Nonparametric Predictive Utility Inference

Brett Houlding[1] and Frank Coolen[2]

[1] *Dept. Statistics, Trinity College Dublin, Ireland*

[2] *Dept. Mathematical Sciences, Durham University, UK*

Outline:

Which to choose?

Known fruits:

Newly discovered fruits:

(Dragon Fruit)

(Mangosteen)

Which to choose?

Known fruits:

- Five previously experienced fruits $f_1, \ldots, f_5$ which, on a $[0,1]$ scale, have ordered utility values $u_{(1)}, \ldots, u_{(5)}$ equal to 0.3, 0.35, 0.4, 0.5 and 0.7:

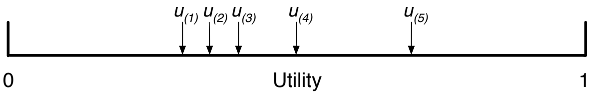

Newly discovered fruits:

- Two alternative and unexperienced fruits f_{new} and f_{new_2}.

What to select in a one off choice? What about a sequential choice?

Bayesian Decision Theory

- In Bayesian statistics, beliefs over an unknown random quantity are typically assigned a parametric model. Learning then occurs following observation of data that has probabilistic dependence with the unknown random quantity.
- The theory is well established (though not undisputed):

$$\text{Posterior} \propto \text{Likelihood} \times \text{Prior}$$

- If the aim of the analysis is to perform statistical inference, then the posterior distribution (or posterior predictive distribution) is all that is of interest.
- If, however, the aim is to aid ('optimal'?) decision making, then the preferences of the decision maker should be taken into account.
- Preferences are modelled via a utility function, which is typically assumed to be fully known, *i.e.*, preferences are known precisely.
- The 'optimal' decision is then (the) one that gives highest expected utility with respect to beliefs over the random quantity involved.

Traditional Utility Theory

- Preference over decisions reconstructed from assumed known utility of decision outcomes and the probability of achieving that outcome.
- Usual to assume a fixed utility form, and/or specific utility values for the available outcomes:
  - $u(\$x) = \log(x + c)$
  - $u(apple) = 0.9$, $u(banana) = 0.5$
- Does not permit inherent uncertainty in preferences over decisions.
- Does not allow the learning of utility and assumes the decision maker will never be surprised by the utility of an outcome.

Adaptive Utility

- In reality people often learn (about) preferences, *e.g.*, by experimenting.
- This requires a generalization of the traditional concept of utility.
- Adaptive Utility, as first suggested by Cyert & DeGroot [3], is one such possibility.
- Basic idea rather simple: Treat utility in the same way that unknown random quantities are typically treated in standard Bayesian statistical inference, *i.e.*, subject them to a parametric belief model, for example:

$$u(saving, speed|\theta) = (1-\theta) \times saving + \theta \times speed$$

Adaptive Utility

This idea was further developed by Houlding [6]:

- Construction of adaptive utility from commensurable options.
- Application in sequential problems, *e.g.*, reliability.
- How is value of sample information affected by uncertainty in preferences.
- Adaptive Utility leads to a concept of trial aversion.

Yet despite the above, there are remaining issues:

- How to determine prior beliefs over an uncertain utility value?
- How to determine a likelihood linking the uncertain utility value with utility data?
- What is utility data?

Nonparametric Predictive Inference

Based on Hill's $A_{(n)}$ assumption [4]:

Let real-valued $x_{(1)} < \ldots < x_{(n)}$ be the ordered values of data $x_1, \ldots, x_n$, and let X_i be the corresponding pre-data random quantities, then:

1. The observable random quantities $X_1, \ldots, X_n$ are exchangeable.
2. Ties have probability 0, so $x_i \neq x_j$ for all $i \neq j$, almost surely.
3. Given data $x_1, \ldots, x_n$ and the definition that $x_{(0)} = -\infty$, $x_{(n+1)} = \infty$, $I_j = (x_{(j-1)}, x_{(j)})$, then for $j = 1, \ldots, n+1$:

$$P(X_{n+1} \in I_j) = \frac{1}{n+1}$$

This generalises to the following predictive probability bounds:

$$\underline{P}(X_{n+1} \in B) = \frac{|\{j : I_j \subseteq B\}|}{n+1}, \overline{P}(X_{n+1} \in B) = \frac{|\{j : I_j \cap B \neq \phi\}|}{n+1}$$

Nonparametric Predictive Inference

- NPI is a low structure statistical technique that is predictive by nature.
- Less restrictive belief model that is closer to resembling a state of ignorance.
- Less presumptious alternative for making inference than the direct specification of conditional independencies and specific distributional forms.
- May be relevant when there is a lack of additional information further to the data itself.
- Coincides with the general framework of a finitely additive prior (Hill [5]) and has been related to the theory of imprecise probability (Augustin & Coolen [1]).
- Subjectivist interpretation of lower and upper bounds on betting price.

NPUI

- The NPI statistical technique offers a simple, yet possibly appealing, solution to the problem of identifying an appropriate utility learning model.
- Particularly useful when decision outcomes form a finite set (with assumed exchangeability over their utility values) which includes the option of novel outcomes, *e.g.*, a new brand becomes available in a consumer selection problem.
- Additional possibilities for comparing decisions over multiple sets of outcomes with exchangeability only assumed within each set (Coolen [2]), though not considered here.

NPUI

- Let $u_{(1)}, \ldots, u_{(n)}$, with $u_{(i)} \in (0,1)$ be the known ordered values of the utilities $u_1, \ldots, u_n$ representing preferences over outcomes $\mathcal{O}_n = \{o_1, \ldots, o_n\}$.
- Let $\mathcal{U}_n = \{U_1, \ldots, U_n\}$ denote the set of random quantities representing the utilities of the elements within $\mathcal{O}_n$ before they are experienced, and suppose that the elements of $\mathcal{U}_n$ are considered exchangeable.
- Given a new and novel outcome o_{new} whose utility value $U_{new} \in (0,1)$ is unknown but considered exchangeable with the elements of $\mathcal{U}_n$, the NPUI model considered here states only the following:

$$P\Big(U_{new} \in (0, u_{(1)}]\Big) = P\Big(U_{new} \in [u_{(i)}, u_{(i+1)}]\Big) = P\Big(U_{new} \in [u_{(n)}, 1)\Big) = \frac{1}{n+1}$$

Expected Utility Bounds

NPUI leads to the following rules:

- Lower expected utility bound:

$$\underline{E}[U_{new}] = \frac{1}{n+1} \sum_{i=1}^{n} u_i$$

- Upper expected utility bound:

$$\overline{E}[U_{new}] = \frac{1}{n+1} \left(1 + \sum_{i=1}^{n} u_i\right)$$

- Difference in utility bounds:

$$\Delta\Big(E[U_{new}]\Big) = \overline{E}[U_{new}] - \underline{E}[U_{new}] = \frac{1}{n+1}$$

Updating

Expected utility bounds of a second novel outcome o_{new_2} once u_{new} is known:

- Lower updated expected utility bound:

$$\underline{E}[U_{new_2}|u_{new}] = \frac{n+1}{n+2}\underline{E}[U_{new}] + \frac{1}{n+2}u_{new}$$

- Upper updated expected utility bound:

$$\overline{E}[U_{new_2}|u_{new}] = \frac{n+1}{n+2}\overline{E}[U_{new}] + \frac{1}{n+2}u_{new}$$

- Difference in updated utility bounds:

$$\Delta\Big(E[U_{new_2}|u_{new}]\Big) = \frac{1}{n+2}$$

Decision Tree

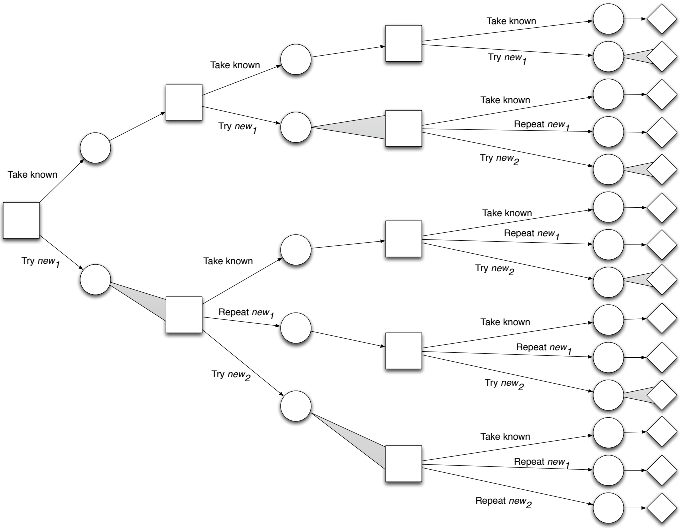

Reduced Decision Tree

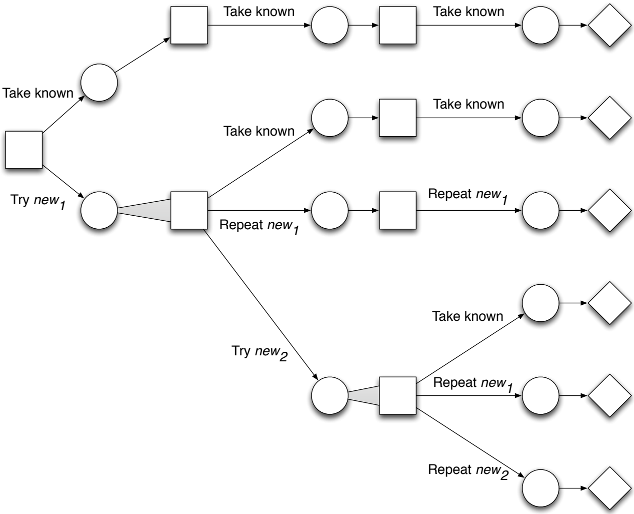

Sequential Choice Rules

In a sequential problem, a rule must be devised for choosing future decisions.

Extreme Pessimism:

The DM will always select the outcome or sequential decision path whose lower expected utility bound is greatest. Furthermore, future uncertain utility realisations will always fall at the infimum of any considered interval formed by the ordering of known utility values.

Extreme Optimism:

The DM will always select the outcome or sequential decision path whose upper expected utility bound is greatest. Furthermore, future uncertain utility realisations will always fall at the supremum of any considered interval formed by the ordering of known utility values.

Conditioning

Expected utility bounds of a second novel outcome o_{new_2} given that only the interval of u_{new} is known:

- Lower conditional expected utility bound:

$$\underline{E}[U_{new_2}|U_{new} \in I_j] = \frac{1}{n+2}\left(\sum_{i=1}^{n} u_i + \inf(I_j)\right)$$

- Upper conditional expected utility bound:

$$\overline{E}[U_{new_2}|U_{new} \in I_j] = \frac{1}{n+2}\left(1 + \sum_{i=1}^{n} u_i + \sup(I_j)\right)$$

- Difference in updated utility bounds:

$$\Delta\left(E[U_{new_2}|U_{new} \in I_j]\right) = \frac{1 + \sup(I_j) - \inf(I_j)}{n+2}$$

- Internal Consistency:

$$\underline{E}[U_{new_2}] = \sum_{j=1}^{n+1} \underline{E}[U_{new_2}|U_{new} \in I_j] P(U_{new} \in I_j)$$

Summary Results Table

Expected Utility for Optimal Decision Strategy

Available	Pessimistic		Optimistic		Select a Novel Option	
	Lower Bound	Upper Bound	Lower Bound	Upper Bound	Pessimistic	Optimistic
f_1	1.298	1.817	1.298	1.817	Yes	Yes
f_2	1.305	1.819	1.305	1.819	Yes	Yes
f_3	1.323	1.785	1.319	1.826	Yes	Yes
f_4	1.500	1.500	1.367	1.855	No	Yes
f_5	2.100	2.100	2.100	2.100	No	No

i	1	2	3	4	5
$u_{(i)}$	0.3	0.35	0.4	0.5	0.7

For the one-period problem:

$$\underline{E}[U_{new}] = 0.375$$

$$\overline{E}[U_{new}] \approx 0.542$$

Discussion

- NPUI appears to offer a simple, yet possibly appealing, model for utility learning.
- There has been limited discussion on the idea that preferences over decision outcomes may be uncertain, even though such scenarios have empirical support.
- How should uncertainty over preferences be incorporated within a normative decision analysis, and what are the implications of utility learning models?
- What sequential choice rule(s) should be employed?
- How to determine scaling within $[0, 1]$ interval, or more generally, how to deal with the problem of induction when the actual value realized can be far better or far worse then anything as yet observed, and when it is the actual value that is important rather than the ordinal ranking.

References

T. Augustin and F.P.A. Coolen.
Nonparametric predictive inference and interval probability.
Journal of Statistical Planning and Inference, 124:251–272, 2004.

F.P.A. Coolen.
Comparing Two Populations Based on Low Stochastic Structure Assumptions.
Statistics & Probability Letters, 29:297–305, 1996.

R.M. Cyert and M.H. DeGroot.
Adaptive Utility.
In R.H. Day and T. Groves, editors, Adaptive Economic Models, pages 223–246. Academic Press, New York, 1975.

B.M. Hill.
Posterior Distribution of Percentiles: Bayes' Theorem for Sampling from a Population.
Journal of the American Statistical Association, 63:677–691, 1968.

B.M. Hill.
De Finetti's Theorem, Induction, and $A_{(n)}$ or Bayesian Nonparametric Predictive Inference (with discussion).
In J.M. Bernardo, M.H. DeGroot, D.V. Lindley, and A.F.M. Smith, editors, *Bayesian Statistics 3*, pages 211–241. Oxford University Press, 1988.

B. Houlding.
Sequential Decision Making with Adaptive Utility.
PhD thesis, Dept. Mathematical Sciences, Durham University, UK, 2008.