Computing linear rankings from trillions of pairwise outranking situations

Raymond Bisdorff

University of Luxembourg
CSC/ILIAS

DA2PL'2016
Paderborn, November 2016

Content

Motivation: Showing a performance tableau

Consider a performance table showing the service quality of 12 commercial cloud providers measured by an external auditor on 14 incommensurable performance criteria.

criterion	upT	dwT	ouT	LB	MTBF	Rcv	Lat	RspT	Thrpt	stoC	snpC	auT	enC	auD
Amz	2	2	2	4	3	3	NA	3	NA	4	NA	4	4	4
Cen	4	4	0	4	4	4	NA	2	NA	3	NA	4	4	4
Cit	2	4	2	4	3	4	NA	2	NA	3	4	4	4	4
Dig	2	1	4	4	3	3	NA	2	NA	3	NA	4	4	4
Ela	4	4	0	4	4	4	NA	4	NA	3	4	4	4	4
GMO	1	3	4	4	3	2	NA	4	NA	3	NA	4	4	4
Ggl	4	2	1	4	2	3	NA	2	NA	4	4	4	4	4
HP	3	3	2	4	4	3	NA	4	NA	3	4	4	4	4
Lux	2	2	2	4	3	3	NA	2	NA	2	NA	4	4	4
MS	4	4	0	4	4	4	NA	4	NA	4	NA	4	4	4
Rsp	NA	NA	NA	4	NA	3	NA	NA	NA	3	4	4	4	4
Sig	4	4	0	4	4	4	NA	3	NA	3	4	4	4	4

Legend: 0 = '*very weak*', 1 = '*weak*', 2 = '*fair*', 3 = '*good*', 4 = '*very good*','*NA*' = missing data; '*green*' and '*red*' mark the **best**, respectively the **worst**, performances on each criterion.

Motivation: showing an ordered heat map

The same performance tableau may be optimistically colored with the highest 7-tiles class of the marginal performances and presented like a heat map,

Ranking of cloud providers by service quality

criteria	dwT	Rcv	MTBF	upT	RspT	stoC	auD	enC	auT	snpC	Thrpt	Lat	LB	ouT
weights	2.00	2.00	2.00	2.00	2.00	3.00	1.00	1.00	1.00	3.00	2.00	2.00	2.00	2.00
tau(*)	0.56	0.44	0.44	0.41	0.33	0.14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	-0.45
MS	4	4	4	4	4	4	4	4	4	NA	NA	NA	4	0
Ela	4	4	4	4	4	3	4	4	4	4	NA	NA	4	0
Sig	4	4	4	4	3	3	4	4	4	4	NA	NA	4	0
Cen	4	4	4	4	2	3	4	4	4	NA	NA	NA	4	0
HP	3	3	4	3	4	3	4	4	4	4	NA	NA	4	2
Cit	4	4	3	2	2	3	4	4	4	4	NA	NA	4	2
GMO	3	2	3	1	4	3	4	4	4	NA	NA	NA	4	4
Ggl	2	3	2	4	2	4	4	4	4	4	NA	NA	4	1
Rsp	NA	3	NA	NA	NA	3	4	4	4	4	NA	NA	4	NA
Amz	2	3	3	2	3	4	4	4	4	NA	NA	NA	4	2
Dig	1	3	3	2	2	3	4	4	4	NA	NA	NA	4	4
Lux	2	3	3	2	2	2	4	4	4	NA	NA	NA	4	2

Color legend:

quantile	20.00%	40.00%	60.00%	80.00%	100.00%

(*) tau: Ordinal (Kendall) correlation between marginal criterion and global ranking relation.

eventually linearly ordered, following for instance the Copeland ranking rule, from the best to the worst performing alternatives (ties are lexicographically resolved).

How to rank big performance tableaux ?

- The Copeland ranking rule is based on crisp net flows requiring the in- and out-degree of each node in the outranking digraph;
- When the order n of the outranking digraph becomes big (several thousand or millions of alternatives), this requires handling a huge set of n^2 pairwise outranking situations;
- We shall present hereafter a sparse model of the outranking digraph, where we only keep a linearly ordered list of diagonal quantiles equivalence classes with local outranking content.

Performance quantiles

- Let X be the set of n potential decision alternatives evaluated on a single real performance criteria.
- We denote $x, y, \ldots$ the performances observed of the potential decision actions in X.
- We call quantile $q(p)$ the performance such that $p\%$ of the observed n performances in X are less or equal to $q(p)$.
- The quantile $q(p)$ is estimated by linear interpolation from the cumulative distribution of the performances in X.

Performance quantile classes

- We consider a series: $p_k = k/q$ for $k = 0, \ldots q$ of $q + 1$ equally spaced quantiles limits like
  - quartiles limits: $0, .25, .5, .75, 1$,
  - quintiles limits: $0, .2, .4, .6, .8, 1$,
  - deciles limits: $0, .1, .2, \ldots, .9, 1$, etc
- The upper-closed q^k class corresponds to the interval $]q(p_{k-1}); q(p_k)]$, for $k = 2, \ldots, q$, where $q(p_q) = \max_X x$ and the first class gathers all data below p_1: $]-\infty; q(p_1)]$.
- The lower-closed q_k class corresponds to the interval $[q(p_{k-1}); q(p_k)[$,for $k = 1, \ldots, q-1$, where $q(p_0) = \min_X x$ and the last class gathers all data above $q(p_{q-1})$: $[q(p_{q-1}), +\infty[$.
- We call q-tiles a complete series of $k = 1, \ldots, q$ upper-closed q^k, resp. lower-closed q_k, quantile classes.

q-tiles sorting on a single criteria

If x is a measured performance, we may distinguish three sorting situations:

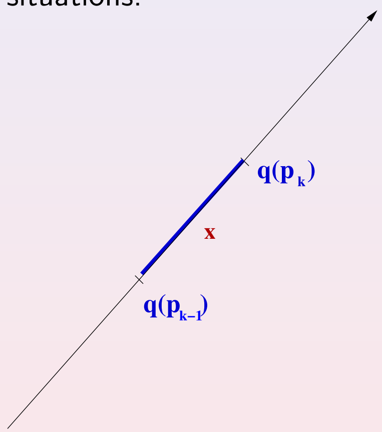

1. $x \leqslant q(p_{k-1})$ and $x < q(p_k)$
   The performance x is lower than the q^k class;
2. $x > q(p_{k-1})$ and $x \leqslant q(p_k)$
   The performance x belongs to the q^k class;
3. $(x > q(p_{k-1})$ and$)$
   $x > q(p_k)$
   The performance x is higher than the p^k class.

If the relation $<$ is the dual of $\geqslant$, it will be sufficient to check that both, $q(p_{k-1}) \not\geqslant x$, as well as $q(p_k) \geqslant x$, are verified for x to be a member of the k-th q-tiles class.

Multiple criteria extension

- $A = \{x, y, z, \ldots\}$ is a finite set of n objects to be sorted.
- $F = \{1, \ldots, m\}$ is a finite and coherent family of m performance criteria.
- Each criterion j in F carries a rational significance w_j such that $0 < w_j < 1.0$ and $\sum_{j \in F} w_j = 1.0$.
- For each criterion j in F, the objects are evaluated on a real performance scale $[0; M_j]$,

  supporting an indifference threshold ind_j,

  a preference threshold pr_j,

  and a veto threshold v_j,

  such that $0 \leqslant ind_j < pr_j < v_j \leqslant M_j$.
- The performance of object x on criterion j is denoted x_j.

Performing marginally "*at least as good as*"

Each criterion j is characterizing a double threshold order $\geqslant_j$ on A in the following way:

$$r(x \geqslant_j y) = \begin{cases} +1 & \text{if} \quad x_j - y_j \geqslant -ind_j \\ -1 & \text{if} \quad x_j - y_j \leqslant -pr_j \\ 0 & \text{otherwise.} \end{cases} \qquad (1)$$

$+1$ signifies x is *performing at least as good as* y on criterion j,

-1 signifies that x is *not performing at least as good as* y on criterion j.

0 signifies that it is *unclear* whether, on criterion j, x is performing at least as good as y.

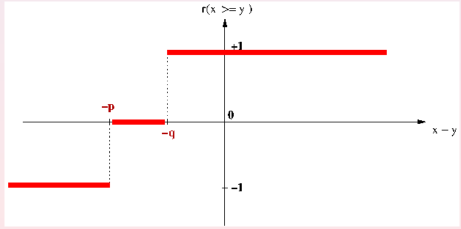

Performing globally "*at least as good as*"

Each criterion j contributes the significance w_j of his "*at least as good as*" characterization $r(\geqslant_j)$ to the global characterization $r(\geqslant)$ in the following way:

$$\boxed{r(x \geqslant y) \;=\; \sum_{j \in F} \left[\, w_j \cdot r(x \geqslant_j y) \,\right]} \qquad (2)$$

$r > 0$ signifies x is *globally performing at least as good as* y,

$r < 0$ signifies that x is *not globally performing at least as good as* y,

$r = 0$ signifies that it is *unclear* whether x is globally performing at least as good as y.

The bipolar outranking relation $\succsim$

From an epistemic point of view, we say that:

1. performance x *outranks* performance y, denoted $(x \succsim y)$, if
   1.1 a significant majority of criteria validates a global outranking situation between x and y, i.e. $(x \geqslant y)$ and
   1.2 no veto $(x \not\lll y)$ is observed on a discordant criterion,
2. performance x *does not outrank* performance y, denoted $(x \not\succsim y)$, if
   2.1 a significant majority of criteria invalidates a global outranking situation between x and y, i.e. $(x \not\geqslant y)$ and
   2.2 no counter-veto $(x \not\ggg y)$ observed on a concordant criterion.

Polarising the global "*at least as good as*" characteristic

The valued bipolar outranking characteristic $r(\succsim)$ is defined as follows:

$$r(x \succsim y) = \begin{cases} 0, & \text{if } \big[\exists j \in F : r(x \lll_j y)\big] \wedge \big[\exists k \in F : r(x \ggg_k y)\big] \\ \big[\, r(x \geqslant y) \circledvee -r(x \lll y) \,\big] & \text{,otherwise.} \end{cases}$$

And in particular,

- $r(x \succsim y) = r(x \geqslant y)$ if no very large positive or negative performance differences are observed,
- $r(x \succsim y) = 1$ if $r(x \geqslant y) \geqslant 0$ and $r(x \ggg y) = 1$,
- $r(x \succsim y) = -1$ if $r(x \geqslant y) \leqslant 0$ and $r(x \lll y) = 1$,

q-tiles sorting with bipolar outrankings

Proposition

The bipolar characteristic of x belonging to upper-closed q-tiles class q^k, resp. lower-closed class q_k, may hence, in a multiple criteria outranking approach, be assessed as follows:

$$\begin{aligned} r(x \in q^k) &= \min\big[\, -r\big(\mathbf{q}(p_{k-1}) \succsim x\big),\; r\big(\mathbf{q}(p_k) \succsim x\big)\,\big] \\ r(x \in q_k) &= \min\big[\, r\big(x \succsim \mathbf{q}(p_{k-1})\big),\; -r\big(x \succsim \mathbf{q}(p_k)\big)\,\big] \end{aligned}$$

Proof.

The bipolar outranking relation $\succsim$, being weakly complete, verifies actually the coduality principle. The dual ($\not\succsim$) of $\succsim$ is, hence, identical to the strict converse outranking $\precnsim$ relation. □

R. Bisdorff (2013), *On polarizing outranking relations with large performance differences*. Journal of Multi-Criteria Decision Analysis **20**:3-12

The multicriteria (upper-closed) q-tiles sorting algorithm

1. **Input**: a set X of n objects with a performance table on a family of m criteria and a set $\mathcal{Q}$ of $k = 1, .., q$ empty q-tiles equivalence classes.
2. **For each** object $x \in X$ **and each** q-tiles class $q^k \in \mathcal{Q}$:
   2.1 $r(x \in q^k) \quad \leftarrow \quad \min\big(\, -r(\mathbf{q}(p_{k-1}) \succsim x), r(\mathbf{q}(p_k) \succsim x)\,\big)$
   2.2 if $r(x \in q^k) \;\geqslant\; 0$:
   add x to q-tiles class q^k;
3. **Output**: $\mathcal{Q}$.

Comment

1. *The complexity of the q-tiles sorting algorithm is $\mathcal{O}(nmq)$; linear in the number of decision actions (n), criteria (m) and quantile classes (q).*
2. *As $\mathcal{Q}$ represents a partition of the criteria measurement scales, i.e. the upper limits of the preceding category correspond to the lower limits of the succeeding ones, there is a potential for run time optimization.*

Example of quintiles sorting result

```
>>> from randomPerfTabs import RandomPerformaceTableau
>>> t = RandomPerformanceTableau(numberOfActons=50,
  ...                                      seed=5)
>>> from sparseOutrankingDigraphs import\
  ...                   PreRankedOutrankingDigraph
>>> pr = PreRankedOutrankingDigraph(t,qintiles)
>>> pr.showSorting()
*--- Sorting results in descending order ---*
]0.8 - 1.0]: [a16, a2, a24, a32]
]0.6 - 0.8]: [a01, a02, a06, a09, a10, a13, a16, a18,
              a22, a25, a27, a28, a31, a32, a36, a37,
              a39, a40, a41, a43, a45, a48]
]0.4 - 0.6]: [a01, a03, a04, a05, a07, a08, a09, a10,
              a11, a12, a13, a14, a15, a17, a18, a20,
              a26, a27, a29, a30, a33, a34, a35, a38,
              a42, a43, a44, a45, a46, a47, a49, a50]
]0.2 - 0.4]: [a04, a11, a12, a17, a19, a21, a23,
              a29, a34, a42, a46, a47, a50]
] < - 0.20]: []
```

Example of quintiles decomposing result

```
>>> pr.showDecomposition()
*--- quantiles decomposition in decreasing order---*
     c1. ]0.8-1.0]: [a24]
     c2. ]0.6-1.0]: [a16,a22, a32]
     c3. ]0.6-0.8]: [a02, a06, a25, a28, a31, a36, a37,
                     a39,  a40, a41, a48]
     c4. ]0.4-0.8]: [a01, a09, a10, a13, a18,
                     a27, a43, a45]
     c5. ]0.4-0.6]: [a03, a05, a07, a08, a14, a15, a20,
                     a26, a30, a33, a35, a38, a44, a49']
     c6. ]0.2-0.6]: [a04,'a11, a12, a17, a29, a34, a42,
                     a46, a47, a50]
     c7. ]0.2-0.4]: [a19, a21, a23]
```


Sparse versus standard outranking digraph of order 50

Symbol legend

- ⊤ outranking for certain
- \+ more or less outranking
- ' ' indeterminate
- − more or less outranked
- ⊥ outranked for certain

Sparse digraph *bg*:
\# Actions : 50
\# Criteria : 7
Sorted by : 5-Tiling
Ranking rule : Copeland
\# Components : 7
Minimal order : 1
Maximal order : 15
Average order : 7.1
fill rate : 20.980%
correlation : +0.7563

Content

1. Pre-ranked sparse outranking digraphs
   - How to rank a big performance tableau ?
   - quantiles sorting of a performance tableau
   - Multiple criteria quantiles sorting
2. Ranking a q-tiled performance tableau
   - Properties of the q-tiles sorting
   - Ordering the q-tiles sorting result
   - q-tiles ranking algorithm
3. HPC-ranking a big performance tableau
   - Multithreading the sorting&ranking procedures
   - Profiling the HPC sorting&ranking procedures

Properties of q-tiles sorting result

1. *Coherence*: Each object is always sorted into a non-empty subset of adjacent q-tiles classes.
2. *Uniqueness*: If $r(x \in q^k) \neq 0$ for $k = 1, ..., q$, then performance x is sorted into exactly one single$q$-tiles class.
3. *Separability*: Computing the sorting result for performance x is independent from the computing of the other performances' sorting results.

Comment

The separability property gives us access to efficient parallel processing of class membership characteristics $r(x \in q^k)$ for all $x \in X$ and q^k in $\mathcal{Q}$.

Ordering the q-tiles sorting result

The q-tiles sorting result leaves us with a more or less refined partition of the set X of n potential decision actions. For linearly ranking from best to worst the resulting parts of the q-tiles partition we may apply three strategies:

1. Optimistic: In decreasing lexicographic order of the upper and lower quantile class limits;
2. Pessimistic: In decreasing lexicographic order of the lower and upper quantile class limits;
3. Average: In decreasing numeric order of the average of the lower and upper quantile limits.

q-tiles ranking algorithm

1. **Input**: the outranking digraph $\mathcal{G}(X, \succsim)$, a partition P_q of k linearly ordered decreasing parts of X obtained by the q-sorting algorithm, and an empty list $\mathcal{R}$.
2. **For each** quantile class $q^k \in P_q$:
   if $\#(q^k) > 1$:
   $R_k \leftarrow$ **locally rank** q^k in $\mathcal{G}_{|q^k}$
   (if ties, render alphabetic order of action keys)
   else: $R_k \leftarrow q^k$
   append R_k to $\mathcal{R}$
3. **Output**: $\mathcal{R}$

q-tiles ranking algorithm – Comments

1. The complexity of the q-tiles ranking algorithm is linear in the number k of components resulting from a q-tiles sorting which contain more than one action.
2. Three local ranking rules are available – *Copeland*'s, *Net-flows*' and *Kohler*'s rule – of complexity $\mathcal{O}\big((\#q^k)^2\big)$ on each restricted outranking digraph $\mathcal{G}_{|q^k}$.
3. In case of local ties (very similar evaluations for instance), the **local ranking** procedure will render these actions in increasing alphabetic ordering of the action keys.

Content

Multithreading the q-tiles sorting & ranking procedures

1. Following from the separability property of the q-**tiles sorting** of each action into each q-tiles class, the q-sorting algorithm may be safely split into as much threads as are multiple processing cores available in parallel.
2. Furthermore, the **ranking** procedure being local to each diagonal component, these procedures may as well be safely processed in parallel threads on each restricted outranking digraph $\mathcal{G}_{|q^k}$.

HPC performance measurements

digraph order	standard model #c.	t_g sec.	τ_g	sparse model #c.	t_{bg}	τ_{bg}
1 000	118	6"	+0.88	8	1.6'	+0.83
2 000	118	15"	+0.88	8	3.5"	+0.83
2 500	118	27"	+0.88	8	4.4"	+0.83
10 000				118	7"	
15 000				118	12"	
25 000				118	21"	
50 000				118	48"	
100 000	(size	=	10^{10})	118	2'	(fill rate = 0.077%)
1 000 000	(size	=	10^{12})	118	36'	(fill rate = 0.028%)
1 732 051	(size	=	3×10^{12})	118	2h17'	(fill rate = 0.010%)
2 236 068	(size	=	5×10^{12})	118	3h15'	(fill rate = 0.010%)

Legend:

- #c. = number of cores;
- g: standard outranking digraph, bg: the sparse outranking digraph;
- t_g, resp. t_{bg}, are the corresponding constructor run times;
- τ_g, resp. τ_{bg} are the ordinal correlation of the Copeland ordering with the given outranking relation.

Fitness of local ranking rules

Sample of 100 random outranking graphs of order 250

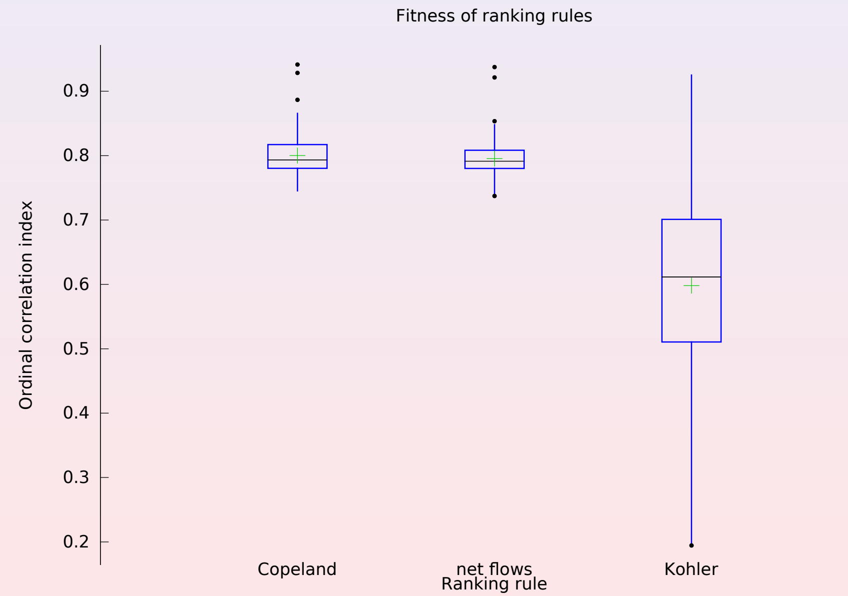

Standard versus 50-tiled sparse outranking digraphs

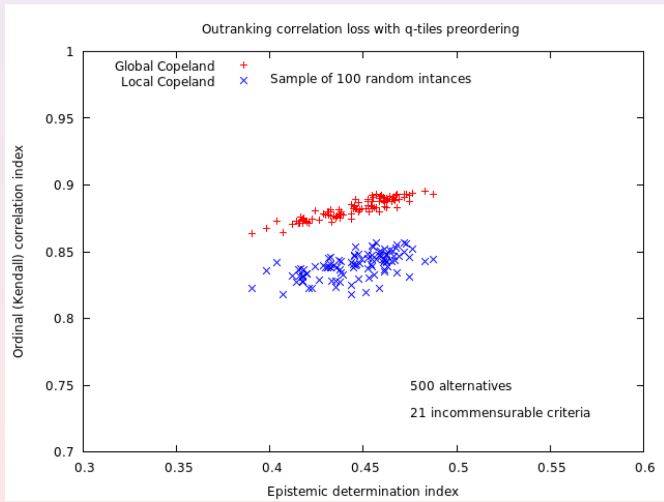

Profiling the local ranking procedure

It is opportune to use Copeland's rule for ranking from the standard outranking digraph, whereas both, Net Flows and Copeland's ranking rule, are equally efficient on the sparse outranking digraph.

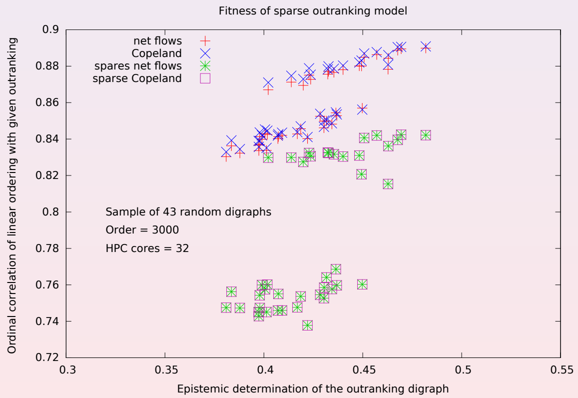

The quality of the sparse model based linear ordering is depending on the alignment of the given outranking digraph, but not on its actual order.

Concluding ...

- We implement a sparse outranking digraph model coupled with a linearly ordering algorithm based on quantiles-sorting & local-ranking procedures;
- Global ranking result fits apparently well with the given outranking relation;
- Independent sorting and local ranking procedures allow effective multiprocessing strategies;
- Efficient scalability allows hence the linear ranking of very large sets of potential decision actions (millions of nodes) graded on multiple incommensurable criteria;
- Good perspectives for further optimization with cPython and HPC ad hoc tuning.

Python and cython HPC modules available under:
http://github.com/rbisdorff/Digraph3

Documentation: http://charles-sanders-peirce.uni.lu/docDigraph3/