Interactive Preference Learning of Utility Functions for Multi-Objective Optimization

Ian Dewancker, Michael McCourt, Samuel Ainsworth

Dec 8 2016, FILM Workshop @ NIPS

SIGOPT

SMBO as a service for helping clients solve :

$$\mathbf{x}_{\text{opt}} = \arg\max_{\mathbf{x} \in \mathcal{X}} \; f(\mathbf{x})$$

SMBO as a service for helping clients solve :

$$\mathbf{x}_{\text{opt}} = \arg\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$$

Require that $f : \mathcal{X} \to \mathbb{R}$

Real-world engineering systems more realistically discussed as a multi-objective optimization problem :

$$\mathbf{f} : \mathcal{X} \to \Omega, \qquad \mathbf{f}(\mathbf{x}) = \begin{pmatrix} f_1(\mathbf{x}) \\ \vdots \\ f_N(\mathbf{x}) \end{pmatrix}$$

Real-world engineering systems more realistically discussed as a multi-objective optimization problem :

$$\mathbf{f} : \mathcal{X} \to \Omega, \qquad \mathbf{f}(\mathbf{x}) = \begin{pmatrix} f_1(\mathbf{x}) \\ \vdots \\ f_N(\mathbf{x}) \end{pmatrix}$$

$$\mathbf{x}_{\text{opt}} = \arg\max_{\mathbf{x} \in \mathcal{X}} \; u(\mathbf{f}(\mathbf{x}))$$

$u : \Omega \to \mathbb{R}$ denotes a *utility function*

Real-world engineering systems more realistically discussed as a multi-objective optimization problem :

$$\mathbf{f} : \mathcal{X} \to \Omega, \qquad \mathbf{f}(\mathbf{x}) = \begin{pmatrix} f_1(\mathbf{x}) \\ \vdots \\ f_N(\mathbf{x}) \end{pmatrix}$$

$$\mathbf{x}_{\text{opt}} = \arg\max_{\mathbf{x} \in \mathcal{X}} \; u(\mathbf{f}(\mathbf{x}))$$

$u : \Omega \to \mathbb{R}$ denotes a *utility function*

Client 1 : *"Love idea of Bayesian global optimization, but I don't have time to develop scalar valued objective"*

Real-world engineering systems more realistically discussed as a multi-objective optimization problem :

$$\mathbf{f} : \mathcal{X} \to \Omega, \qquad \mathbf{f}(\mathbf{x}) = \begin{pmatrix} f_1(\mathbf{x}) \\ \vdots \\ f_N(\mathbf{x}) \end{pmatrix}$$

$$\mathbf{x}_{\text{opt}} = \arg\max_{\mathbf{x} \in \mathcal{X}} \; u(\mathbf{f}(\mathbf{x}))$$

$u : \Omega \to \mathbb{R}$ denotes a *utility function*

Client 1 : *"Love idea of Bayesian global optimization, but I don't have time to develop scalar valued objective"*

Client 2 : *"How can I quickly get buy-in from all stakeholders to use scalar-valued objective?"*

Propose learning a model for utility functions.

$$u(\mathbf{f}(\mathbf{x})) = \prod_{i=1}^{N} u_i(f_i(\mathbf{x}))$$

Propose learning a model for utility functions.

$$u(\mathbf{f}(\mathbf{x})) = \prod_{i=1}^{N} u_i(f_i(\mathbf{x}))$$

Utility is product of cumulative distribution functions of beta random variables :

$$u_i(f_i(\mathbf{x});\, \alpha_i, \beta_i) = \int_0^{f_i(\mathbf{x})} \frac{t^{\alpha_i - 1}(1-t)^{\beta_i - 1}}{B(\alpha_i, \beta_i)}\, \mathrm{d}t$$

$$\log(\alpha_i) \sim \mathcal{N}(\mu_{\alpha_i}, \sigma_{\alpha_i}), \quad \log(\beta_i) \sim \mathcal{N}(\mu_{\beta_i}, \sigma_{\beta_i})$$

Can we learn utilities using just binary preference data from user?

Can we learn utilities using just binary preference data from user?

Pairs of multi-objective values with clear utility preference ($\mathcal{D}_P$)
Pairs having utilties seen as equivalent ($\mathcal{D}_E$),

$$\mathcal{D}_P = \{(\mathbf{f}_1^{p_1} \prec \mathbf{f}_2^{p_1}), \ldots, (\mathbf{f}_1^{p_M} \prec \mathbf{f}_2^{p_M})\}$$
$$\mathcal{D}_E = \{(\mathbf{f}_1^{e_1} \prec\succ \mathbf{f}_2^{e_1}), \ldots, (\mathbf{f}_1^{e_L} \prec\succ \mathbf{f}_2^{e_L})\}$$

Can we learn utilities using just binary preference data from user?

Pairs of multi-objective values with clear utility preference ($\mathcal{D}_P$)
Pairs having utilties seen as equivalent ($\mathcal{D}_E$),

$$\mathcal{D}_P = \{(\mathbf{f}_1^{p_1} \prec \mathbf{f}_2^{p_1}), \ldots, (\mathbf{f}_1^{p_M} \prec \mathbf{f}_2^{p_M})\}$$
$$\mathcal{D}_E = \{(\mathbf{f}_1^{e_1} \prec\succ \mathbf{f}_2^{e_1}), \ldots, (\mathbf{f}_1^{e_L} \prec\succ \mathbf{f}_2^{e_L})\}$$

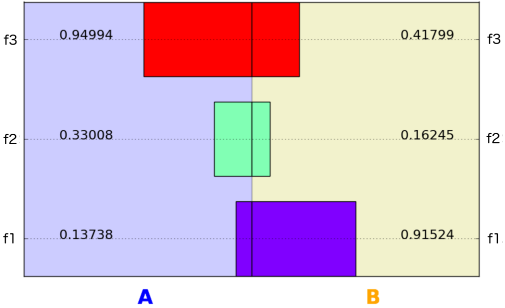

Figure : A sample comparison card for preference solicitation showing 3 metrics (f_1, f_2, f_3) in two configurations A and B.

Specify two likelihoods for observed preference data

Specify two likelihoods for observed preference data

$$
\begin{aligned}
p(\mathcal{D}_P \mid \boldsymbol{\theta}) &= \prod_{i=1}^{M} p(u(\mathbf{f}_1^{p_i}) \prec u(\mathbf{f}_2^{p_i}) \mid \boldsymbol{\theta}) \\
&= \prod_{i=1}^{M} \iint h(u_d(\mathbf{f}_1^{p_i}, \mathbf{f}_2^{p_i} \; ; \; \boldsymbol{\alpha}, \boldsymbol{\beta})) \; p(\boldsymbol{\alpha}, \boldsymbol{\beta} \mid \boldsymbol{\theta}) \; d\boldsymbol{\beta} \; d\boldsymbol{\alpha} \\
p(\mathcal{D}_E \mid \boldsymbol{\theta}) &= \prod_{j=1}^{L} p(u(\mathbf{f}_1^{e_j}) \prec\succ u(\mathbf{f}_2^{e_j}) \mid \boldsymbol{\theta}) \\
&= \prod_{j=1}^{L} \iint 2 \; p(u_E \leq -\left|u_d(\mathbf{f}_1^{e_j}, \mathbf{f}_2^{e_j} \; ; \; \boldsymbol{\alpha}, \boldsymbol{\beta})\right|) p(\boldsymbol{\alpha}, \boldsymbol{\beta} \mid \boldsymbol{\theta}) d\boldsymbol{\beta} \; d\boldsymbol{\alpha}
\end{aligned}
$$

where h is the Heaviside function and $u_E \sim \mathcal{N}(0, \sigma_E)$ We estimate both $p(\mathcal{D}_P \mid \boldsymbol{\theta})$ and $p(\mathcal{D}_E \mid \boldsymbol{\theta})$ using Monte Carlo techniques.

How to generate preference queries?

How to generate preference queries?

Algorithm 2 Active Utility Function Pref. Learning

Input: Ω

$\mathcal{D}_P, \mathcal{D}_E \leftarrow \textsc{InitUserPrefs}(\Omega)$

for $i \leftarrow 1$ **to** T **do**

  $\boldsymbol{\theta}_{MLE} \leftarrow \arg\max_{\boldsymbol{\theta}} \; p(\mathcal{D}_P \mid \boldsymbol{\theta}) \; p(\mathcal{D}_E \mid \boldsymbol{\theta})$

  $\mathbf{f}_A, \mathbf{f}_B \leftarrow \arg\max_{\mathbf{f}_1, \mathbf{f}_2 \in \Omega} \; a(\mathbf{f}_1, \mathbf{f}_2 \; ; \; \boldsymbol{\theta}_{MLE})$

  $p \leftarrow \textsc{GetUserPref}(\mathbf{f}_A, \mathbf{f}_B)$ ({ A,B,E })

  if $p == \text{E}$

    $\mathcal{D}_E \leftarrow \mathcal{D}_E \cup \; (\mathbf{f}_A \prec\succ \mathbf{f}_B)$

  else

    $\mathcal{D}_P \leftarrow \mathcal{D}_P \cup \; (\mathbf{f}_o \prec \mathbf{f}_p)$

end for

How to generate preference queries?

Algorithm 3 Active Utility Function Pref. Learning

Input: Ω
$\mathcal{D}_P, \mathcal{D}_E \leftarrow \textsc{InitUserPrefs}(\Omega)$
for $i \leftarrow 1$ **to** T **do**
  $\boldsymbol{\theta}_{MLE} \leftarrow \arg\max_{\boldsymbol{\theta}}\ p(\mathcal{D}_P \mid \boldsymbol{\theta})\ p(\mathcal{D}_E \mid \boldsymbol{\theta})$
  $\mathbf{f}_A, \mathbf{f}_B \leftarrow \arg\max_{\mathbf{f}_1, \mathbf{f}_2 \in \Omega}\ a(\mathbf{f}_1, \mathbf{f}_2\ ;\ \boldsymbol{\theta}_{MLE})$
  $p \leftarrow \textsc{GetUserPref}(\mathbf{f}_A, \mathbf{f}_B)$ ({ A,B,E })
  if $p == \mathrm{E}$
    $\mathcal{D}_E \leftarrow \mathcal{D}_E \cup\ (\mathbf{f}_A \prec\succ \mathbf{f}_B)$
  else
    $\mathcal{D}_P \leftarrow \mathcal{D}_P \cup\ (\mathbf{f}_o \prec \mathbf{f}_p)$
end for

$$a(\mathbf{f}_1, \mathbf{f}_2\ ;\ \boldsymbol{\theta}_{\mathrm{MLE}}) = \mathrm{Var}(u_d(\mathbf{f}_1, \mathbf{f}_2 \mid \boldsymbol{\alpha}_{\mathrm{MLE}},\ \boldsymbol{\beta}_{\mathrm{MLE}})),$$

Since the learned utility is a product of N cdfs, these distributions can be independently visualized as an introspection mechanism.

Since the learned utility is a product of N cdfs, these distributions can be independently visualized as an introspection mechanism.

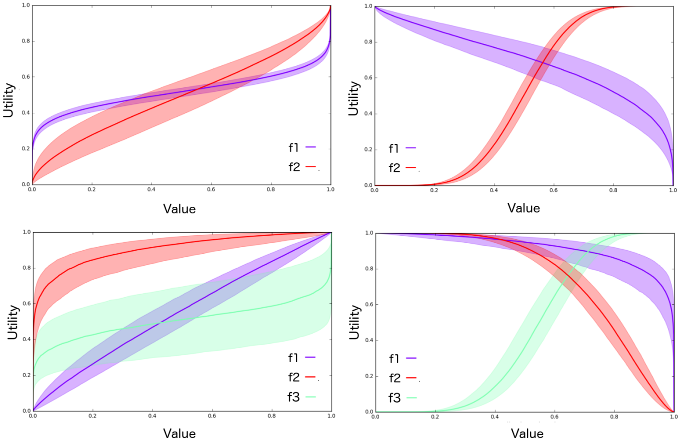

Figure : Plots of learned independent utilities with mean and IQR

Explicit test utility functions were used to simulate implicit human utility functions. Hold-out set of multi-objective configurations and

Kendall rank correlation used to quantify the ordinal association between the test utility values and the mean learned utility function values

Table : Kendall-Tau Correlation using Different Search Policies

Test Utility Function	Rnd Search	Pair Entropy
max $f_1 + 2f_2$	**0.8756**	0.8618
max $f_1 + 10f_2$	0.9422	**0.9615**
min f_1 s.t. $f_2 > 0.6$	0.6507	**0.6893**
max $2\, f_1 f_2 \,/\, (f_1 + f_2)$	0.8844	**0.9039**
max $5\, f_1 f_2 \,/\, (4f_1 + f_2)$	0.8949	**0.9120**
max $f_1 + 2f_2 + f_3$	**0.8490**	0.7805
max $5f_1 + 2f_2 + f_3$	**0.8738**	0.8311
min f_1 s.t. $f_2 > 0.6, f_3 < 0.2$	0.2949	**0.3257**