) \quad (18)$$

573 for any  $i \in R$ . □

574 Now we can characterise the form of any stochastic game as follows:

575 **Proposition 16.** *All s-games with finite carrier can be written as a linear combination of s-symmetric*  
 576 *games,*

$$\nu = \sum_{R \subseteq N, R \neq \emptyset} \nu_{c_R(\nu), R} \quad (19)$$

577 where

$$C_R(\nu) = \sum_{T \subseteq R} (-1)^{r-t} \nu(T) \quad (20)$$

578 *Proof.* We start by verifying

$$\nu(S) = \sum_{R \subseteq N, R \neq \emptyset} \nu_{c_R(\nu), R}(S) \quad (21)$$

579 holds for all  $S \subseteq U$ , and for any finite carrier  $N$  of  $\nu$ . If  $S \subseteq N$ , then we can rewrite the expression  
 580 as,

$$\nu(S) = \sum_{R \subseteq S} \sum_{T \subseteq R} (-1)^{r-t} \nu(T) \quad (22)$$

$$= \sum_{T \subseteq S} \sum_{T \subseteq R \subseteq S} (-1)^{r-t} \nu(T) \quad (23)$$

$$= \sum_{T \subseteq S} \nu(T) \sum_{r=t}^s (-1)^{r-t} \binom{s-t}{r-t} \quad (24)$$

$$= \nu(S) \quad (25)$$

581 where in the last equation we used the fact that  $\sum_{r=t}^s (-1)^{r-t} \binom{s-t}{r-t}$  is a binomial expansion of  
 582  $(1 + (-1))^{s-t}$ , therefore the only non-zero expression is when  $t = s$ .

583 □

584 We can now prove the uniqueness of stochastic Shapley values,

585 **Theorem 4** (Stochastic Shapley values). *The only stochastic value allocation  $\phi$  of  $\nu$  satisfying*  
 586 *s-symmetry, s-efficiency, and s-linearity takes the following form,*

$$\phi_i(\nu) = \sum_{S \subseteq N \setminus \{i\}} c_{|S|} (\nu(S \cup i) - \nu(S)) \quad (1)$$

587 where  $N$  is the smallest carrier set of  $\Omega$ ,  $c_{|S|} = \frac{1}{|N|} \binom{|N|-1}{|S|}^{-1}$  and  $\phi_i(\nu)$  is the  $i^{\text{th}}$  SSV of  $s$ -game  $\nu$ .

*Proof.* First, let us denote

$$\gamma_i(S) := \sum_{\substack{R \subseteq N \\ S \cup \{i\} \subseteq R}} (-1)^{r-s} \frac{1}{r}.$$

588 Applying the s-linearity axiom on  $\phi$  to the characterisation of  $\nu$  from the previous propositions leads  
 589 us to the following,

$$\phi_i(\nu) = \phi_i \left( \sum_{R \subseteq N, R \neq \emptyset} \nu_{C_R(\nu), R} \right) \quad (26)$$

$$= \sum_{R \subseteq N, R \neq \emptyset} \phi_i(\nu_{C_R(\nu), R}) \quad (27)$$

$$= \sum_{R \subseteq N, i \in R} c_R(\nu) \frac{1}{r} \quad (28)$$

$$= \sum_{R \subseteq N, i \in R} \frac{1}{r} \left( \sum_{S \subseteq R} (-1)^{r-s} \nu(S) \right) \quad (29)$$

$$= \sum_{S \subseteq N} \sum_{\substack{R \subseteq N \\ S \cup \{i\} \subseteq R}} (-1)^{r-s} \nu(S) \frac{1}{r} \quad (30)$$

$$= \sum_{S \subseteq N} \gamma_i(S) \nu(S) \quad (31)$$

$$= \sum_{\substack{S \subseteq N \\ i \in S}} \gamma_i(S) \nu(S) + \gamma_i(S - \{i\}) \nu(S - \{i\}) \quad (32)$$

$$= \sum_{\substack{S \subseteq N \\ i \in S}} \gamma_i(S) (\nu(S) - \nu(S - \{i\})) \quad (33)$$

$$= \sum_{\substack{S \subseteq N \\ i \in S}} \frac{(s-1)!(n-s)!}{n!} (\nu(S) - \nu(S - \{i\})) \quad (34)$$

$$= \sum_{S \subseteq N \setminus \{i\}} c_{|S|} (\nu(S \cup i) - \nu(S)) \quad (35)$$

590 where in (32) we used the following observation: given  $i \notin S' \subseteq N$ , and  $S = S' \cup \{i\}$ , then  
 591  $\gamma_i(S) = -\gamma_i(S')$ .

592 It satisfies uniqueness by construction.  $\square$

593 **Proposition 5.** *Given the player set  $\Omega$ , let  $\nu$  be a stochastic game,  $\phi$  a stochastic Shapley value*  
 594 *allocation, and  $\bar{\phi}$  a deterministic Shapley value allocation. Suppose that  $\mathbb{E}[\nu]$  and  $\mathbb{V}[\nu]$  are the*  
 595 *corresponding mean and variance  $d$ -games, respectively. Then,  $\mathbb{E}[\phi(\nu)] = \bar{\phi}(\mathbb{E}[\nu])$ , but  $\mathbb{V}[\phi(\nu)] \neq$*   
 596  *$\bar{\phi}(\mathbb{V}[\nu])$ . In particular, the SSV variance is given by*

$$\mathbb{V}[\phi_i(\nu)] = \sum_{S \subseteq N \setminus \{i\}} \sum_{S' \subseteq N \setminus \{i\}} c_{|S|} c_{|S'|} (\mathbb{C}[\nu_{S \cup i}, \nu_{S' \cup i}] - \mathbb{C}[\nu_{S \cup i}, \nu_{S'}] - \mathbb{C}[\nu_S, \nu_{S' \cup i}] + \mathbb{C}[\nu_S, \nu_{S'}]),$$

597 where  $\nu_S = \nu(S)$  and  $\mathbb{C}$  is the covariance function between the stochastic payoffs.

598 *Proof.* The equivalence between mean of stochastic Shapley values and deterministic Shapley values  
599 of mean game is trivial to show leveraging the linearity of expectation. The variance of  $\mathbb{V}[\phi_i(\nu)]$  can  
600 be shown by repeatedly applying the standard identity  $\mathbb{V}[X + Y] = \mathbb{V}[X] + \mathbb{V}[Y] + 2\mathbb{C}[X, Y]$  for  
601 random variables  $X, Y$ . Now consider the deterministic Shapley values of variance game  $\mathbb{V}[\nu]$ ,

$$\bar{\phi}_i[\mathbb{V}[\nu(\cdot)]] = \sum_{S \subseteq N \setminus \{i\}} c_{|S|} (\mathbb{V}[\nu(S \cup i)] - \mathbb{V}[\nu(S)]) \quad (36)$$

602 Comparing to the expression of  $\mathbb{V}[\phi_i(\nu)]$  from the lemma,

$$\mathbb{V}[\phi_i(\nu)] = \sum_{S \subseteq N \setminus \{i\}} \sum_{S' \subseteq N \setminus \{i\}} c_{|S|} c_{|S'|} (\mathbb{C}[\nu_{S \cup i}, \nu_{S' \cup i}] - \mathbb{C}[\nu_{S \cup i}, \nu_{S'}] - \mathbb{C}[\nu_S, \nu_{S' \cup i}] + \mathbb{C}[\nu_S, \nu_{S'}]),$$

603 even if we assume mutual independence across all payoff random variables, leading to  $\mathbb{C}[\nu(S \cup$   
604  $i), \nu(S)] = 0$  for all  $S$ , we still would not subtract but instead sum the variance of  $\mathbb{V}[\nu(S \cup i)]$  and  
605  $\mathbb{V}[\nu(S)]$ . Therefore the variances of stochastic Shapley values is not the same as the deterministic  
606 Shapley values of the variance game.  $\square$

## 607 B.2 Section 3.1 proofs on the stochastic Shapley values for induced stochastic game from GP

608 **Proposition 6** (Stochastic game  $\nu_f$  as induced GP). *Let  $f \sim \mathcal{GP}(\tilde{m}, \tilde{k})$  with integrable sample paths,*  
609 *i.e.  $\int_{\mathcal{X}} |f| dp_X < \infty$  almost surely. The stochastic payoff function  $\nu_f$  induced by  $f$  is a Gaussian*  
610 *process with the following mean and covariance functions:*

$$m_\nu(\mathbf{x}, S) := \mathbb{E}_X[\tilde{m}(X) \mid X_S = \mathbf{x}_S], \quad (4)$$

$$k_\nu((\mathbf{x}, S), (\mathbf{x}', S')) := \mathbb{E}_{X, X'}[\tilde{k}(X, X') \mid X_S = \mathbf{x}_S, X'_{S'} = \mathbf{x}'_{S'}]. \quad (5)$$

611 *Proof.* This is a direct application of Chau et al. [18, Proposition 3.2] to the distribution  $P(X \mid X_S =$   
612  $\mathbf{x}_S)$ .  $\square$

613 **Theorem 7** (Stochastic Shapley values of  $\nu_f$ ). *Let  $\nu_f$  be an induced stochastic game from the GP*  
614  *$f \sim \mathcal{GP}(\tilde{m}, \tilde{k})$  and denote  $\mathbf{v}_\mathbf{x} := [\nu_f(\mathbf{x}, S_1), \dots, \nu_f(\mathbf{x}, S_{2^d})]^\top$  the vector of stochastic payoffs across*  
615 *all coalitions, then the corresponding stochastic Shapley values  $\phi(\nu_f(\mathbf{x}, \cdot))$  follows a  $d$ -dimensional*  
616 *multivariate Gaussian distribution,*

$$\phi(\nu_f(\mathbf{x}, \cdot)) \sim \mathcal{N}(\mathbf{A}\mathbb{E}[\mathbf{v}_\mathbf{x}], \mathbf{A}\mathbb{V}[\mathbf{v}_\mathbf{x}]\mathbf{A}^\top) \quad \text{with} \quad \mathbf{A} := (\mathbf{Z}^\top \mathbf{W} \mathbf{Z})^{-1} \mathbf{Z}^\top \mathbf{W}, \quad (6)$$

617 where  $\mathbb{E}[\mathbf{v}_\mathbf{x}] \in \mathbb{R}^{2^d}$  and  $\mathbb{V}[\mathbf{v}_\mathbf{x}] \in \mathbb{R}^{2^d \times 2^d}$  are the corresponding mean vector and covariance matrix  
618 of the payoffs.

619 *Proof.* Recall from Lundberg and Lee [2, Theorem 2], for deterministic Shapley values, given a  
620 deterministic payoff  $\bar{\nu}_\mathbf{x}$  for all  $2^d$  coalitions, the expression of Shapley values for each  $i \in [d]$ ,

$$\bar{\phi}_{\mathbf{x}i} = \sum_{S \subseteq [d] \setminus \{i\}} c_{|S|} (\bar{\nu}_f(S \cup i) - \bar{\nu}_f(S)) \quad (37)$$

621 can be written compactly as the following vector,

$$\bar{\phi}_\mathbf{x} = \mathbf{A} \bar{\mathbf{v}}_\mathbf{x}. \quad (38)$$

622 We can therefore similarly write down the form of the stochastic Shapley values using this linear  
623 operator  $\mathbf{A}$ , acting now on a vector of random variable output stochastic payoff vector  $\mathbf{v}_\mathbf{x}$ ,

$$\phi_\mathbf{x} = \mathbf{A} \mathbf{v}_\mathbf{x}. \quad (39)$$

624 Nonetheless, as Proposition 8 implies that  $\mathbf{v}_\mathbf{x}$  is a multivariate Gaussian, therefore  $\phi_\mathbf{x}$  is also  
625 multivariate Gaussian with mean and covariance the following,

$$\mathbf{v}_\mathbf{x} \sim \mathcal{N}(\mathbf{A}\mathbb{E}[\mathbf{v}_\mathbf{x}], \mathbf{A}\mathbb{V}[\mathbf{v}_\mathbf{x}]\mathbf{A}^\top). \quad (40)$$

626  $\square$

627 **B.3 Section 3.2 proofs on estimation**

628 To proceed, we first introduce the concepts of conditional mean embedding as a tool to estimate  
629 conditional expectation of functions living in their corresponding RKHSs,

630 **Definition 17** (Conditional mean embedding [38]). *Let  $X, Y$  be random variables and  $k : \mathcal{X} \rightarrow \mathcal{X} \rightarrow$   
631  $\mathbb{R}$  a kernel on  $X$ , then we define the following as the conditional mean embedding of  $p(X | Y = y)$ ,*

$$\mu_{X|Y=y} := \int k(\cdot, X) d\mathbb{P}(X | Y = y) \quad (41)$$

632 **Proposition 18** (Conditional Mean estimation). *For random variable  $X, Y$ , and a kernel  $k : \mathcal{X} \rightarrow$   
633  $\mathcal{X} \rightarrow \mathbb{R}$  on  $\mathcal{X}$  and a kernel  $l : \mathcal{Y} \rightarrow \mathcal{Y} \rightarrow \mathbb{R}$  on  $\mathcal{Y}$ . Given observations  $\mathbf{D} = \{\mathbf{X}, \mathbf{y}\}$ , the empirical  
634 conditional mean embedding can be estimated as*

$$\hat{\mu}_{X|Y=y} = l(y, \mathbf{y}) (\mathbf{L}_{\mathbf{y}\mathbf{y}} + \lambda I)^{-1} k(\mathbf{X}, \cdot), \quad (42)$$

635 where  $l(y, \mathbf{y}) = [l(y, y_1), \dots, l(y, y_n)]^\top$  and  $k(\cdot, \mathbf{X}) = [k(\cdot, \mathbf{x}_1), \dots, k(\cdot, \mathbf{x}_n)]^\top$ , the parameter  
636  $\lambda > 0$  is there to stabilise the inversion. Now for  $f \in \mathcal{H}_k$ , the conditional expectation can then be  
637 estimated as,

$$\hat{\mathbb{E}}[f(X) | Y = y] = \langle \hat{\mu}_{X|Y=y}, f \rangle \quad (43)$$

$$= l(y, \mathbf{y}) (\mathbf{L}_{\mathbf{y}\mathbf{y}} + \lambda I)^{-1} \mathbf{f}, \quad (44)$$

638 where  $\mathbf{f} = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^\top$ .

639 *Proof.* This is standard result from literature, please read Song et al. [49], Muandet et al. [38] for  
640 more details.  $\square$

641 Now we can apply these propositions to estimate the mean and covariance functions of the induced  
642 stochastic game from GP,

643 **Proposition 8** (Estimating  $\nu_f$ ). *Given  $\mathbf{D} = (\mathbf{X}, \mathbf{y})$  and the posterior GP  $f | \mathbf{D} \sim \mathcal{GP}(\tilde{m}, \tilde{k})$ , the  
644 mean and covariance function of the stochastic cooperative game  $\nu_f$  can be estimated as,*

$$\hat{m}_\nu(\mathbf{x}, S) = \mathbf{b}(\mathbf{x}, S)^\top \tilde{m}(\mathbf{X}), \quad \hat{k}_\nu((\mathbf{x}, S), (\mathbf{x}', S')) = \mathbf{b}(\mathbf{x}, S)^\top \tilde{\mathbf{K}}_{\mathbf{X}\mathbf{X}} \mathbf{b}(\mathbf{x}', S'), \quad (7)$$

645 where  $\mathbf{b}(\mathbf{x}, S) := (\mathbf{K}_{\mathbf{X}_S \mathbf{X}_S} + \lambda I)^{-1} k_S(\mathbf{X}_S, \mathbf{x}_S)$ ,  $\tilde{m}(\mathbf{X}) = [\tilde{m}(\mathbf{x}_1), \dots, \tilde{m}(\mathbf{x}_n)]^\top$ , and  $k_S :$   
646  $\mathcal{X}_S \times \mathcal{X}_S \rightarrow \mathbb{R}$  is the kernel defined on the sub-feature space of  $\mathcal{X}$  and we write  $k_S(\mathbf{x}_S, \mathbf{X}_S) :=$   
647  $[k_S(\mathbf{x}_S, \mathbf{x}_{1S}), \dots, k_S(\mathbf{x}_S, \mathbf{x}_{nS})]$  and  $\mathbf{K}_{\mathbf{X}\mathbf{X}}$  and  $\tilde{\mathbf{K}}_{\mathbf{X}\mathbf{X}}$  as the gram matrix of  $\mathbf{X}$  using kernel  $k$  and  $\tilde{k}$   
648 respectively. The parameter  $\lambda > 0$  is a fixed hyperparameter to stabilise the inversion.

649 *Proof.* Without loss of generality, we will demonstrate this proposition with  $\tilde{m}, \tilde{k}$  obtained via  
650 standard GP regression, i.e.,

$$\tilde{m}(\mathbf{x}) = k(\mathbf{x}, \mathbf{X}) (\mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma^2 I)^{-1} \mathbf{y} \quad (45)$$

$$\tilde{k}(\mathbf{x}, \mathbf{x}') = k(\mathbf{x}, \mathbf{x}') - k(\mathbf{x}, \mathbf{X}) (\mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma^2 I)^{-1} k(\mathbf{X}, \mathbf{x}'). \quad (46)$$

651 Starting with the mean function,

$$\mathbb{E}[\tilde{m}(X) | X_S = \mathbf{x}_S] = \mathbb{E}_X [k(X, \mathbf{X}) (\mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma^2 I)^{-1} \mathbf{y} | X_S = \mathbf{x}_S] \quad (47)$$

$$= \langle k(\cdot, \mathbf{X}) (\mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma^2 I)^{-1} \mathbf{y}, \mu_{X|X_S=\mathbf{x}_S} \rangle_{\mathcal{H}_k}. \quad (48)$$

652 We can replace the population conditional mean embedding with the empirical version, and expand,

$$\hat{\mathbb{E}}[\tilde{m}(X) | X_S = \mathbf{x}_S] = \langle k(\cdot, \mathbf{X}) (\mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma^2 I)^{-1} \mathbf{y}, \hat{\mu}_{X|X_S=\mathbf{x}_S} \rangle_{\mathcal{H}_k} \quad (49)$$

$$= k_S(\mathbf{X}_S, \mathbf{x}_S) (\mathbf{K}_{\mathbf{X}_S \mathbf{X}_S} + \lambda I)^{-1} \mathbf{K}_{\mathbf{X}\mathbf{X}} (\mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma^2 I)^{-1} \mathbf{y} \quad (50)$$

$$= \mathbf{b}(\mathbf{x}, S)^\top \tilde{m}(\mathbf{X}). \quad (51)$$

653 Analogously, the conditional expectation of the posterior covariance function, i.e.,  $\mathbb{E}[\tilde{k}(X, X') |$   
654  $X_S = \mathbf{x}_S, X'_S = \mathbf{x}'_S]$ , can be estimated following the steps above,

$$\mu_{X|X_S=\mathbf{x}_S}^\top \mu_{X'|X'_S=\mathbf{x}'_S} - \mu_{X|X_S=\mathbf{x}_S}^\top k(\cdot, \mathbf{X}) (\mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma^2 I)^{-1} k(\mathbf{X}, \cdot) \mu_{X'|X'_S=\mathbf{x}'_S}. \quad (52)$$

655 After replacing the population conditional mean embedding as their empirical estimates, we can  
656 arrive at the solution.  $\square$

657 **Proposition 9 (GP-SHAP).** *Let the matrix  $\mathbf{A}$  be defined as in Theorem 7. The mean and covariance*  
 658 *for the multivariate stochastic Shapley values can be estimated as,*

$$\phi(\hat{\nu}_f(\mathbf{x}, \cdot)) = \mathcal{N}\left(\mathbf{A}\mathbf{B}(\mathbf{x}, [d])^\top \tilde{m}(\mathbf{X}), \mathbf{A}\mathbf{B}(\mathbf{x}, [d])^\top \tilde{\mathbf{K}}_{\mathbf{X}\mathbf{X}}\mathbf{B}(\mathbf{x}, [d])\mathbf{A}^\top\right) \quad (8)$$

659 where  $\mathbf{B}(\mathbf{x}, [d]) = [\mathbf{b}(\mathbf{x}, [d]_1), \dots, \mathbf{b}(\mathbf{x}, [d]_{2^d})]^\top$ .

660 *Proof.* The result follows directly from the previous proposition. Recall  $\phi(\hat{\nu}_f(\mathbf{x}, \cdot)) = \mathbf{A}\hat{\mathbf{v}}_{\mathbf{x}}$  for  $\hat{\mathbf{v}}_{\mathbf{x}}$   
 661 the vector of stochastic payoffs for each coalition. To estimate the mean, we

$$\mathbb{E}[\phi(\hat{\nu}_f(\mathbf{x}, \cdot))] = \mathbf{A}\mathbb{E}[\hat{\mathbf{v}}_{\mathbf{x}}] \quad (53)$$

$$= \mathbf{A} \begin{bmatrix} \hat{m}_{\nu}(\mathbf{x}, S_1) \\ \vdots \\ \hat{m}_{\nu}(\mathbf{x}, S_{2^d}) \end{bmatrix} \quad (54)$$

$$= \mathbf{A} \begin{bmatrix} \mathbf{b}(\mathbf{x}, S_1)^\top \tilde{m}(\mathbf{X}) \\ \vdots \\ \mathbf{b}(\mathbf{x}, S_{2^d})^\top \tilde{m}(\mathbf{X}) \end{bmatrix} \quad (55)$$

$$= \mathbf{A}\mathbf{B}(\mathbf{x}, [d])^\top \tilde{m}(\mathbf{X}). \quad (56)$$

662 Recall  $\mathbb{V}[\mathbf{v}_{\mathbf{x}}]_{i,j} = \hat{k}_{\nu}((\mathbf{x}, S_i), (\mathbf{x}, S_j)) = \mathbf{b}(\mathbf{x}, S_i)^\top \tilde{\mathbf{K}}_{\mathbf{X}\mathbf{X}}\mathbf{b}(\mathbf{x}, S_j)$ , the derivation for the covariance  
 663 matrix then follows analogously as the derivation for the mean,

$$\mathbb{V}[\phi(\hat{\nu}_f(\mathbf{x}, \cdot))] = \mathbf{A}\mathbb{V}[\hat{\mathbf{v}}_{\mathbf{x}}]\mathbf{A}^\top \quad (57)$$

$$= \mathbf{A} \left[ \mathbf{b}(\mathbf{x}, S_i)^\top \tilde{\mathbf{K}}_{\mathbf{X}\mathbf{X}}\mathbf{b}(\mathbf{x}, S_j) \right]_{i=1, j=1}^{2^d, 2^d} \mathbf{A}^\top \quad (58)$$

$$= \mathbf{A}\mathbf{B}(\mathbf{x}, [d])^\top \tilde{\mathbf{K}}_{\mathbf{X}\mathbf{X}}\mathbf{B}(\mathbf{x}, [d])\mathbf{A}^\top. \quad (59)$$

664

□

665 **Proposition 10 (BayesSHAP [20]).** *Given the data generation above, the posterior distribution on  $\bar{\phi}$*   
 666 *and  $\sigma^2$  follows:*

$$\bar{\phi} \mid \sigma^2, \mathbf{Z}_\ell, f, \mathbf{x}, \mathbf{D} \sim \mathcal{N}(\mathbf{A}_\ell \bar{\mathbf{v}}_{\mathbf{x}}, (\mathbf{Z}_\ell^\top \mathbf{W}_\ell \mathbf{Z}_\ell)^{-1} \sigma^2) \quad (11)$$

$$\sigma^2 \mid \mathbf{Z}_\ell, f, \mathbf{x}, \mathbf{D} \sim \text{Scaled-Inv-}\chi^2 \left( \ell_0 + \ell, \frac{\ell_0 \sigma_0^2 + \ell s^2(\bar{\mathbf{v}}_{\mathbf{x}})}{\ell_0 + \ell} \right) \quad (12)$$

667 where  $\ell$  is the number of coalitions  $\mathcal{S} = \{S_j\}_{j=1}^\ell$  we sample uniformly from  $2^{[d]}$ ,  $\mathbf{Z}_\ell$  is the binary  
 668 matrix representing  $\mathcal{S}$ , and  $\mathbf{W}_\ell$  is the corresponding weight matrix, and  $\mathbf{A}_\ell = (\mathbf{Z}_\ell^\top \mathbf{W}_\ell \mathbf{Z}_\ell)^{-1} \mathbf{Z}_\ell^\top \mathbf{W}_\ell$   
 669 is the WLS matrix,  $\bar{\mathbf{v}}_{\mathbf{x}} = [\bar{v}_f(\mathbf{x}, S_1), \dots, \bar{v}_f(\mathbf{x}, S_\ell)]^\top$  is the vector of deterministic payoffs, and

$$s^2(\bar{\mathbf{v}}_{\mathbf{x}}) = \frac{1}{\ell} [(\bar{\mathbf{v}}_{\mathbf{x}} - \mathbf{Z}_\ell \mathbf{A}_\ell \bar{\mathbf{v}}_{\mathbf{x}})^\top \mathbf{W}_\ell (\bar{\mathbf{v}}_{\mathbf{x}} - \mathbf{Z}_\ell \mathbf{A}_\ell \bar{\mathbf{v}}_{\mathbf{x}}) + (\mathbf{A}_\ell \bar{\mathbf{v}}_{\mathbf{x}})^\top (\mathbf{A}_\ell \bar{\mathbf{v}}_{\mathbf{x}})] \quad (13)$$

670 measures the average weighted error in the regression and the norm of the mean explanations.

671 *Proof.* See Slack et al. [20, Section. 3.1]. □

672 **Proposition 11 (BayesGP-SHAP).** *Continuing from Propositions 9 and 10, the posterior distribution*  
 673 *of the stochastic Shapley values can be estimated using the Bayesian WLS approach as,*

$$\phi \mid \sigma^2, \mathbf{Z}_\ell, \mathbf{x}, \mathbf{D} \sim \mathcal{N}\left(\mathbf{A}_\ell \mathbf{B}(\mathbf{x}, \mathcal{S})^\top \tilde{m}(\mathbf{X}), \mathbf{A}_\ell \mathbf{B}(\mathbf{x}, \mathcal{S})^\top \tilde{\mathbf{K}}_{\mathbf{X}\mathbf{X}}\mathbf{B}(\mathbf{x}, \mathcal{S})\mathbf{A}_\ell^\top + (\mathbf{Z}_\ell^\top \mathbf{W}_\ell \mathbf{Z}_\ell)^{-1} \sigma^2\right)$$

674 where  $\sigma^2$  is sampled from  $\sigma^2 \mid \mathbf{Z}_\ell \sim \text{Scaled-Inv-}\chi^2 \left( \ell_0 + \ell, \frac{\ell_0 \sigma_0^2 + \ell s^2(\mathbb{E}[\mathbf{v}_{\mathbf{x}}])}{\ell_0 + \ell} \right)$ .

675 *Proof.* We drop the bar notation of  $\bar{\phi}$  to unify notations. Given the posterior GP  $f \mid \mathbf{D} \sim \mathcal{GP}(\tilde{m}, \tilde{k})$

$$p(\phi \mid \sigma^2, \mathbf{Z}_\ell, \mathbf{x}, \mathbf{D}) = \int p(\phi \mid \sigma^2, \mathbf{Z}_\ell, f, \mathbf{x}, \mathbf{D}) p(f \mid \mathbf{D}) df \quad (60)$$

676 Using a standard Gaussian conjugacy procedure, we can derive the variance as the sum of variances  
677 from GP-SHAP and BayesSHAP. While it is possible to integrate  $p(\sigma^2 \mid \mathbf{Z}_\ell, f, \mathbf{x}, \mathbf{D})$  with respect to  
678 the posterior, this leads to a complex scaled mixture of normals that is difficult to model. Instead,  
679 we construct a scaled inverse chi-square distribution with  $s^2 \mathbb{E}[\mathbf{b}_\mathbf{x}]$ , which represents the error of  
680 the weighted regression with respect to the mean payoffs  $\mathbb{E}[\mathbf{v}_\mathbf{x}]$ . We sample  $\sigma^2$  from the following  
681 distribution:

$$\sigma^2 \mid \mathbf{Z}_\ell, \mathbf{x}, \mathbf{D} \sim \text{Scale-Inv-}\chi^2 \left( \ell_0 + \ell, \frac{\ell_0 \sigma_0^2 + \ell s^2 (\mathbb{E}[\mathbf{v}_\mathbf{x}])}{\ell_0 + \ell} \right). \quad (61)$$

682

□

#### 683 B.4 Proofs for section 4 on predictive explanation and Shapley prior

684 **Proposition 12** (The Shapley prior over  $\phi$ ). *The prior  $f \sim \mathcal{GP}(0, k)$  and the corresponding stochastic  
685 game  $\nu_f(\mathbf{x}, S) = \mathbb{E}[f(X) \mid X_S = \mathbf{x}_S]$  induce a vector-valued GP prior over the explanation  
686 functions  $\phi \sim \mathcal{GP}(0, \kappa)$  where  $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^{d \times d}$  is the matrix-valued covariance kernel*

$$\kappa(\mathbf{x}, \mathbf{x}') = \mathcal{A}(\mathbf{x})^\top \mathcal{A}(\mathbf{x}'), \quad \mathcal{A}(\mathbf{x}) = \Psi(\mathbf{x}) \mathbf{A}^\top \quad (14)$$

687 where  $\Psi(\mathbf{x}) = [\mathbb{E}[k(\cdot, X) \mid X_{S_1} = x_{S_1}], \dots, \mathbb{E}[k(\cdot, X) \mid X_{S_{2d}} = x_{S_{2d}}]]$ .

688 *Proof.* The proof is similar to how we proved previous propositions but applied to prior GP  $f \sim$   
689  $\mathcal{GP}(0, k)$  instead. If we set,

$$\nu_f(\mathbf{x}, S) = \mathbb{E}[f(X) \mid X_S = \mathbf{x}_S], \quad (62)$$

690 then  $\nu_f$  is a GP on the joint space of data and coalitions with mean 0, and covariance function,

$$\text{cov}(\nu_f(\mathbf{x}, S), \nu_f(\mathbf{x}', S')) = \mathbb{E}[k(X, X') \mid X_S = \mathbf{x}_S, X'_{S'} = \mathbf{x}'_{S'}] \quad (63)$$

$$= \mu_{X|X_S=\mathbf{x}_S}^\top \mu_{X'|X_{S'}=\mathbf{x}'_{S'}}. \quad (64)$$

691 Since  $\phi = \mathbf{A} \mathbf{v}_\mathbf{x}$  for  $\mathbf{v}_\mathbf{x}$  the vector of stochastic payoff from the game induced by the GP prior, the  
692 mean stays 0, and the covariance is,

$$\kappa(\mathbf{x}, \mathbf{x}') = \mathbf{A} \left[ \mu_{X|X_{S_i}=\mathbf{x}_{S_i}}^\top \mu_{X'|X_{S'_j}=\mathbf{x}'_{S'_j}} \right]_{i=1, j=1}^{2^d, 2^d} \mathbf{A}^\top \quad (65)$$

$$= \mathbf{A} \Psi(\mathbf{x})^\top \Psi(\mathbf{x}') \mathbf{A}^\top \quad (66)$$

$$= \mathcal{A}(\mathbf{x})^\top \mathcal{A}(\mathbf{x}'), \quad (67)$$

693 therefore we have a matrix-valued covariance kernel  $\kappa$  to build a prior over the induced Shapley  
694 values. □

695 **Proposition 13** (Predictive explanations as multi-output GPs). *Given  $\mathbf{D}_\phi = \{(\mathbf{x}_i, \phi_i)\}_{i=1}^n =$   
696  $(\mathbf{X}, \Phi_\mathbf{X})$  where  $\phi_i \in \mathbb{R}^d$  are the Shapley values computed under predictive model  $f$  and  $\Phi_\mathbf{X} =$   
697  $[\phi_1, \dots, \phi_n]^\top$ , the predictive explanations for new data  $\mathbf{x}'$  is distributed as,*

$$\phi(\mathbf{x}') \mid \mathbf{D}_\phi \sim \mathcal{N}(\tilde{m}_\phi(\mathbf{x}'), \quad \kappa(\mathbf{x}', \mathbf{x}') - \kappa(\mathbf{x}', \mathbf{X}) b_\kappa(\mathbf{x}', \mathbf{X})) \quad (15)$$

698 where  $\tilde{m}_\phi(\mathbf{x}') = b_\kappa(\mathbf{x}', \mathbf{X})^\top \text{vec}(\Phi_\mathbf{X})$ ,  $b_\kappa(\mathbf{x}', \mathbf{X}) := (\mathcal{K}_{\mathbf{X}\mathbf{X}} + \sigma_\phi^2 I)^{-1} \kappa(\mathbf{X}, \mathbf{x}')$ ,  $\mathcal{K}_{\mathbf{X}\mathbf{X}}$  is the gram  
699 matrix for kernel  $\kappa$  of size  $nd \times nd$ ,  $\kappa(\mathbf{x}', \mathbf{X}) = [\kappa(\mathbf{x}', \mathbf{x}_1), \dots, \kappa(\mathbf{x}', \mathbf{x}_n)]$  is of size  $d \times nd$  and  $\sigma_\phi^2$   
700 is the noise parameter for regression.

701 *Proof.* Follows from standard vector-valued Gaussian process regression results. See Alvarez et al.  
702 [50] for a detailed discussion on regression with matrix-valued kernels. □

703 **Proposition 14** (Posterior mean as Shapley values for payoff vector  $\tilde{\mathbf{v}}_{\mathbf{x}'}$ ). *The posterior mean  
704  $\tilde{m}_\phi(\mathbf{x}')$  corresponds to Shapley values for the payoff vector  $\tilde{\mathbf{v}}_{\mathbf{x}'}$ , i.e.,  $\tilde{m}_\phi(\mathbf{x}') = \mathbf{A} \tilde{\mathbf{v}}_{\mathbf{x}'}$ , where  
705  $\tilde{\mathbf{v}}_{\mathbf{x}'} = \sum_{i=1}^n \Psi(\mathbf{x}')^\top \Psi(\mathbf{x}_i) \mathbf{A}^\top \alpha_i$  and  $\alpha_i \in \mathbb{R}^d$  is the  $[i, \dots, i + (d - 1)]$  subvector of  $(\mathcal{K}_{\mathbf{X}\mathbf{X}} +$   
706  $\sigma_\phi^2 I)^{-1} \text{vec}(\Phi_\mathbf{X})$ .*

707 *Proof.* There are two ways to see this. First is by brute force and rearranging the terms in the posterior  
 708 mean expression. The other is to leverage the vector-valued representer theorem [51] and write the  
 709 posterior mean as,

$$\tilde{m}_\phi(\mathbf{x}') = \sum_{i=1}^n \mathcal{A}(\mathbf{x}')^\top \mathcal{A}(\mathbf{x}_i) \alpha_i, \quad \alpha_i \in \mathbb{R}^d \quad (68)$$

$$= \sum_{i=1}^n \mathbf{A} \Psi(\mathbf{x}')^\top \Psi(\mathbf{x}_i) \mathbf{A}^\top \alpha_i \quad (69)$$

$$= \mathbf{A} \left( \sum_{i=1}^n \Psi(\mathbf{x}')^\top \Psi(\mathbf{x}_i) \mathbf{A}^\top \alpha_i \right) \quad (70)$$

$$= \mathbf{A} \tilde{\mathbf{v}}_{\mathbf{x}'} \quad (71)$$

710 after some linear algebra exercises, we can see that  $\alpha_i$  is the  $[i : i + (d - 1)]$  sub-vector of  
 711  $(\mathcal{K}_{\mathbf{X}\mathbf{X}} + \sigma_\phi^2 I)^{-1} \text{vec}(\Phi_{\mathbf{X}})$   $\square$

## 712 C Implementation details and further illustrations.

713 All illustrations are run locally on a MacbookPro 2021 with Apple M1 pro chip.

### 714 C.1 Ablation study on different notions of uncertainties captured

715 To demonstrate the difference between the uncertainties captured by GP-SHAP, BayesSHAP, and  
 716 BayesGP-SHAP, we utilise the California housing dataset [41]. This dataset was derived from the  
 717 1990 U.S. census, each observation represent a census block group. A block group is the smallest  
 718 geographical unit for which the U.S. Census Bureau publishes sample data (a block group typically  
 719 has a population of 600 to 3,000 people). The dataset includes 20640 instances with 8 numerical  
 720 features measuring the following:

- 721 • **MedInc:** Median income in block group
- 722 • **HouseAge:** Median house age in block group
- 723 • **AveRooms:** Average number of rooms per household
- 724 • **AveBedrms:** Average number of bedrooms per household
- 725 • **Population:** Block group population
- 726 • **AveOccup:** Average number of household members
- 727 • **Latitude:** Block group latitude
- 728 • **Longitude:** Block group longitude

729 The target variable is the median house value for California districts, expressed in hundreds of  
 730 thousands of dollars. In the following, we train a GP model and extract explanations using GP-SHAP,  
 731 BayesSHAP, and BayesGP-SHAP, for 4 different configurations:

- 732 1. trained on 25% of data, estimate the Shapley values using 50% of coalitions.
- 733 2. trained on 25% of data, estimate the Shapley values using 100% of coalitions.
- 734 3. trained on 100% of data, estimate the Shapley values using 50% of coalitions.
- 735 4. trained on 100% of data, estimate the Shapley values using 100% of coalitions.

736 To fit the GP model, we employ a sparse Variational GP approach with 200 learnable inducing point  
 737 locations. The evidence lower bound is optimized using batch gradient descent with a batch size of  
 738 64, a learning rate of 0.01, and 100 iterations. The RBF kernel with learnable bandwidths initialized  
 739 using the median heuristic approach is used for the sparse GP. The inducing locations are initialized  
 740 using a standard clustering approach to obtain a representative set of inducing points.

741 After training the model, we reuse the learned inducing points and kernel bandwidths for the  
 742 explanation algorithms. The explanations are obtained using the procedure described in Algorithm 1  
 743 of our work.

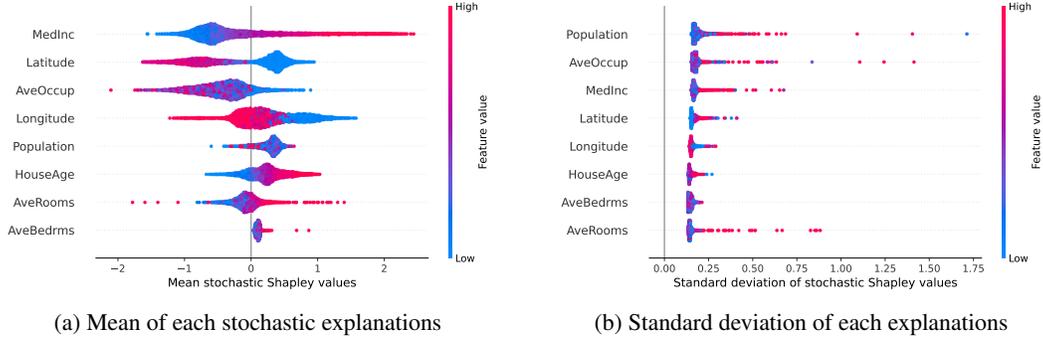


Figure 4: We plot the beeswarm plot of the mean and standard deviations of each stochastic explanations from BayesGP-SHAP fitted on the housing dataset. The features are ranked according to the distance span by the largest and smallest mean (std) stochastic Shapley values.

744 In Figure 1 of our paper, we present the stochastic Shapley values for the 11th observation, computed  
 745 using the three explanation algorithms. The plot includes the 95% credible interval to visualize the  
 746 uncertainties associated with the explanations.

747 **Further illustration:** In Figure 4, we plot the beeswarm plot on the mean and standard deviation  
 748 of each stochastic explanations respectively. We color the point based on the relative size of the  
 749 feature value compared to the rest. We see that in Figure 4a, which plotted the mean stochastic  
 750 shapley values for each observation, the relationship between most features’ explanation to the target  
 751 variable is quite linear. For example, the higher the median income (**MedInc**), the more positive those  
 752 feature contribute to predicting the respective median house value. On the other hand, Figure 4b  
 753 illustrated the standard deviation of each stochastic explanations. In general, we see that the larger  
 754 the feature values are, the more uncertain the explanation becomes. Nonetheless, we see that the  
 755 feature contributing the most, defined as the feature having largest distance spanned by their most  
 756 positive and most negative mean stochastic Shapley values, does not necessarily have the largest  
 757 variation respectively.

## 758 C.2 Exploratory analysis of the stochastic explanations

759 For this illustration, we utilise the breast cancer dataset [42], containing 569 patients with 30 numeric  
 760 features. They are computed from a digitized image of a fine needle aspirate (FNA) of a breast mass  
 761 and describe characteristics of the cell nuclei present in the image:

- 762 • radius (mean of distances from center to points on the perimeter)
- 763 • texture (standard deviation of gray-scale values)
- 764 • perimeter
- 765 • area
- 766 • smoothness (local variation in radius lengths)
- 767 • compactness ( $\frac{\text{perimeter}^2}{\text{area}-1}$ )
- 768 • concavity (severity of concave portions of the contour)
- 769 • concave points (number of concave portions of the contour)
- 770 • symmetry
- 771 • fractal dimension (“coastline approximation” - 1)

772 The goal is to predict whether a tumour is malignant or benign. We first fit a GP model with RBF  
 773 kernel using again the sparse Variational GP formulation with 200 learnable inducing locations. We  
 774 initialise the inducing points using standard clustering techniques on the data. The evidence lower  
 775 bound objective is optimised with a learning rate of  $1e^{-4}$  and 1000 iterations using batch gradient  
 776 descent of batch size 64. To obtain the explanations, we run the BayesGP-SHAP algorithm with  $2^{16}$

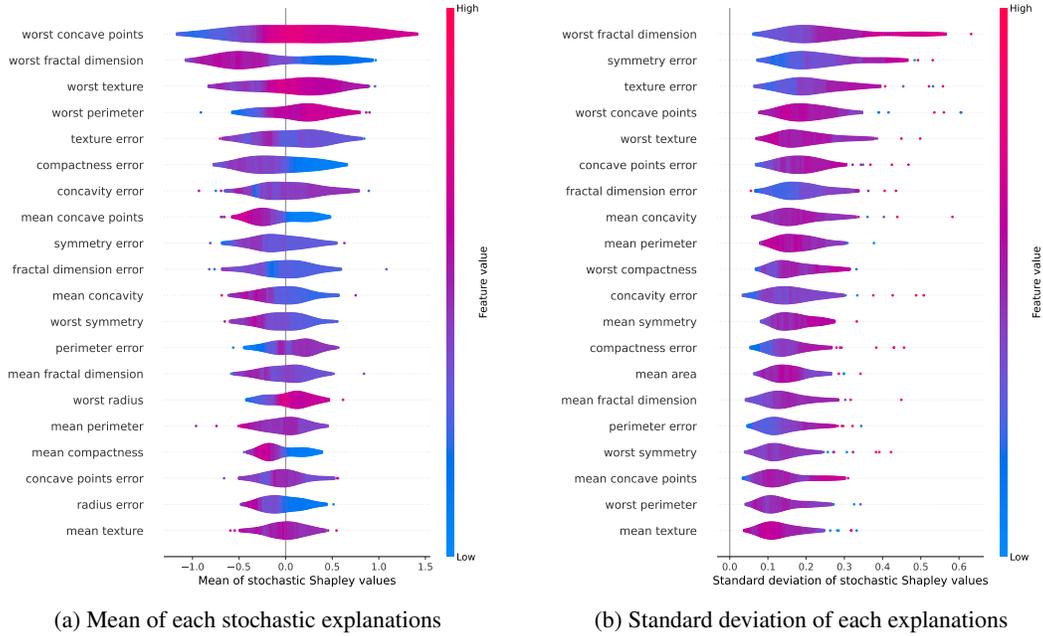


Figure 5: We plot the violin plot of the mean and standard deviations of each stochastic explanations from BayesGP-SHAP fitted on the breast cancer. The features are ranked according to the distance span by the largest and smallest mean (std) stochastic Shapley values.

777 number of coalitions. We do not compare GP-SHAP and BayesSHAP here because the BayesSHAP  
 778 uncertainties have shrunk to almost zero, i.e., the mean standard deviations from the BayesSHAP  
 779 uncertainties across all features and data is 0.0002. This reconfirms the fact from Slack et al. [20]  
 780 that as we increase the sample size the estimation error goes to zero, thus the uncertainties from  
 781 BayesSHAP goes to zero as well. On the other hand, GP-SHAP uncertainties still remain valid  
 782 because it represents the GP predictive uncertainties, which do not shrink to zero as we increase the  
 783 number of coalitions we use to estimate the SVs.

784 **Further illustrations:** In Figure 5, we plot two violin plots to illustrate the relationship between  
 785 mean and standard deviation of the stochastic values with respect to the size of the original feature.  
 786 We see that the feature “worst fractal dimension” are the second most influential feature in terms of  
 787 mean stochastic explanations and also the feature that has highest uncertainty around its explanations.  
 788 In comparison with the housing prediction problem illustrated in Figure 4, the higher the feature  
 789 value doesn’t necessary give higher uncertainty around its explanation.

### 790 C.3 Predictive explanations

791 For this illustration, we utilise the Diabetes dataset [47] with 442 patient data and 10 numeric features  
 792 measuring the following:

- 793 • age: age in years
- 794 • sex
- 795 • bmi: body mass index
- 796 • bp: average blood presuure
- 797 • s1: total serum cholesterol
- 798 • s2: low-density lipoproteins
- 799 • s3: high-density lipoproteins
- 800 • s4: total cholesterol
- 801 • s5: Log of serum triglycerides level

802 • s6: blood sugar level

803 The experiment is to assess the effectiveness of the Shapley prior we proposed in predicting explanations  
804 estimated using SHAP algorithms for general models, including GP-SHAP, TreeSHAP, and  
805 DeepSHAP. We use the implementation of TreeSHAP and DeepSHAP from the `shap` package [2].

806 While algorithms such FastSHAP [22] also learn a vector-valued function that returns explanations  
807 given instances, the algorithm require access to the underlying model  $f$  during training while ours  
808 required previously computed explanations. Due to this importance difference in the problem setup,  
809 we do not compare the two algorithm.

810 We first generate three sets of explanations to set up three regression problems:

- 811 1. Fit a Gaussian process model and then run GP-SHAP to obtain explanations.
- 812 2. Fit a random forest model and then run TreeSHAP to obtain explanations.
- 813 3. Fit a neural network model and then run DeepSHAP to obtain explanations.

814 After obtaining explanations as groundtruths for this experiment, we randomly divide 70% of them  
815 as training data and 30% of them as testing data. We then do the following,

- 816 1. We fit a multi-output GP using the proposed Shapley prior on the training data and predict  
817 the explanations of the unseen test data.
- 818 2. We fit a multi-output random forest model on the training data and predict the explanations  
819 of the unseen test data.
- 820 3. We fit a multi-output neural network model on the training data and predict the explanations  
821 of the unseen test data.

822 We repeat this experiment 10 times using different seeds and compute the RMSE between the  
823 predicted and groundtruths explanations. The results are then plotted in Figure 3.

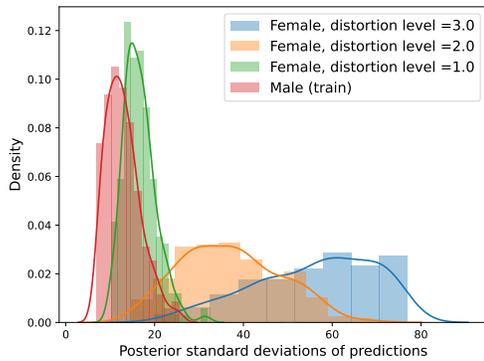
#### 824 **C.4 Further ablation study: Impact of increased posterior prediction uncertainty on** 825 **explanation uncertainties**

826 In this ablation study, we aim to examine the effect of increasing the uncertainty in posterior  
827 predictions on the corresponding uncertainty in stochastic Shapley values. To demonstrate this, we  
828 utilize the diabetic dataset [47] and split the data based on recorded sex. We train our GP model on  
829 the male data and employ BayesGP-SHAP to explain the prediction results for both the male training  
830 data and the female testing data. We adopt this split because we expect the biological characteristics  
831 between males and females to be distinct enough to treat the female data as out-of-sample data,  
832 thereby naturally resulting in increased predictive uncertainty for the female data. To further amplify  
833 this uncertainty, we multiply each instance in the female testing data by distortion factors of two and  
834 three, respectively, and assess the corresponding uncertainties in the explanations.

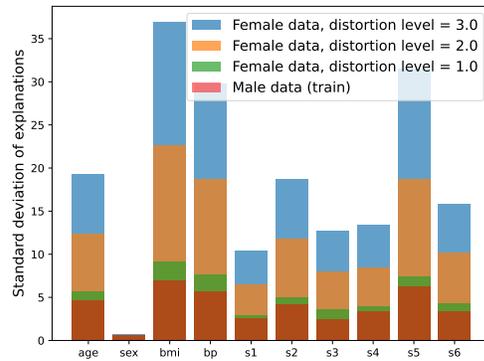
835 We begin by illustrating the relationship between the out-of-sampleness of the data and the corre-  
836 sponding increase in predictive posterior uncertainties. This is depicted in Figure 6a, where we  
837 observe that as the data becomes more out-of-sample, the predictive uncertainties consistently rise.  
838 Even at distortion level 1, which represents the original female data, we can already observe increased  
839 uncertainties compared to the uncertainties derived from male data prediction.

840 Furthermore, these increased uncertainties in the predictive posterior are reflected in the associated  
841 feature explanations. This is evident in Figure 6b, where we visualize the uncertainties associated  
842 with the feature explanations. For instance, the green bars representing the average uncertainties in  
843 explaining female data with no distortion are consistently larger than the red bars, which represent the  
844 average uncertainties of male data explanations. This observation aligns with the higher predictive  
845 uncertainties observed in Figure 6a for the female data compared to the male training data.

846 It is worth noting that the uncertainty for the feature “sex” remains consistently close to zero. This is  
847 because the feature “sex” is constant within both the female and male datasets. As a result, it acts as  
848 the null player in each dataset and obtains an almost Dirac zero as its stochastic Shapley value.



(a) Predictive posterior standard deviation



(b) Mean of standard deviations of explanations

Figure 6: Ablation study: (left) We begin by training a Gaussian Process (GP) model on the male data. We then make predictions using this trained model on both the male data and out-of-sample female data. To assess the impact of increasing posterior uncertainties, we multiply the female data by distortion levels of 1.0, 2.0, and 3.0. We visualize the results by plotting the density plot of the standard deviations obtained from the predictive posterior distributions. (right) Next, we focus on analyzing the average standard deviations of explanations per feature from the male and female data, considering different distortion levels. We observe that as we progressively increase the posterior uncertainties in the sample, these uncertainties are reflected in the uncertainties of the explanations provided.