Closure Operators, Classifiers and Desirability

Alessio Benavoli[1] Alessandro Facchini[2] Marco Zaffalon[2]

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

ISIPTA 2023, Oviedo

July 11th, 2023

Motivation

Desirability/probability theory assumes linearity of the utility-scale in which rewards are measured.

Miranda and Zaffalon:

1. gaining money is desirable;
2. losing money is undesirable;
3. the value of money is measured on a logically consistent utility-scale determined by a closure operator

Casanova, Benavoli, Zaffalon:

1. classify nonnegative gambles as $+1$;
2. classify negative gambles as -1;
3. the value of money is measured on a utility-scale represented by a "nonlinear" classifier.

Are these approaches different/equivalent? What is the difference between utility and utility-scale of a closure operator/classifier?

Belief Structures (de Cooman 2005)

$$g_1 = \begin{bmatrix} -1 \\ +2 \end{bmatrix},\ g_2 = \begin{bmatrix} -0.5 \\ +3 \end{bmatrix},\ g_3 = \begin{bmatrix} +2 \\ -1 \end{bmatrix}$$

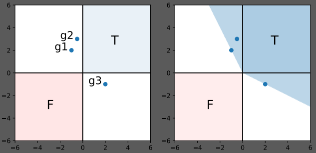

A map $K : \mathcal{P}(\mathcal{L}) \to \mathcal{P}(\mathcal{L})$ is a *closure operator* on $\mathcal{L}$ if:

(K1–Extensiveness) $A \subseteq K(A)$;

(K2–Monotonicity) if $A \subseteq A'$ then $K(A) \subseteq K(A')$;

(K3–Idempotency) $K(K(A)) = K(A)$.

A belief model is an element of $A \in \mathcal{P}(\mathcal{L})$ and is closed if $K(A) = A$.

Definition

A *belief structure* is a triple $\mathfrak{B} := (\mathcal{P}(\mathcal{L}), K, \mathcal{C})$ where

- $K : \mathcal{P}(\mathcal{L}) \to \mathcal{P}(\mathcal{L})$ is a closure operator
- $\mathcal{C}$ is a consistency set.

Generalised Desirability Theories

Let $\mathfrak{B} := (\mathcal{P}(\mathcal{L}), K, \mathcal{C})$ be a belief structure over $\mathcal{L}$.
We call it a *generalised almost-desirability theory* (GADT for short) whenever the following properties are satisfied:

(G1) $K(\mathcal{L}_0^+) = \mathcal{L}_0^+$ is the minimal element of its consistency set $\mathcal{C}$, and in particular $K(A) \supseteq \mathcal{L}_0^+$, for every $A \in \mathcal{P}(\mathcal{L})$

(G2) $K_{order} \leq K$,

(G3) $K(A) \in \mathcal{C}$ if and only if $K(A) \cap \mathcal{L}^- = \emptyset$.

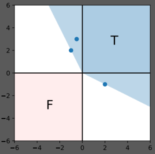

Conic hull

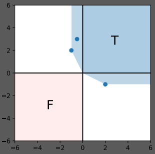

Convex hull

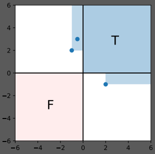

Pareto hull

Weak order/Utility $\Rightarrow$ Classifier

A *weak order* $\succeq$ is a *transitive* and *total* binary relation.
We define the *support function* of a set $A \subseteq \mathcal{L}$ with respect to $\succeq$ as the collection of all order-equivalence infima of A w.r.t. $\succeq$:

$$s_{\succeq}(A) := \{h \in \underline{A} \mid h \succeq g, \forall g \in \underline{A}\}, \tag{1}$$

where $\underline{A} := \{g \in \mathcal{L} \mid \ f \succeq g, \ \ \forall f \in A\}$. We then define the *support half-space* of the set $A \neq \emptyset$ as

$$S_{\succeq}(A) := \{g \in \mathcal{L} \mid \ \ g \succeq f \ , \ \forall f \in s_{\succeq}(A)\}. \tag{2}$$

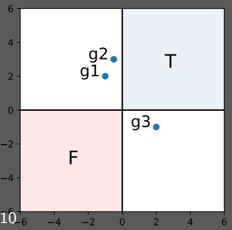

$g_2 \succeq g_1 \succeq g_3,$
$s_{\succeq}(A) = g_3, \quad S_{\succeq}(A) = \{g \in \mathcal{L} : g \succeq g_3\}.$

Closure Operators $\Leftrightarrow$ Classifiers

Let K be a closure operator over $\mathcal{P}(\mathcal{L})$, which satisfies dominance (resp. continuity). Then there exists a family of order-preserving (resp. order-continuous) weak-orders such that:

$$K(A) = \bigcap_{i \in \mathcal{I}} (S_{\succeq_i}(A) \cup \mathrm{T}), \tag{3}$$

for all $A \subseteq \mathcal{L}$. Conversely, for any family $\{\succeq_i \mid i \in \mathcal{I}\}$ of order-preserving (resp. order-continuous) weak-orders, and a sequence $(X_i : i \in \mathcal{I})$ where each X_i satisfies dominance (and is continuous), the map $\kappa : \mathcal{P}(\mathcal{L}) \to \mathcal{P}(\mathcal{L})$ defined as

$$\kappa(A) := \bigcap_{i \in \mathcal{I}} (S_{\succeq_i}(A) \cup X_i) \tag{4}$$

is a closure operator that satisfies dominance (resp. continuity) and such that $\mathrm{T} = \bigcap_{i \in \mathcal{I}} X_i$.

Closure Operators $\Leftrightarrow$ Utility

If $\succeq$ is an order-preserving order-continuous weak-order on $\mathcal{L}$. Then there is a non-decreasing order-continuous utility function $u : \mathcal{L} \to \mathbb{R}$ that represents $\succeq$ and vice versa, that is

$$f \succeq g \quad \text{iff} \quad u(f) \geq u(g).$$

This leads to the following equivalent definition of support half-space:

$$S_{u_i}(A) = \left\{ g \in \mathcal{L} \mid \quad u_i(g) \geq \sup_{h \in \underline{A}} u_i(h) \right\} \overset{example}{=} \{ g \in \mathcal{L} \mid \quad u(g) \geq u(g_3) \}$$

where $\underline{A} = \{ g \in \mathcal{L} \mid \; u_i(f) \geq u_i(g), \;\; \forall f \in A \}$. Hence, the weak-order plays the role of the utility-scale associated to the closure-operator $S_{\succeq_i}$.

By changing the utility function, we can derive different models of nonlinear desirability proposed in literature.

Particular cases

Standard almost-desirability has linear-utility $u_i(g) := p_i^\top g$, where p_i is a probability vector, leading to $K(A) := \bigcap_{i \in \mathcal{I}} (S_{\succeq_i}(A) \cup \mathcal{L}_0^+)$.

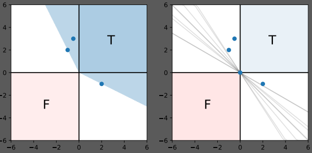

Almost-desirability with convex-hull closure operator has also linear utility:

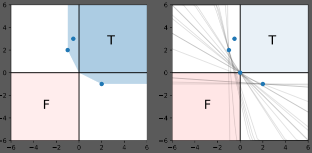

$$S_{u_i}(A) = \left\{ g \in \mathcal{L} \mid \quad u_i(g) \geq \sup_{h \in \underline{A}} u_i(h) \right\},$$

Particular cases

Using Chebyshev-utility $u_i(g) := \max_{j=1,\dots,n} p_{ij}(g_j - c_{ij})$ results in a closure operator that only preserves the order of the vector-space:

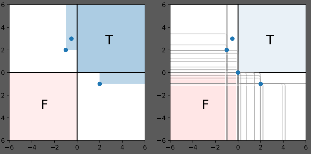

Almost-desirability with utility $u_i(g) := u_i(g) := p_i^\top (g - c_i)^d$:

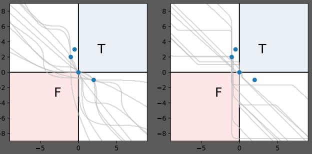

for $d = 3$ (left) and $d = 31$ (right).
A lower prevision defines a single utility-function and a single support half-space. Imprecision and utility blend in nonlinear desirability.

We now accept Euros!

Isipta 2023

Oviedo, Spain

VOL.XIII…No.2 | JULY 11, 2023 | 1 LOTTERY TICKET

Closure Operators, Classifiers and Desirability

By BENAVOLI, FACCHINI, ZAFFALON

At the core of desirability theory lies an assumption of linearity of the utility-scale in which rewards are measured. Recently, Miranda and Zaffalon proposed a unifying theoretical framework to extend linear-desirability theory to the nonlinear case by letting the utility-scale be represented by a general closure operator. This framework retains the overall logical structure of linear-desirability theory, which is based on the following axioms: (i) gaining money is desirable; (ii) losing money is undesirable; but replaces the linearity axiom with: (iii) the value of money is measured on a logically consistent utility-scale determined by a closure operator.

A more operational approach to extend linear-desirability to the nonlinear case was pursued by Casanova, Benavoli and Zaffalon. This approach starts from the observation that the logical consistency of a set of linearly-desirable gambles can be checked by solving a binary linear classification problem. Then the authors extend desirability to the nonlinear case by instead considering a binary nonlinear classification problem. This framework imposes the logical constraints of desirability theory by forcing the classifier to separate the non-negative gambles (gaining money is desirable) from the negative ones (losing money is undesirable).

The present article reviews and compares these two methods to extend desirability to the nonlinear case. It shows how they are related and how they can be used to represent various nonlinear variants of desirability. It also uncovers the utility-scale implied by the closure operator.

This is obtained in three steps. First, since this connection follows by standard basic results in lattice theory and algebraic logic, we formalise these results in the framework of *belief structures* introduced by de Cooman in 2005. To deal with nonlinearity, we slightly need to extend this framework by observing that:

"In the Belief Structures framework, the notion of closure operator and that of consistency can be conceptually separated." Same separation can be used to define the morphisms.

For instance in linear almost-desirability, the closure operator is the conic hull and the consistency set (predicate) corresponds to the set of coherent closed convex cones. This enables us to provide a generalised definition of *almost-desirability theory*.

"A Belief Structure is called a generalised almost-desirability theory, whenever (1) the set of nonnegative gambles is the minimal element of its consistency set; (2) the closure operator preserves the order of the underlying vector space; (3) for any set of gambles its closure, by the closure operator, is consistent provided that does not include negative gambles." Strict and other variants of desirability can be defined analogously.

Second, we introduce weak-orders $\succeq$ to connect closure operators to classifiers. A weak-order on a set (of gambles) is a binary relation which is transitive and complete. This lets us define the *support functions* $s_{\succeq}(A)$ of a non-empty set of gambles $A \subseteq \mathcal{L}$ as the largest element in $\mathcal{L}$ that is no greater than any element of A under the weak-order. We then define the *support half-space* of the set A as

$$S_{\succeq}(A) := \{g \in \mathcal{L} \mid \; g \succeq f \;, \forall f \in s_{\succeq}(A)\},$$

and $S_{\succeq}(\emptyset) = \emptyset$.

It can then be proven that, for any closure operator K over set of gambles, which preserves the order of the vector space $\mathcal{L}$ (and satisfies continuity), there exists a family of order-preserving (resp. order-continuous) weak-orders s.t.

$$K(A) = \bigcap_{i \in \mathcal{I}} (S_{\succeq_i}(A) \cup \mathsf{T}),$$

where $K(\emptyset) = \mathsf{T}$ (e.g., $\mathsf{T} = \mathcal{L}_0^+$). Vice-versa, any intersections of support half-spaces defines a closure operator. This implies that $S_{\succeq}$ is also a closure operator. We can think of this result as a "generalisation" of the standard *separating hyperplane theorem* from convex geometry. It states that any set $K(A)$ (convex or not convex) can be expressed as the intersection of support half-spaces.

Since $S_{\succeq}(A)$ is a support half-space, we can call the closure operator $S_{\succeq}$ defined by the weak-order $\succeq$ a **binary (nonlinear) classifier**.

Particular cases

If $\succeq$ is an order-preserving order-continuous weak-order on $\mathcal{L}$. Then there is a non-decreasing order-continuous utility function $u : \mathcal{L} \to \mathbb{R}$ that represents $\succeq$ and vice versa, that is

$$f \succeq g \;\text{ iff }\; u(f) \geq u(g).$$

This leads to the following equivalent definition of support half-space:

$$S_{\succeq_i}(A) = \left\{ g \in \mathcal{L} \mid \; u_i(g) \geq \sup_{h \in \underline{A}} u_i(h) \right\},$$

where $\underline{A} = \{g \in \mathcal{L} \mid \; u_i(f) \succeq u_i(g), \;\; \forall f \in A\}$. Hence, the weak-order plays the role of the utility-scale associated to the closure-operator $S_{\succeq_i}$. By changing the utility function, we can derive different models of nonlinear desirability proposed in literature.

Standard almost-desirability has linear-utility $u_i(g) := p_i^\top g$, where p_i is a probability vector, leading to $K(A) := \bigcap_{i \in \mathcal{I}} (S_{\succeq_i}(A) \cup \mathcal{L}_0^+)$.

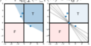

Almost-desirability with convex-hull closure operator has also linear utility:

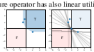

Using Chebyshev-utility $u_i(g) := \max_{j=1,\dots,n} p_{ij}(g_j - c_{ij})$ results in a closure operator that only preserves the order of the vector-space:

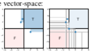

Almost-desirability with utility $u_i(g) := u_i(g) := p_i^\top (g - c_i)^d$ leads to:

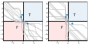

for $d = 3$ (left) and $d = 31$ (right).

Note that, a lower prevision defines a single utility-function and a single support half-space. Thus, imprecision and utility blend in nonlinear desirability.