A Similarity Model for Subjective Assessment

Sarra Tajouri   Alexis Tsoukias

LAMSADE-CNRS, Univ Paris Dauphine – PSL, Paris, France

DA2PL 2026

Similarity is at the heart of consequential decisions

Hiring

Is this candidate similar enough to our past successful hires?

Credit scoring

Does this borrower's profile resemble those who repaid?

Admissions

Which applicants are comparable to our admitted students?

Housing allocation

Who among the applicants has the most comparable needs?

Two comparison modes

Historical: new individual vs. past accepted cases

Cross-candidate: applicants ranked against one another

Why it matters

The similarity measure *encodes* which differences count and how much, it is a **normative choice**, not a neutral technical tool.

Different measures ⇒ different rankings ⇒ different outcomes.

Similarity is a subjective experience

When individuals *judge* similarity, several psychological phenomena arise:

- People do not weight all features equally
- Perceived similarity need not be **reciprocal**
- Upward and downward social comparisons are experienced **asymmetrically**
- The same pair can be judged more or less similar **depending on who is doing the judging**

The Problem with Metric Similarity

Metric axioms

- Non-negativity: $d(x, y) \geq 0$
- Symmetry: $d(x, y) = d(y, x)$
- Triangle inequality: $d(x, z) \leq d(x, y) + d(y, z)$

The Problem with Metric Similarity

Metric axioms

- Non-negativity: $d(x, y) \geq 0$
- *Symmetry:* $d(x, y) = d(y, x)$
- Triangle inequality: $d(x, z) \leq d(x, y) + d(y, z)$

Symmetry fails — admission example

Two candidates: x: score $= 20$, school A;
y score $= 20$, school B.

- y perceives herself as *similar to* x: "same results, same skills — we are comparable."
- x does *not* perceive y as similar: school prestige is a salient distinctive feature *from her reference point*.

Similarity is **directed**.

The Problem with Metric Similarity

Metric axioms

- Non-negativity: $d(x, y) \geq 0$
- Symmetry: $d(x, y) = d(y, x)$
- *Triangle inequality:* $d(x, z) \leq d(x, y) + d(y, z)$

Triangle inequality fails — admissions example

Three students: i: high school A, physics;
j: high school A, mathematics;
k: high school B, mathematics.

- $i \approx j$: high school is a salient shared feature for i
- $j \approx k$: field is a salient shared feature for j
- $i \not\approx k$: *no shared attribute* — different schools, different fields

Each pair considers a **different** set of features salient.

The Problem with Metric Similarity

Empirical conclusion [4]

Human similarity judgements systematically violate symmetry and the triangle inequality.

Theoretical conclusion [1]

Metric axioms are *necessary and sufficient* for a geometric embedding — therefore *no metric* can faithfully represent human judgements.

A similarity measure over human judgements should handle

- *Non-symmetry*
- *Context-dependence / Feature saliency*

Tversky's ratio model [4]

Individuals represented as **feature sets** A, B. Similarity defined through **feature matching**:

$$S(a,b) = \frac{\theta\,\sigma \overbrace{(A\cap B)}^{\text{shared features}}}{\theta\,\sigma \underbrace{(A\cap B)}_{\text{shared}} + \alpha\,\sigma \underbrace{(A\setminus B)}_{\text{distinctive to } a} + \beta\,\sigma \underbrace{(B\setminus A)}_{\text{distinctive to } b}}$$

- $\sigma(A\cap B)$: shared features
- $\sigma(A\setminus B)$: features of A absent from B
- $\sigma(B\setminus A)$: features of B absent from A

Key property

Non-symmetric when $\alpha \neq \beta$:
$S(A,B) \neq S(B,A)$

Limitation: designed for binary/set-valued features only.
Real profiles mix *nominal* attributes (field of study) and *ordinal* ones (years of experience, education level).

Our contribution

This paper proposes a principled extension of the ratio model to mixed-type profiles.

(i) Define a **directed similarity measure** for nominal + ordinal profiles

(ii) Establish **formal properties**: self-similarity, boundedness, monotonicity

(iii) Provide a **behavioral interpretation** of (α, β, θ) in terms of social orientation

(iv) Show α is **identifiable from ordinal ranking data** and propose three practical **elicitation protocols**

A Directed Similarity Measure for Mixed-Type Profiles

$\mathcal{F}$

set of all features

partition for (i, j) $\longrightarrow$

$$\mathcal{F}_{=}(i,j) = \{l \in \mathcal{F} : f^l(i) = f^l(j)\}$$
$$\mathcal{F}_{\neq}^{\text{ord}}(i,j) = \{l \in \mathcal{F} : f^l(i) \neq f^l(j),\ \mathcal{V}_l \text{ ordinal}\}$$
$$\mathcal{F}_{\neq}^{\text{nom}}(i,j) = \{l \in \mathcal{F} : f^l(i) \neq f^l(j),\ \mathcal{V}_l \text{ nominal}\}$$

A Directed Similarity Measure for Mixed-Type Profiles

Feature partition of i and j: $\mathcal{F}_{=}(i,j) = \{l : f^l(i) = f^l(j)\}$; $\mathcal{F}_{\neq}^{\text{ord}}(i,j) = \{l : f^l(i) \neq f^l(j),\ \mathcal{V}_l \text{ ord.}\}$; $\mathcal{F}_{\neq}^{\text{nom}}(i,j) = \{l : f^l(i) \neq f^l(j),\ \mathcal{V}_l \text{ nom.}\}$

Intersection term

$$\mathcal{I}(i,j) = \underbrace{|\mathcal{F}_{=}(i,j)|}_{\substack{\text{Shared features}\\ \text{with same values}}} + \sum_{l \in \mathcal{F}_{\neq}^{\text{ord}}(i,j)} \frac{\overbrace{\min(f^l(i), f^l(j))}^{\text{overlap}} - \min \mathcal{V}_l}{\max \mathcal{V}_l - \min \mathcal{V}_l}$$

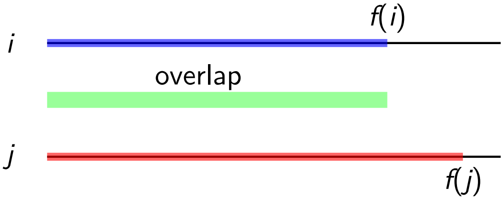

Directed gaps

$$D^{+}(i,j) = |\mathcal{F}_{\neq}^{\text{nom}}| + \sum_{l \in \mathcal{F}_{\neq}^{\text{ord}}} \max\left(0, \overbrace{\frac{f^l(i) - f^l(j)}{\text{range}_l}}^{\text{gap}}\right)$$

$$D^{-}(i,j) = |\mathcal{F}_{\neq}^{\text{nom}}| + \sum_{l \in \mathcal{F}_{\neq}^{\text{ord}}} \max\left(0, \overbrace{\frac{f^l(j) - f^l(i)}{\text{range}_l}}^{\text{gap}}\right)$$

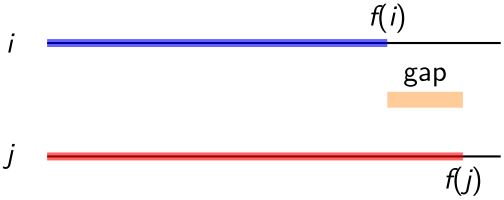

Similarity formula and properties

$$\mathrm{sim}(x,y) = \frac{\theta\,\mathcal{I}(x,y)}{\theta\,\mathcal{I}(x,y) + \alpha\,D^{+}(x,y) + \beta\,D^{-}(x,y)}$$

with $\theta, \alpha, \beta \geq 0$ and $\alpha + \beta = 1$.

Why $\alpha + \beta = 1$?
A nominal mismatch (contributing $\alpha + \beta$ to the denominator) is always penalized more than a maximally different ordinal feature (contributing $\max(\alpha, \beta) \cdot \delta \leq \max(\alpha, \beta)$).

The constraint makes penalties *commensurable* across feature types (in $[0, 1]$).

Properties

- $\mathrm{sim}(x,x) = 1$
- $\mathrm{sim}(x,y) \in [0,1]$
- Non-symmetric when $\alpha \neq \beta$
- Monotonicity: if $\mathcal{I}(i,j) \nearrow$, $\mathrm{sim}(i,j) \nearrow$; if $D^{+}(i,j) \nearrow$ or $D^{-}(i,j) \nearrow$, $\mathrm{sim}(i,j) \searrow$

Two standard comparable measures

Gower similarity [3]

Weighted average of per-feature similarities:

$$G(x, y) = \frac{\sum_{j \in \mathcal{F}} w_j \, s_j(x, y)}{\sum_{j \in \mathcal{F}} w_j}$$

Shared with our model: weights w_j allow *feature saliency*.

Fuzzy Jaccard (Ruzicka) [5]

Generalizes Jaccard to continuous memberships:

$$J(x, y) = \frac{\sum_{j \in \mathcal{F}} \min(f^j(x), f^j(y))}{\sum_{j \in \mathcal{F}} \max(f^j(x), f^j(y))}$$

Shared with our model: the min/max structure mirrors our *intersection* $\mathcal{I}$ and *union* $\mathcal{I} + D^+ + D^-$ terms.

Both measures are **symmetric** by construction — they cannot capture the directionality of similarity judgments.

Illustrative example

8 individuals, 4 features (2 nominal, 2 ordinal).

Parameters: $\theta = 1$, $\alpha = 0.4$, $\beta = 0.6$.

Feature	**Type**	**Domain**	**Bounds (if ordinal)**
F_1	Nominal	$\{H, F, N\}$	–
F_2	Nominal	$\{B, I\}$	–
F_3	Ordinal	$\{1, 2, 3\}$	$\min = 1, \ \max = 3$
F_4	Ordinal	$[0, 4]$	$\min = 0, \ \max = 4$

Table: Feature types and domains.

Individual profiles are given by:

$$x_1 = \langle H, B, 2, 1.5 \rangle, \qquad x_2 = \langle F, B, 1, 1 \rangle, \tag{1}$$
$$x_3 = \langle N, I, 3, 3.2 \rangle, \qquad x_4 = \langle H, I, 2, 2.9 \rangle, \tag{2}$$
$$x_5 = \langle F, B, 3, 1.5 \rangle, \qquad x_6 = \langle N, B, 3, 3.9 \rangle, \tag{3}$$
$$x_7 = \langle H, I, 1, 2.1 \rangle, \qquad x_8 = \langle F, I, 2, 1.2 \rangle. \tag{4}$$

Illustrative example

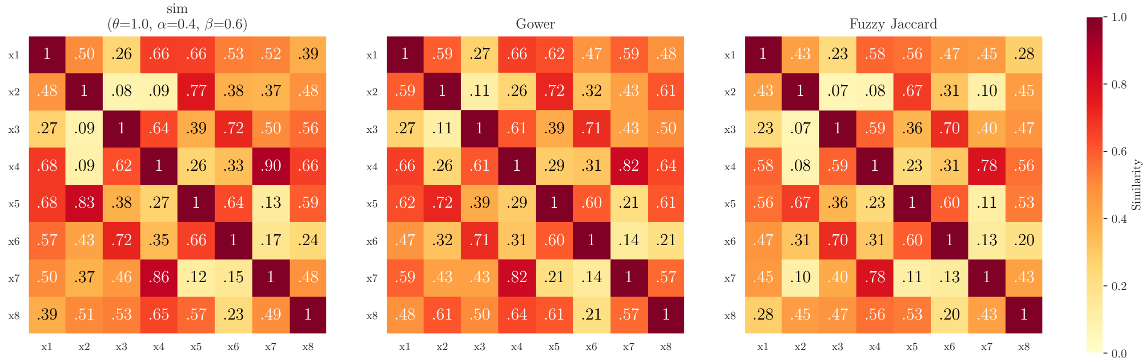

Pair (x_2, x_5)**:** x_5 exceeds x_2 on both ordinal features.

- x_5's perspective: x_2 falls short $\Rightarrow D^-$ is small
- x_2's perspective: x_5 exceeds $\Rightarrow$ large D^-, penalized by β

$$\mathrm{sim}(x_5, x_2) = 0.83 > \mathrm{sim}(x_2, x_5) = 0.77$$

Behavioral Interpretation of Parameters

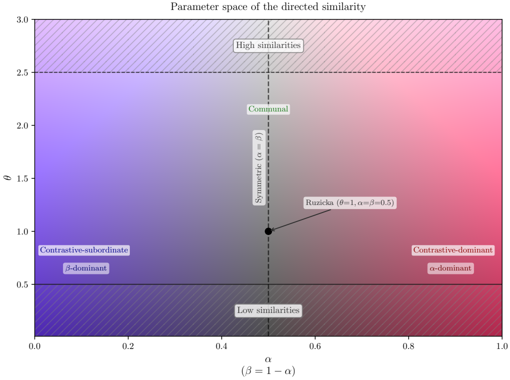

Parameter space (θ, α, β) with $\alpha + \beta = 1$.

Light regions: shared features dominate.

Dark: distinctive features dominate.

Three perception profiles [2]:

- *Communal* ($\theta \gg \alpha, \beta$)
  Similarity driven by what is shared. Differences matter little — characteristic of group-oriented, collectivist orientations.
- *Contrast-dominant* (high α, low β)
  Sensitive to own advantages over others. Perceives those below as dissimilar — status-conscious, hierarchical orientation.
- *Contrast-subordinate* (low α, high β)
  Discounts own advantages but strongly perceives those who surpass her as dissimilar — upward comparison sensitivity, inferiority aversion.

Identifiability from ordinal data

Which parameters can be recovered from ordinal similarity judgments?

θ is *not* identifiable

θ_i cancels from all ordinal comparisons.
Recovering θ requires **cardinal** judgments.
θ controls the amplitude of similarity scores.

α *is* identifiable

Only α_i (and $\beta_i = 1 - \alpha_i$) drives ordinal comparisons.
This is exactly the parameter encoding perception profiles.

Elicitation protocols

Three protocols to recover α from similarity judgments:

Full Rankings

Rank K others by similarity. $\binom{K}{2}$ pairwise constraints per query.

2-stage estimation: grid search (Kendall τ) + hinge-loss refinement (L-BFGS-B).

Best-Worst Scaling

Identify only most and least similar from group of K. $2K-3$ constraints per query (vs. $\binom{K}{2}$).

Less cognitive burden, same estimation procedure.

Diagnostic Reference Profiles

Synthetic profiles A, B designed to discriminate specific α thresholds. Single binary question per query.

Soft voting estimator, $O(1/Q)$ worst-case error.

Elicitation Results: Three Protocols Compared

Monte Carlo validation (n=100 individuals, diverse $\alpha_i \in [0, 1]$)

Method	Queries	Pairs	MAE
Full rank.	5	950	0.017
BWS	5	185	0.111
Ref. profiles	20	20	**0.018**

Reference profiles: robust to noise — at 10% flipped responses, MAE rises only to 0.054.

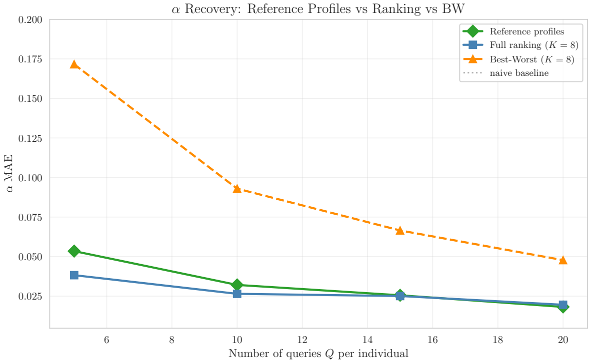

MAE as a function of number of queries Q for each protocol.

Takeaway

Reference profiles achieve near-full-ranking accuracy with only *20 binary questions* per individual, making population-scale elicitation practically feasible.

Conclusion & Applications

Summary

- Directed similarity for mixed-type profiles
- α encodes social orientation, identifiable from rankings
- Reference profiles: efficient elicitation with 20 binary questions

Frameworks where *individual perspective matters*:

- **Individual fairness**
- **Social network analysis**

Future directions

- Decision-level evaluation (top-k ranking disagreements across measures)
- Graded nominal penalties via ontology or learned inter-category similarities
- Robustness under real human inconsistencies and intransitivities

Thank you!

References I

- Richard Beals, David H Krantz, and Amos Tversky.
  Foundations of multidimensional scaling.
  Psychological review, 75(2):127, 1968.

- Bram P. Buunk and Jan F. Ybema.
  Social comparisons and occupational stress: The identification-contrast model.
  In Bram P. Buunk and Frederick X. Gibbons, editors, *Health, Coping, and Well-Being: Perspectives from Social Comparison Theory*, pages 359–388. Lawrence Erlbaum Associates, Mahwah, NJ, 1997.

- John C Gower.
  A general coefficient of similarity and some of its properties.
  Biometrics, 27(4):857–871, 1971.

References II

Amos Tversky.
Features of similarity.
Psychological review, 84(4):327, 1977.

Lotfi A Zadeh.
Fuzzy sets.
Information and control, 8(3):338–353, 1965.