Gaussian Processes for choice data

Dario Azzimonti[1] Alessio Benavoli[2] Dario Piga[1]

[1]Dalle Molle Institute for Artificial Intelligence (IDSIA), Lugano, Switzerland.
[2]School of Computer Science and Statistics, Trinity College, Dublin, Ireland

LIKE23 -Lifting Inference with Kernel Embeddings
28th June 2023 - Bern, Switzerland

Outline

How to model a user's behavior?

Many real-world optimization problems do not provide explicit objectives.

Examples:

- Choosing a cake for an event;
- Buying a new laptop;
- Refine the industrial production machined pieces.

Modelling choices of a user/consumer can be seen as a real world optimisation problem involving multiple conflicting objectives

How to model a user's behavior?

Example: choosing the best cake

Possible objectives:

- taste;
- softness;
- aestetics.

Input features (ingredients):

- flour;
- butter;
- sugar;
- ...

The user only sees the final cakes.

$A = \{$  ,  ,  ,  ,  $\}$

How to select the best item?

A: Ask the user for **ratings** (e.g. 1-10)

B: Ask the user for **preferences**

C: Ask the user for **choice sets**

Given

$A = \{$  $\}$

The user chooses

$C(A) = \{$  $\}$

How to select the best item?

Option	**Pros**	**Cons**
A ratings	* easy modelling	* often inconsistent * noisy
B preferences	* easier for the user	* can be inconsistent * many iterations
C choices	* easiest for the user in case of inconsistencies	* no model available

Choice function

Consider

- $\mathcal{X}$, a finite set of items;

$\mathcal{X} = \{$  $\}$

- $\mathcal{Q}$, the set of all (finite) subsets of $\mathcal{X}$.

$\mathcal{Q} = \{ \{ \ldots \} , \ldots, \{ \ldots \}, \ldots, \{ \ldots \} \}$

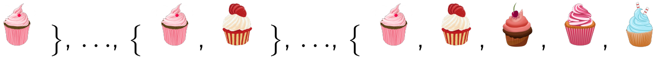

A choice function is a map $C : \mathcal{Q} \to \mathcal{Q}$ such that

$$C : A \in \mathcal{Q} \;\mapsto\; C(A) \in \mathcal{Q}$$

$C : A = \{$  $\} \in \mathcal{Q} \;\mapsto\; C(A) = \{$  $\} \in \mathcal{Q}$

Learning user's behavior $\Rightarrow$ learning choice function from history.

Choice function - a few considerations

We model each item in A with a vector $\mathbf{x} \in \mathbb{R}^{n_x}$ containing its features.

Set of rejected items: $R(A) = A \setminus C(A)$, for any $A \in \mathcal{Q}$.

- If $\mathbf{x}_j \in R(A)$, there is at least one object in $C(A)$ better than $\mathbf{x}_j$

Incomparability: if $\{\mathbf{x}_j, \mathbf{x}_k\} \subseteq C(A)$ then $\mathbf{x}_j$ and $\mathbf{x}_k$ are incomparable

- The user may have multiple utilities
- Lack of knowledge

How to model choice functions?

Vector of utility functions $\mathbf{u} = [u_1(\mathbf{x}), \ldots, u_{n_o}(\mathbf{x})]^T$.

Pareto-dominant option $\mathbf{x_1}$ Pareto-dominates $\mathbf{x_2}$ ($\mathbf{x_1} \succ \mathbf{x_2}$) if

i) for all $j = 1, \ldots, n_o$, $u_j(\mathbf{x_1}) \geq u_j(\mathbf{x_2})$

ii) $\exists j \in \{1, \ldots, n_o\}$ s.t. $u_j(\mathbf{x_1}) > u_j(\mathbf{x_2})$

Non-dominated Pareto set Given $A = \{\mathbf{x_1}, \ldots, \mathbf{x_n}\}$, the set of non-dominated options is

$$A' = \{\mathbf{x} \in A : \nexists \mathbf{x}' \in A \text{ s.t. } \mathbf{x}' \succ \mathbf{x}\}$$

Pareto-rationalisable choice functions

u describes the choice function C if, for each $A \subset \mathcal{X}$,

- $C(A)$ is the *non-dominated set* in the *strong Pareto sense* for **u**;
- $R(A)$ is the set of dominated objects.

Note: not all choice functions are Pareto-rationalisable.

Pareto-rationalisable choice functions: example

Example: choose best cake

- utilities: taste, aspect
-  dominates 
- set of non-dominated options

$A' = \{$  ,  ,  $\}$

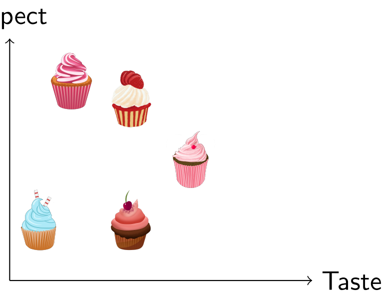

We do not observe the vector of utility functions

An exact link between choices and utilities

Assume: latent vector of utilities $\mathbf{u} = [u_1(\mathbf{x}), \ldots, u_{n_o}(\mathbf{x})]^T$.

For each $A \subset \mathcal{X}$, we can link choices and utilities with

$$\neg \left(\min_{i \in \{1,\ldots,n_o\}} (u_i(\mathbf{o}) - u_i(\mathbf{v})) < 0, \ \forall \mathbf{o} \in C(A) \right), \forall \mathbf{v} \in R(A), \qquad (1)$$

(For each $\mathbf{v} \in R(A)$, there is at least a object in $C(A)$ not worse than $\mathbf{v}$)

$$\min_{i \in \{1,\ldots,n_o\}} (u_i(\mathbf{o}) - u_i(\mathbf{v})) < 0, \ \forall \mathbf{o}, \mathbf{v} \in C(A), \ \mathbf{o} \neq \mathbf{v}. \qquad (2)$$

(For each object in $C(A)$, there is no better object in $C(A)$)

Using the link to build a likelihood

Given a choice dataset

$$\mathcal{D}_m = \{(C(A_s), A_s): \text{ for } s = 1, \dots, m\}, \quad A_s \subset \mathcal{X} \text{ for each } s,$$

$X = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t]^\top$ features associated with t objects in $\mathcal{X}$

We define a likelihood

$$p_{exact}(\mathcal{D}_m|\mathbf{u}(X)) = \prod_{k=1}^{m} p_{exact}(C(A_k), A_k|\mathbf{u}(X)),$$

where $\mathbf{u}(X) = [\mathbf{u}(\mathbf{x}_1), \mathbf{u}(\mathbf{x}_2), \dots, \mathbf{u}(\mathbf{x}_t)]^\top$ and

$$p_{exact}(C(A_k), A_k|\mathbf{u}(X)) = \begin{cases} 1 & \text{if both conditions are satisfied} \\ 0 & \text{otherwise} \end{cases}$$

Case of non-Pareto rational choices

Consider $\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3, \mathbf{o}_4$ with

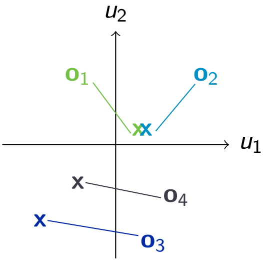

Assume we are given the following choices:

$$C(\{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3\}) = \{\mathbf{o}_1, \mathbf{o}_2\}, \quad C(\{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_4\}) = \{\mathbf{o}_1\},$$

These choices are not Pareto rational, $p_{exact}(\mathcal{D}_m|\mathbf{u}(X))$ is zero.

Likelihood accounting for errors

On the choice dataset $\mathcal{D}_m$ we define the likelihood

$$p(\mathcal{D}_m|\mathbf{u}(X)) = \prod_{k=1}^{m} p(C(A_k), A_k|\mathbf{u}(X))$$

$$= \prod_{k=1}^{m} \prod_{\{\mathbf{o},\mathbf{v}\}\in C_\sharp(A_k)} \left(1 - \prod_{i=1}^{n_o} \Phi\left(\frac{u_i(\mathbf{o}) - u_i(\mathbf{v})}{\sigma}\right) - \prod_{i=1}^{n_o} \Phi\left(\frac{u_i(\mathbf{v}) - u_i(\mathbf{o})}{\sigma}\right)\right)$$

$$\prod_{\mathbf{v}\in R(A_k)} \left(1 - \prod_{\mathbf{o}\in C(A_k)} \left(1 - \prod_{i=1}^{n_o} \Phi\left(\frac{u_i(\mathbf{o}) - u_i(\mathbf{v})}{\sigma}\right)\right)\right)$$

The blue part is a probabilistic relaxation of the first condition
(For each $\mathbf{v} \in R(A)$, there is at least a object in $C(A)$ not worse than $\mathbf{v}$)

The green part is a probabilistic relaxation of the second condition
(For each object in $C(A)$, there is no better object in $C(A)$)

Non-Pareto rational choices

Consider the same objects $\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3, \mathbf{o}_4$

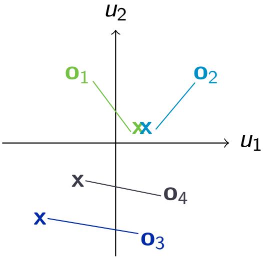

With the following choices:
$C(\{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3\}) = \{\mathbf{o}_1, \mathbf{o}_2\}$,
$C(\{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_4\}) = \{\mathbf{o}_1\}$

Probabilistic relaxation likelihood

$p(\{\mathbf{o}_1, \mathbf{o}_2\}, \{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3\} | \mathbf{u}(X)) \approx 0.48$

$p(\{\mathbf{o}_1\}, \{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_4\} | \mathbf{u}(X)) \approx 0.12$

$p(\mathcal{D}_m | \mathbf{u}(X)) \approx 0.48 \cdot 0.12 = 0.057$

ChoiceGP

Data: $\mathcal{D}_m = \{(C(A_s), A_s) : \text{ for } s = 1, \ldots, m\}, \quad A_s \subset \mathcal{X}$ for each s;

Prior: each latent utility function in $\mathbf{u}(\mathbf{x}) = [u_1(\mathbf{x}), \ldots, u_{n_o}(\mathbf{x})]^\top$ is modelled as an independent GP:

$$u_i(\mathbf{x}) \sim \mathrm{GP}_i(0, k_i(\mathbf{x}, \mathbf{x}')), \qquad i = 1, 2, \ldots, n_o.$$

Likelihood: Error accounting likelihood

$$p(\mathcal{D}_m | \mathbf{u}(X)) = \prod_{k=1}^{m} p(C(A_k), A_k | \mathbf{u}(X))$$

ChoiceGP - Inference

Posterior: $p(\mathbf{u}(X)|\mathcal{D}_m) = \frac{p(\mathbf{u}(X))}{p(\mathcal{D}_m)} \prod_{k=1}^{m} p(C(A_k), A_k|\mathbf{u}(X))$

Model parameters:

- ARD lengthscales of each $k_i(\cdot,\cdot)$;
- scale parameter σ in the likelihood.

The posterior is not a GP, its computation requires an approximation.

ChoiceGP - Inference

Approximate $p(\mathbf{u}(X)|\mathcal{D}_m)$ with **variational inference**.

Variational density: $q(\mathbf{u}(X)) \sim N(\boldsymbol{\mu}, S)$, with S block-diagonal.

Find the q and the model parameters (lengthscales and σ) by maximizing

$$ELBO(q) = \underbrace{\int q(\mathbf{u}(X)) \log p(\mathcal{D}_m|\mathbf{u}(X)) d\mathbf{u}}_{\text{likelihood term}} - \underbrace{KL[q(\mathbf{u}(X))||p(\mathbf{u}(X))]}_{\text{KL between priors}}$$

We approximate the likelihood term with Monte Carlo integration.

ChoiceGP - Prediction

For a new vector of p objects $X^* = \{\mathbf{x}_1^*, \dots, \mathbf{x}_p^*\}$

$$p(\mathbf{u}(X^*)|\mathcal{D}_m) = \int p(\mathbf{u}(X^*)|\mathbf{u}(X))q(\mathbf{u}(X)|\mathcal{D}_m)d\mathbf{u}(X)$$

where $q(\mathbf{u}|\mathcal{D}_m)$ is the approximate VI posterior.

For a new set A^* and a possible choice $C(A^*)$

$$p(C(A^*), A^*|\mathcal{D}_m) = \int p(C(A^*), A^*|\mathbf{u}(X^*))p(\mathbf{u}(X^*)|\mathcal{D}_m)d\mathbf{u}(X^*),$$

computed via Monte Carlo sampling from $p(\mathbf{u}(X^*)|\mathcal{D}_m)$.

A. Benavoli, D. Azzimonti and D. Piga. 2023. Preference Learning with Gaussian Processes. Accepted *UAI 2023* (Pittsburgh, USA)

A simple example

Let us create choice data from a known utility

$\mathbf{u}(x) = [\cos(2x), -\sin(2x)]$, $x \in \mathbb{R}$.

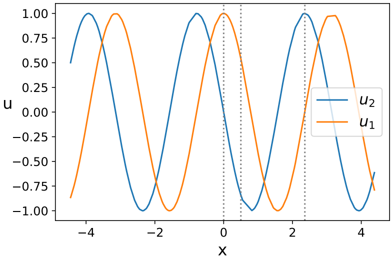

Given the set $A_k = \{0, 0.5, 2.36\}$,

$\mathbf{u}(0) = [1, 0]$
$\mathbf{u}(0.5) = [0.54, -0.84]$
$\mathbf{u}(2.36) = [0, 1]$ and

$C(A_k) = \{0, 2.36\}$

We sample 200 inputs x_i uniformly in $[-4.5, 4.5]$ and evaluate $\mathbf{u}$.

A simple example

From $\{x_i, \mathbf{u}(x_i)\}_{i=1}^{200}$ we generate $m = 150$ choices from sets A_k

A_k has fixed size $|A_k| = 3$

choiceGP trained on
$m = 150$ choices

RBF with ARD for each GP

Predict $\mathbf{u}(X^*)$ at X^*

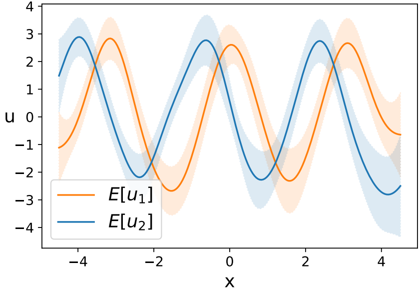

Note:

- we cannot estimate the range of $\mathbf{u}$ from choices;
- only Pareto-dominance is learned.

A simple example

From $\left\{x_i, \mathbf{u}(x_i)\right\}_{i=1}^{200}$ we generate $m = 150$ choices from sets A_k

Increase set size: $|A_k| = 5$

choiceGP trained on
$m = 150$ choices

RBF with ARD for each GP

Predict $\mathbf{u}(X^*)$ at X^*

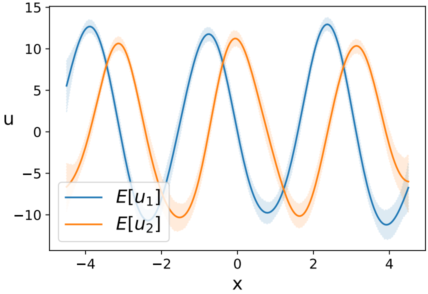

Note:

- by allowing choices from larger sets we gather more information;
- smaller credible intervals than with $|A_k| = 3$.

Latent Dimension Estimation

Until now: n_o was given, but if we observe only choices n_o is not known.

Idea: fit ChoiceGP$_{n_o}$ with increasing n_o.

Compare performances with *Pareto Smoothed Importance Sampling Leave-One-Out* cross-validation (PSIS-LOO, Vehtari et. al., 2017)

Example:

Find the true latent dimension ($n_o = 2$) on the previous example.

PSIS-LOO values

d	$\lvert A_k \rvert = 3$	$\lvert A_k \rvert = 5$
1	-3213	-6108
2	**-69**	**-84**
3	-80	-95
4	-91	-109

Comparison with state-of-the-art

ChoiceNN: a neural network based method for choice learning (Pfannschmidt and Hüllermeier, 2020)

Experiment:

- utilities DTLZ, ZTD multi-objective opt. benchmarks
- generate choices from $|A| = 10$ from utilities;
- object $\mathbf{x}_i \in \mathbb{R}^6$
- number of objectives given, $n_o = 5$

K. Pfannschmidt and E. Hüllermeier. 2020 Learning Choice Functions via Pareto-Embeddings. In *German Conference on Artificial Intelligence. Springer.*

Comparison with state-of-the-art

For each test function

- $90 - 10$ train/test split;
- 5 MC repetitions;
- metric: A-mean

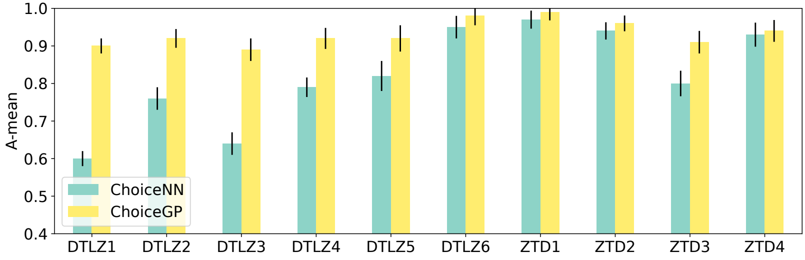

Comparison with preference learning

We compare choiceGP against

- *Preferential GP* (PGP) (Chu and Ghahramani, 2005),
- *General Preferential GP* (GPGP) (Chau et al., 2022)
- *GP with data augmentation* (PairGP) (Chau et al., 2022)

Note:

- preference learning only allows for binary comparisons;
- in case of inconsistencies we need a forced choice;

Comparison with preference learning

Create choices from five benchmark multi-obj opt. datasets (AM, EDM, Jura, Slump, Vehicle)

$|A_k| = 2$, use benchmark outputs to make $\{(C(A_k), A_k) : \ k = 1, \ldots, m\}$

For each $A_k = \{\mathbf{x}_i, \mathbf{x}_j\}$ we can have

- $C(A_k) = \{\mathbf{x}_i\}$ then $\mathbf{x}_i \succ \mathbf{x}_j$
- $C(A_k) = \{\mathbf{x}_i, \mathbf{x}_j\}$, we generate preferences with

  Random: coin flip: $\mathbf{x}_i \succ \mathbf{x}_j$ if Heads; $\mathbf{x}_j \succ \mathbf{x}_i$ if Tails
  Maj. rule: $\mathbf{x}_i \succ \mathbf{x}_j$ if $\mathbf{x}_i$ is better than $\mathbf{x}_j$ w.r.t. majority of outputs.

Comparison with preference learning

Average accuracy

	ChoiceGP		PGP	GPGP	PairGP
AM	0.90	maj.	0.84	0.86	0.87
		rand.	0.74	0.73	0.738
EDM	0.88	maj.	0.83	0.80	0.83
		rand.	0.84	0.82	0.82
Jura	0.91	maj.	0.87	0.87	0.87
		rand.	0.84	0.82	0.82
Slump	0.91	maj.	0.93	0.90	0.90
		rand.	0.83	0.79	0.79
Vehicle	0.93	maj.	0.89	0.90	0.90
		rand.	0.80	0.80	0.80

Bayesian Optimization on Choice data

BO: find the global maximum of an unknown expensive function.

For scalar real-valued function g the objective is $\mathbf{x}^o = \arg\max_{\mathbf{x}\in\Omega} g(\mathbf{x})$

BO makes this a sequential decision problem:

0) collect N couples $\mathcal{D} = \{(\mathbf{x}_i, g(\mathbf{x}_i)) : \quad i = 1, 2, \ldots, N\}$;
1) create a surrogate for g from data;
2) employ *acquisition function* to select next candidate $\mathbf{x}_{N+1}$
3) evaluate g, add $(\mathbf{x}_{N+1}, g(\mathbf{x}_{N+1}))$ to $\mathcal{D}$, update GP;

Peter I. Frazier. 2018. A Tutorial on Bayesian Optimization. arXiv:1807.02811

BO on Choice data

On choice data we do not observe the function.

The user is asked to choose preferred objects in a set A_k.

0) collect N choices $\mathcal{D}_N = \{(C(A_s), A_s) : \text{ for } s = 1, \ldots, N\}$;
1) learn choiceGP from data;
2) employ *acquisition function* to select next object $\mathbf{x}_{N+1}$ to compare;
3) query the agent for new choice(s), update choiceGP;

Two acquisition functions

We consider two acquisition functions

choiceUCB: $\gamma\%$ Upper Credible Bound (UCB) of $p(C(A^*), A^*|\mathbf{u}^*)$ with $\mathbf{u}^* \sim p(\mathbf{u}^*|\mathcal{D}_m)$, $A^* = \mathbf{x} \cup \hat{\mathcal{X}}^{nd}$ and $C(A^*) = \{\mathbf{x}\}$

Note: $C(A^*) = \{\mathbf{x}\}$ is a strong requirement.

choiceThompson: Pareto front hyper-volume increase for a posterior realization when $\mathbf{x}$ is added;

Note: we use a predictive process for the posterior realization. More

A. Benavoli, D. Azzimonti and D. Piga. 2023. Bayesian Optimization For Choice Data. In GECCO '23. Association for Computing Machinery.

Given $\mathbf{x}_{new}$, how to query for new choices?

Consider $\hat{\mathcal{X}}^{nd}$, current Pareto set implied by choices.

We would like to query a choice for $A^* = \{\mathbf{x}_{new}\} \cup \hat{\mathcal{X}}^{nd}$.

Note:

- if $\hat{\mathcal{X}}^{nd}$ is large (> 5) the user might have issues choosing;
- often the expensive part of the pipeline is producing the object, not the comparison;

Given $\mathbf{x}_{new}$, we produce the item and we query the user for choices on

- sets A_j^* of fixed size $|A_j^*|$;
- each set contains 5 combinations of $\mathbf{x}_{new}$ with $|A_j^*| - 1$ items of $\hat{\mathcal{X}}^{nd}$.

Experiments

We consider standard multi-objective benchmark functions:
DTLZ1, ZDT1, Kursawe

choiceUCB, choiceThompson versus Oracle with access to true utilities.

GP model: ARD Matérn, $\nu = 3/2$;

Latent dimension: n_o assumed known;

Initialization:

- 20 randomly selected inputs to initialise Oracle-qEHVI.
- 10 pairs $\{C(A_k), A_k\}$ of size $|A_k| = 3$ for ChoiceBO

Total budget: 100 iterations

Experiments

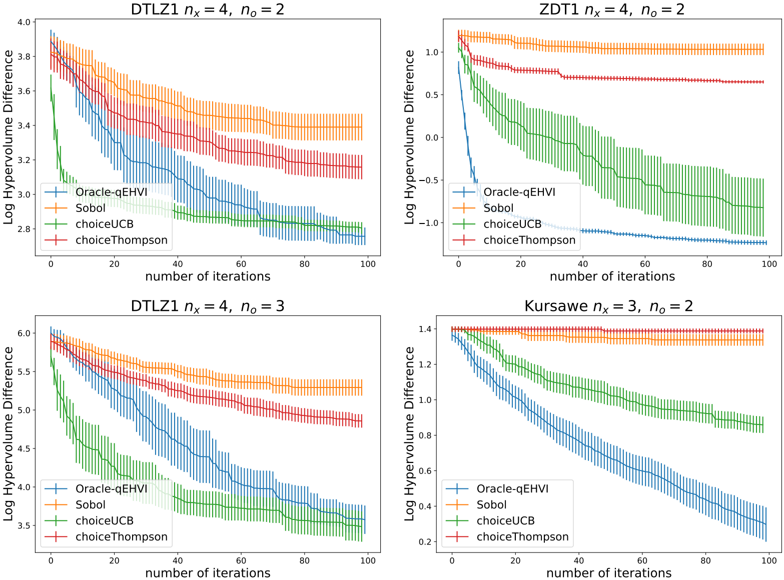

Conclusions

- Generalization of preference learning to choices
- Useful when it is hard to express coherent preferences
- ChoiceGP models user choices without access to latent utilities
- ChoiceBO provides a framework to give the user better choices

Future work:

- Study the case of non Pareto-rationalisable choices;
- Sparse implementation to reduce computational load.

References

H. Moulin. 1985. Choice Functions Over a Finite Set: A Summary. *Social Choice and Welfare* 2(2): 147-60.

A. Vehtari, A. Gelman, and J. Gabry. 2017. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and computing*, 27(5):1413-1432.

A. Benavoli, D. Azzimonti and D. Piga. 2023. Preference Learning with Gaussian Processes. Accepted *UAI 2023* (Pittsburgh, USA)

K. Pfannschmidt and E. Hüllermeier. 2020. Learning Choice Functions via Pareto-Embeddings. In *German Conference on Artificial Intelligence. Springer.*

W. Chu and Z. Ghahramani. 2005. Preference Learning with Gaussian Processes. In *Proceedings of the 22nd International Conference on Machine Learning* (Bonn, Germany) (ICML '05) 137-144.

S. L. Chau, J. Gonzalez, and D. Sejdinovic. 2022. Learning Inconsistent Preferences with Gaussian Processes. In *25th International Conference on Artificial Intelligence and Statistics*, 2266-2281. PMLR.

References (2)

A. Benavoli, D. Azzimonti and D. Piga. 2023. Bayesian Optimization For Choice Data. *In GECCO '23. Association for Computing Machinery.*

S. Daulton, M. Balandat, and E. Bakshy. 2020. Differentiable Expected Hypervolume Improvement for Parallel Multi-Objective Bayesian Optimization. In *Advances in Neural Information Processing Systems* Vol. 33. Curran Associates, Inc., 9851-9864.

William R. Thompson. 1933. On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika 25, 3/4, 285.*

Alessio Benavoli, D. A., Dario Piga. 2021. A unified framework for closed-form nonparametric regression, classification, preference and mixed problems with Skew Gaussian Processes. *Machine Learning* 110, 3095–3133.

Javier González, Zhenwen Dai, Andreas Damianou, and Neil D Lawrence. 2017. Preferential Bayesian optimization. In *Proceedings of the 34th International Conference on Machine Learning*. Volume 70. JMLR. 1282–1291

H. Nickisch and C. E. Rasmussen. 2008. Approximations for Binary Gaussian Process Classification. *Journal of Machine Learning Research* 9 2035-2078