A brief introduction to algorithmic decision theory

Olivier Spanjaard
olivier.spanjaard@lip6.fr
Sorbonne Université

UTC talk - april 2018

2018/04/12

Decision theory

Genesis

Forties and fifties : study of mathematical models able to account for the observed choice behaviors of rational agents (descriptive approach), or to suggest the best decision to made (normative approach).

Risk and uncertainty

Proposition of models
Study of their properties

Social choice

Voting rules
Impossibility theo.

Game theory

Non-cooperative games
Cooperative games

Decision theory and algorithms

Need for algorithmic tools to implement the decision models proposed by economists and mathematicians.

Algorithmic decision theory

Revisiting the work done in decision theory, which focused mainly on the existence (or not) of models satisfying certain conditions, with emphasis on the computational aspects.

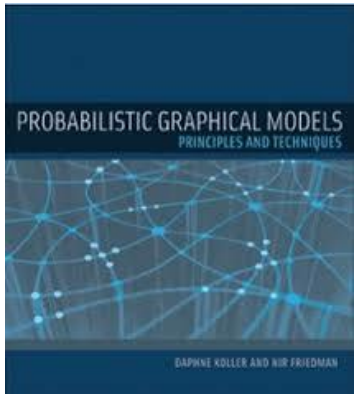

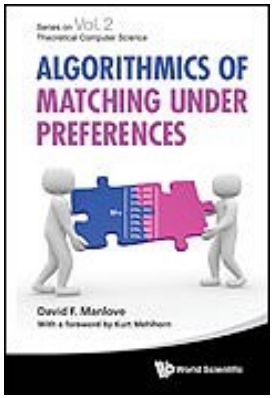

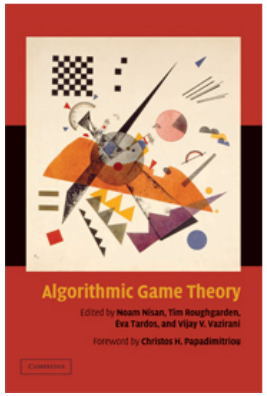

We will illustrate on a few examples

A first example: Who wants to be a millionaire?

Sequential decision under uncertainty

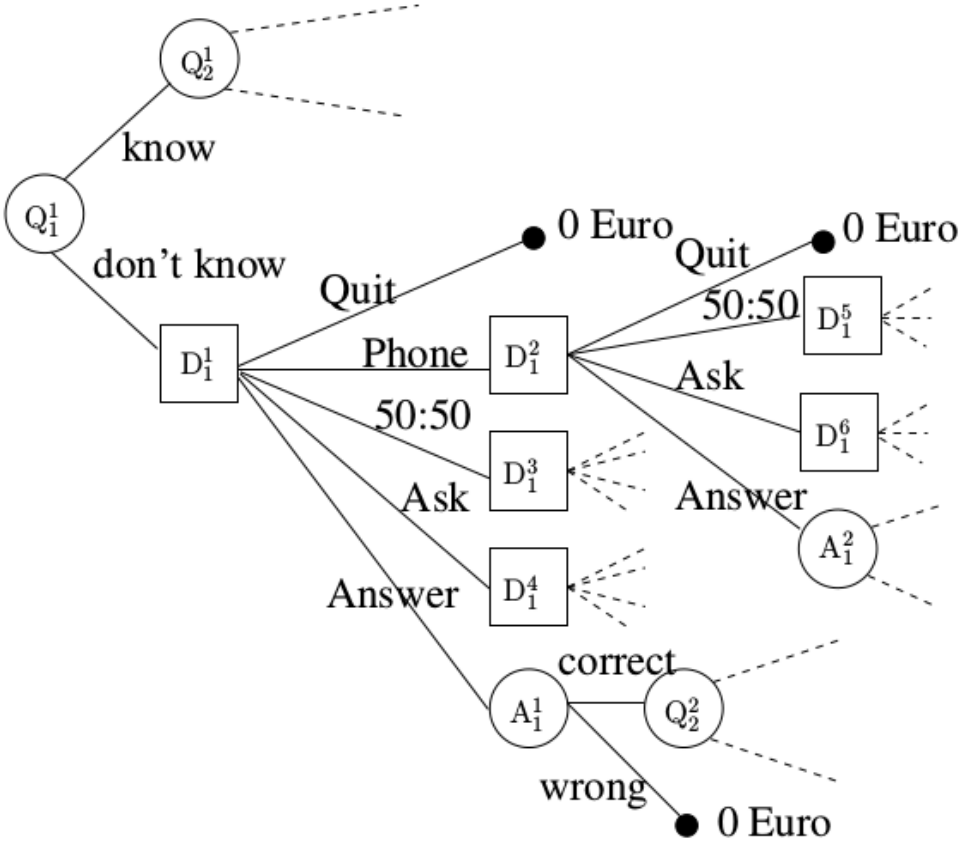

75 millions of nodes

A strategy

The objective is to determine a preferred strategy from the root

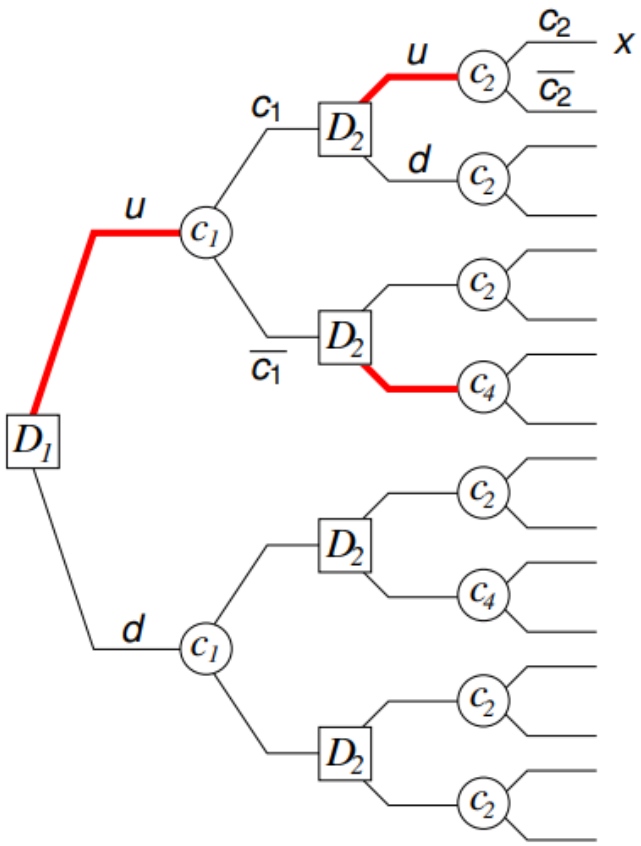

$$x = \text{outcome if } D_1 = u, C_1 = c_1, D_2 = u, C_2 = c_2$$

Evaluating a strategy

Every strategy induces a lottery over the outcomes

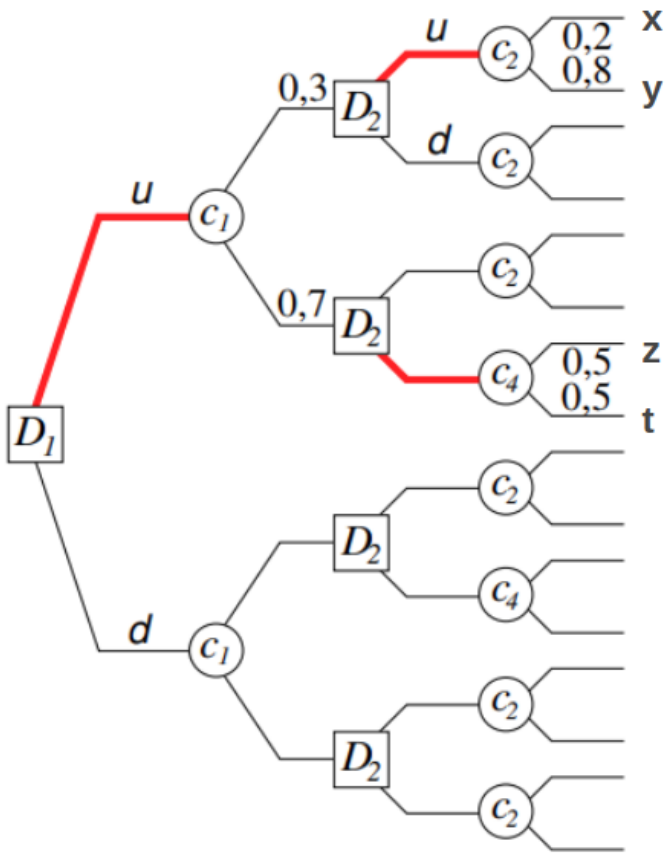

Lottery $\langle x, 0.06; y, 0.24; z, 0.35; t, 0.35 \rangle$

The economists' approach

Hammond[1] characterized dynamic choice behaviors for which there exists a utility function on the outcomes (at the leaves of the tree) such that a best strategy is one optimizing the expected utility.

THEOREM 7:

(1) Let β be consequentialist behaviour satisfying dynamic consistency and continuity condition (CB) for the almost unrestricted domain $\mathcal{T}(Y)$ of finite decision trees with only positive probabilities at all chance nodes. Then there exists a unique cardinal equivalence class of NMUFs $v : Y \to \mathbb{R}$ such that β maximizes expected utility.

(2) Conversely, let $v : Y \to \mathbb{R}$ be any NMUF. Then there exists consequentialist behaviour β satisfying dynamic consistency and condition (CB) on the domain $\mathcal{T}(Y)$ with the property that the implicit preference ordering R_β is represented by the expected value of v.

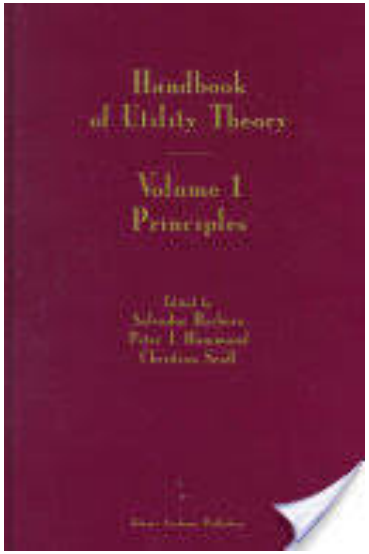

Consequentialism + Dynamic consistency

[1] Peter J Hammond. "Consequentialist foundations for expected utility". In: *Theory and decision* 25.1 (1988), pp. 25–78.

Consequentialism

Definition

The preferences in the subtree rooted in B do not depend on the remainder of the tree.

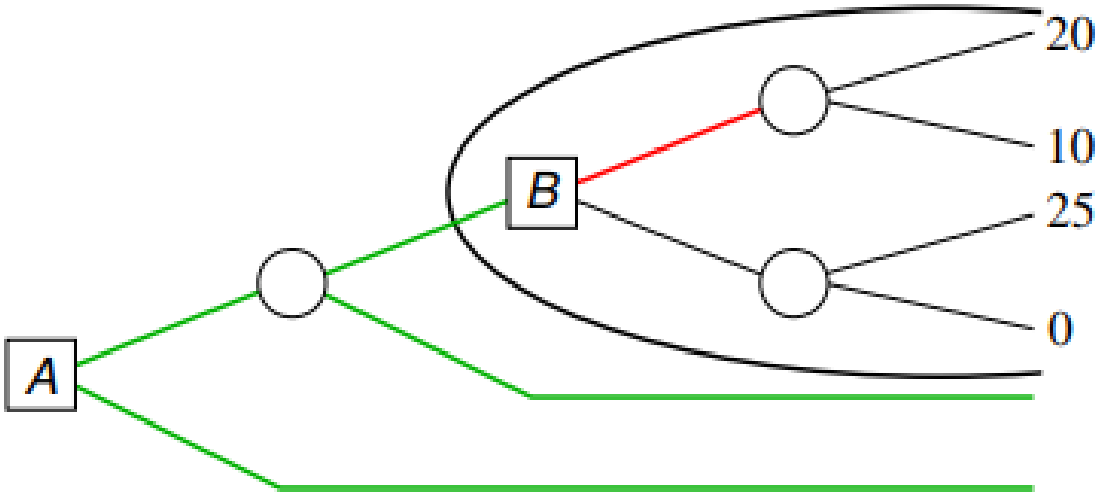

Dynamic consistency

Definition

The preferred strategy from A generates a substrategy from B that is the preferred strategy in the subtree rooted in B.

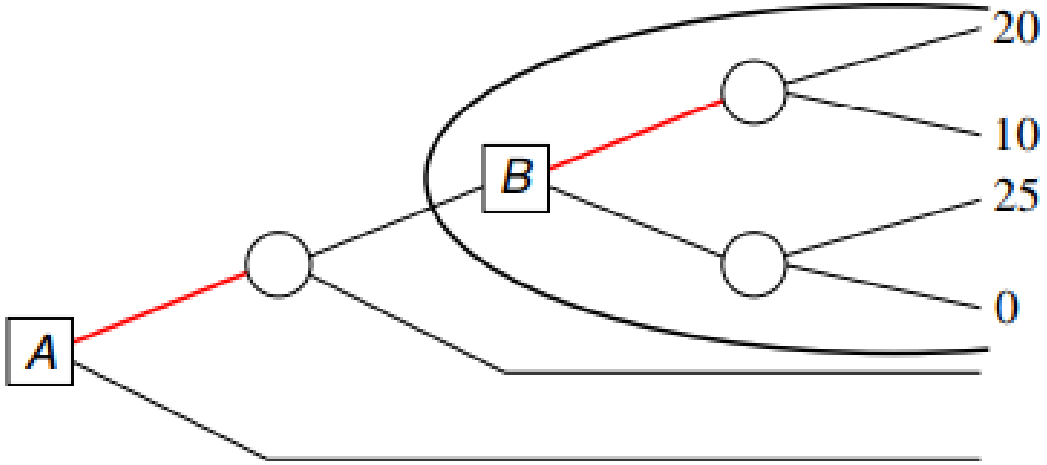

The computer scientists' approach

How to compute an optimal strategy?

Exponential number of strategies
⇒ Impractical to use an enumerate & evaluate method!

Necessity to design implicit enumeration algorithms
⇒ An EU-optimal strategy can be determined in polynomial time by dynamic programming.

This approach goes back at least to the 60s[2], therefore nothing really new...

[2] Howard Raiffa. *Decision Analysis: Introductory Lectures on Choices und Uncertainty*. Addison-Wesley, 1968.

A lot still remains to be done

Allais' paradox[3]

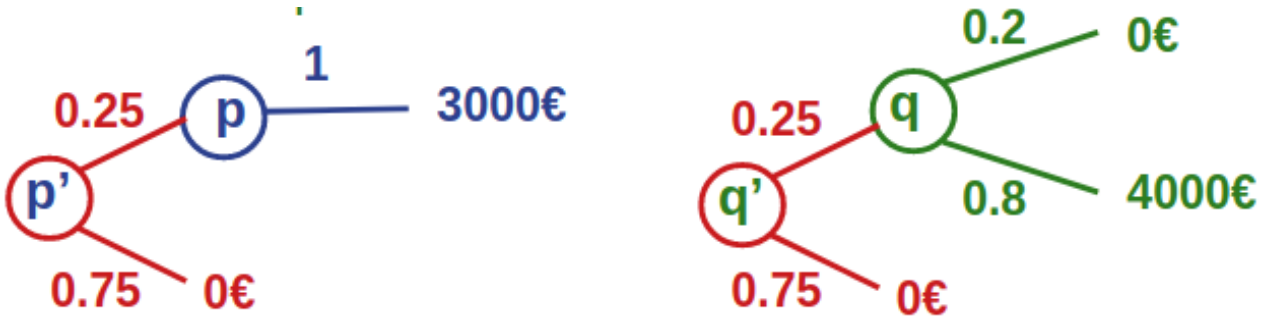

Most DMs prefer **p** to **q** and **q**′ to **p**′, which is not compatible with EU.

Other decision criteria under risk: RDU, SSB, WEU...

Other representations: Markov decision processes, influence diagrams...

[3] Maurice Allais. "Le comportement de l'homme rationnel devant le risque : Critique des postulats et axiomes de l'ecole americaine". In: *Econometrica* 21.4 (1953), pp. 503–546.

Some works on risk sensitive planning

RDU in decision trees

- Gildas Jeantet and Olivier Spanjaard. "Rank-dependent Probability Weighting in Sequential Decision Problems under Uncertainty". In: *International Conference on Automated Planning and Scheduling*. Sydney, Australia: AAAI Press, 2008, pp. 148–155
- Gildas Jeantet, Patrice Perny, and Olivier Spanjaard. "Sequential Decision Making with Rank Dependent Utility: a Minimax Regret Approach". In: *26th AAAI Conference on Artificial Intelligence*. Toronto, Canada, 2012, pp. 1931–1937

SSB and WEU in decision trees

- Hugo Gilbert and Olivier Spanjaard. "Complexity of Solving Decision Trees with Skew-Symmetric Bilinear Utility". In: *Conference on Uncertainty in Artificial Intelligence (UAI-2017)*. Sydney, Australia, 2017

EU in Markov decision processes

- Y. Liu and S. Koenig. "An Exact Algorithm for Solving MDPs Under Risk-sensitive Planning Objectives with One-switch Utility Functions". In: *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 1*. AAMAS '08. 2008, pp. 453–460
- Y. Liu. "Risk-sensitive planning with one-switch utility functions: Value iteration". In: *In AAAI*. 2005, pp. 993–999
- Y. Liu and S. Koenig. "Functional Value Iteration for Decision-Theoretic Planning with General Utility Functions". In: *AAAI*. AAAI Press, 2006, pp. 1186–1193
- Y. Liu and S. Koenig. "Existence and Finiteness Conditions for Risk-Sensitive Planning: Results and Conjectures". In: *In UAI*. 2005, pp. 354–363

SSB in Markov decision processes

- Hugo Gilbert, Olivier Spanjaard, et al. "Solving MDPs with Skew Symmetric Bilinear Utility Functions". In: *24th International Joint Conference on Artificial Intelligence (IJCAI-15)*. Buenos Aires, Argentina, 2015, pp. 1989–1995

Detecting structured preference profiles

Social choice theory

$$\begin{array}{ccccc} \begin{bmatrix}8\\7\\6\\5\\4\\3\\2\\1\end{bmatrix} & \begin{bmatrix}8\\2\\6\\7\\1\\3\\5\\4\end{bmatrix} & \begin{bmatrix}2\\5\\6\\7\\4\\3\\1\\8\end{bmatrix} & \begin{bmatrix}1\\3\\6\\5\\4\\8\\2\\7\end{bmatrix} & \begin{bmatrix}4\\8\\2\\1\\6\\7\\5\\3\end{bmatrix} \\ V_1 & V_2 & V_3 & V_4 & V_5 \end{array}$$

Aggregation procedure:
Collective ranking

$$\begin{bmatrix}6\\5\\2\\8\\7\\4\\1\\3\end{bmatrix}$$

The economists' approach

Arrow's theorem

For three or more candidates, there does not exist an aggregation rule that satisfy universality, non-dictatorship, unanimity and independence of irrelevant alternatives.

Kenneth J. Arrow. *Social Choice and Individual Values*. Yale University Press, 1951, 2nd ed., 1963

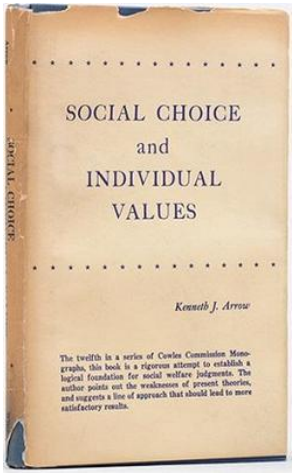

Structured preferences make it possible to overcome this impossibility result.

The most well-known structure: single-peaked preferences[4].

[4] Duncan Black. "On the rationale of group decision-making". In: *Journal of political economy* 56.1 (1958), pp. 23–34.

Single-peaked preferences

PRÉSIDENTIELLE 2017

PP	JLM	BH	EM	FF	NDA	MLP
NPA	FI	PS	EM	LR	DLF	FN

Single-peaked preferences

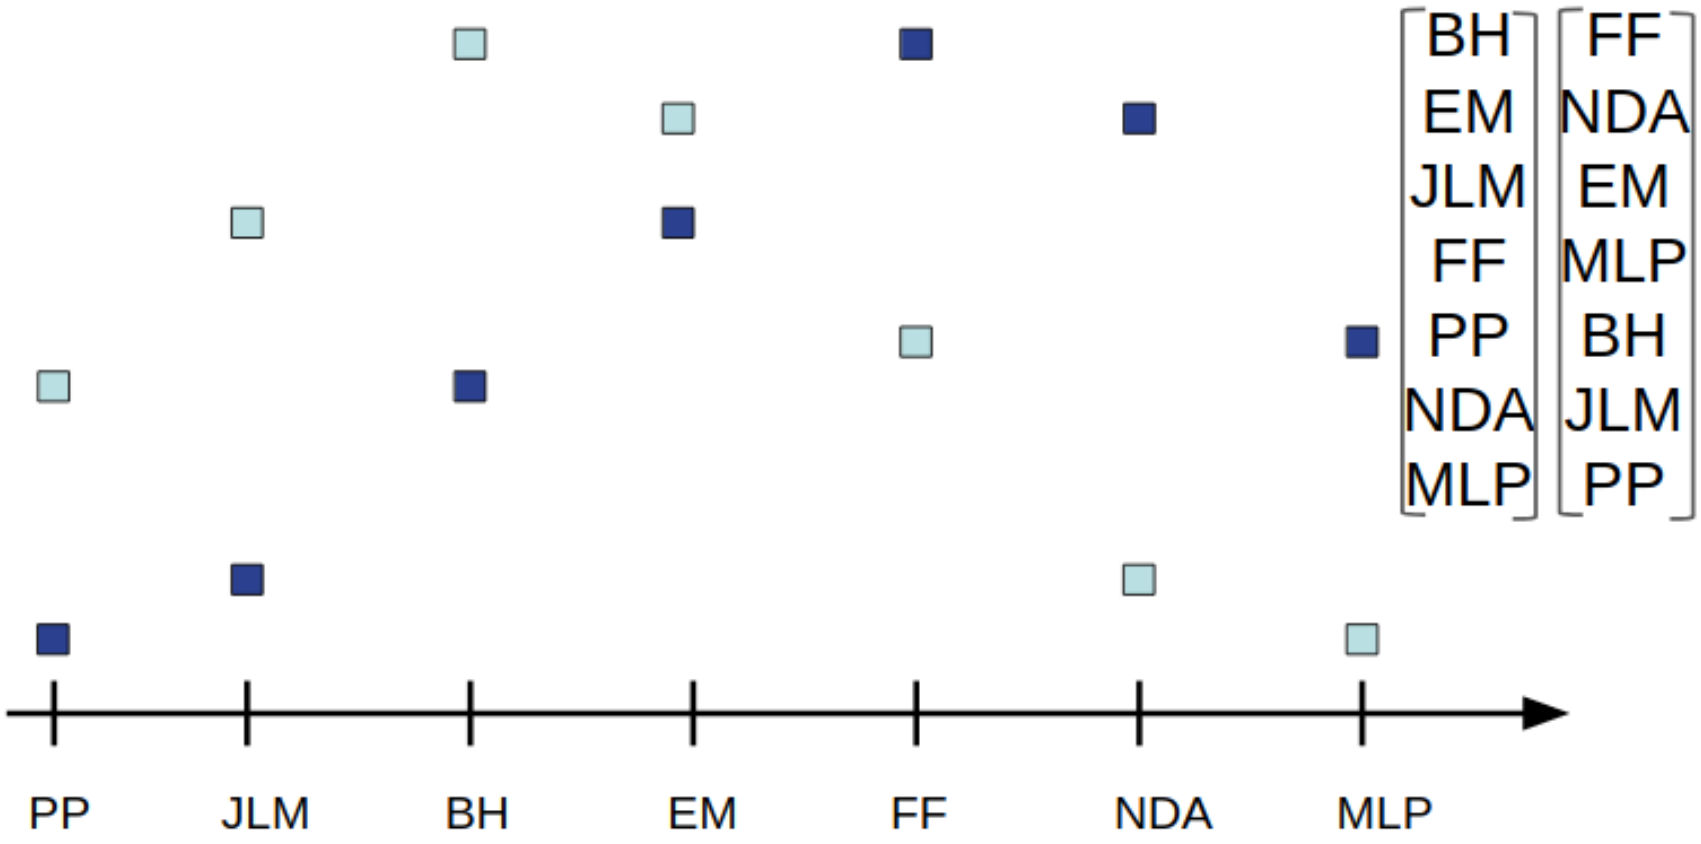

The computer scientists' approach

Given a profile, is it hard to detect (nearly) single-peaked preferences, or other structured preferences?

- Robert Bredereck, Jiehua Chen, and Gerhard J Woeginger. "Are there any nicely structured preference profiles nearby?" In: *Mathematical Social Sciences* 79 (2016), pp. 61–73
- Denis Cornaz, Lucie Galand, and Olivier Spanjaard. "Bounded Single-Peaked Width and Proportional Representation." In: *ECAI*. vol. 12. 2012, pp. 270–275
- Edith Elkind and Piotr Faliszewski. "Recognizing 1-euclidean preferences: An alternative approach". In: *International Symposium on Algorithmic Game Theory*. 2014, pp. 146–157
- Edith Elkind and Martin Lackner. "On Detecting Nearly Structured Preference Profiles." In: *AAAI*. vol. 14. 2014, pp. 661–667
- Bruno Escoffier, Jérôme Lang, and Meltem Öztürk. "Single-peaked consistency and its complexity." In: *ECAI*. vol. 8. 2008, pp. 366–370

Given a voting problem, can we take advantage of (near) single-peakedness, or of other structured preferences, to speed up the solving of the problem?

- Yongjie Yang and Jiong Guo. "Controlling elections with bounded single-peaked width". In: *Proceedings of AAMAS 2014*. 2014, pp. 629–636
- Piotr Skowron et al. "The complexity of fully proportional representation for single-crossing electorates". In: *Theoretical Computer Science* 569 (2015), pp. 43–57
- Dominik Peters. "Single-Peakedness and Total Unimodularity: New Polynomial-Time Algorithms for Multi-Winner Elections". In: *AAAI 2018*. 2018

The stable roommates problem

The stable roommates problem is to match each of $2n$ people so that no two people prefer each other to their assigned partners (stable matching).

Consider 8 people with single-peaked and narcissistic preferences:

1: $1 \succ 2 \succ 3 \succ 4 \succ 5 \succ 6 \succ 7 \succ 8$
2: $2 \succ 3 \succ 4 \succ 5 \succ 1 \succ 6 \succ 7 \succ 8$
3: $3 \succ 4 \succ 2 \succ 5 \succ 1 \succ 6 \succ 7 \succ 8$
4: $4 \succ 3 \succ 5 \succ 6 \succ 2 \succ 7 \succ 1 \succ 8$
5: $5 \succ 4 \succ 3 \succ 2 \succ 6 \succ 7 \succ 8 \succ 1$
6: $6 \succ 5 \succ 4 \succ 3 \succ 2 \succ 7 \succ 1 \succ 8$
7: $7 \succ 6 \succ 5 \succ 4 \succ 3 \succ 2 \succ 8 \succ 1$
8: $8 \succ 7 \succ 6 \succ 5 \succ 4 \succ 3 \succ 2 \succ 1$

A stable matching: $\{3,4\}$, $\{2,5\}$, $\{6,7\}$, $\{1,8\}$.

If preferences are narcissistic and single-peaked, a stable matching always exists, and it can be computed in $O(n)$.

John Bartholdi III and Michael A Trick. "Stable matching with preferences derived from a psychological model". In: *Operations Research Letters* 5.4 (1986), pp. 165–169

For the record... (as noted by Ch. Papadimitriou)

1944

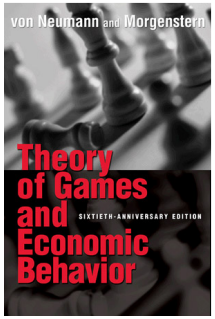

J. von Neumann

1946

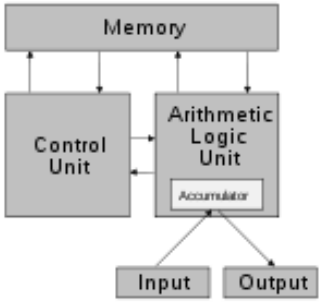

Algorithmic decision theory

ADT 2017

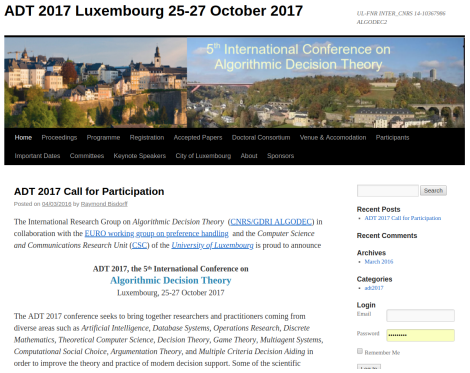

COMSOC 2018

SAGT 2018

Algorithms and explainability

La transparence des algorithmes en question

Les décisions des machines auto-apprenantes sont de plus en plus opaques, même pour leurs concepteurs.

LE MONDE | 09.04.2018 à 08h00 |

Par Denis Delbecq

Abonnez vous à partir de 1 € | Réagir | Ajouter | Partager (74) | Tweeter

La sélection qui ne dit pas son nom

Les enseignants décryptent l'algorithme "ubuesque" de Parcoursup

Par Alexandra Saviana

Publié le 09/04/2018 à 18:07

Macron a promis la publication de l'algorithme national de Parcoursup. Mais l'avenir des étudiants se jouera plutôt dans la pléthore de variantes locales, estampillées "outils d'aide à la décision". "Opaques" pour les uns, "incompréhensibles", pour les autres, ces algorithmes locaux sont "un vrai moyen de sélection" à l'entrée de l'université.

Comprendre le fonctionnement des algorithmes est "essentiel"

Par Le Figaro.fr avec AFP | Mis à jour le 29/03/2018 à 12:46 / Publié le 29/03/2018 à 12:46

LE FIGARO PREMIUM › 1€ le premier mois | Commentez

Comprendre le fonctionnement des algorithmes, qui déterminent par exemple les publicités ou les contenus sur internet, est "essentiel", a estimé jeudi le secrétaire d'Etat au numérique Mounir Mahjoubi lors du colloque "Intelligence artificielle pour l'humanité". La loi prévoit déjà pour les acteurs numériques une "obligation" de "rendre transparents les critères qui ont permis la présentation d'une information et d'expliquer comment l'ordre a été établi", a-t-il rappelé.

Société / Tech & internet

Comment évaluer l'éthique d'un algorithme?

Slate.fr — 29 mars 2018 à 11h09 — mis à jour le 29 mars 2018 à 16h47

Au lendemain de le publication du rapport Villani sur l'intelligence artificielle, il est temps de s'interroger sur l'éthique des algorithmes.

35

Choices based on decision-theoretic criteria are easier to explain.

THANK YOU!