Preferences

Toby Walsh
NICTA and University of NSW
Sydney, Australia

AAAI 07, Vancouver

Acknowledgements

- My co-authors
  - Carmel Domshlak, Jerome Lang, Maria Silvia Pini, Steve Prestwich, Francesca Rossi, Brent Venable ..
- Many colleagues
  - Mike Trick, Vincent Conitzer, Tuomas Sandholm, Craig Boutilier, Ronen Brafman, David Poole, Holger Hoos, ...

Apology

- Work in progress
  - But this is what Adele and Rob wanted as opposed to finished and polished work
- Many open questions
  - AAAI workshop, Preference Handling
  - AAAI tutorial, Brafman & Domshlak

Let’s start!

Why preferences?

- Over constrained problems
  - How do we choose between solutions?
- Multiple agents
  - How do we deal with their conflicting desires?
- ...

Running questions

- How do we combine together preferences?
  - Anne prefers Thai food, Bob prefers Indian, ..
- How do we reason about incomplete preferences?
  - Anne's preferences are only partially known

Voting

- Social choice's method to combine preferences
  - Run an election!
  - Anne, Bob & Carol rank the cuisines
  - Use a voting rule (e.g. plurality or STV) to compute "winner"

An AI perspective!

- Elections typically have a few candidates (except in Italy!)
- Preferences can be over large domains:
  - All restaurants in Vancouver's yellow pages
  - All songs in iTunes

An AI perspective!

- Computational perspective
  - How do we compute if we have elicited enough preferences to declare the winner?
  - Can we prevent strategic voting by making it computationally intractable?
  - ..

An AI perspective!

- Preferences & constraints
  - I prefer a cheap car
  - I prefer a Ferrari
  - But there are no cheap Ferraris!

So what are preferences and how do we represent them?

Quantitative preferences

- Thai food = 0.8, English = 0.1
  - But what do the numbers mean?
  - How do we combine them?
  - What about conditional preferences (e.g. if meal is expensive ..)?

Qualitative preferences

- Anne prefers Thai to English food
- Binary preference relation:
  - thai > english
- Transitive
  - indian > thai and thai > english then indian > english

What's a preference?

- Three "I"s
  - Indifference: thai ≥ indonesian and indonesian ≥ thai
  - Incompleteness: thai ? italian

    [Konczak, Lang, 05]
  - Incomparability: cheap indian ☢ fancy thai

    [Pini, Rossi, Venable, self TARK05, ECAI06, IJCAI07]

Preference domain

- AI (unlike social choice) faces large domains
  - Lunch domain:
    - cuisine x cost x distance x noise-level x ...
    - thai, expensive, near, noisy, ..
      indian, cheap, distance,quiet,..

CP-nets

- Decompose complex preference relation into conditionally independent parts
  - Much like Bayes nets for a complex probability function
- Ceteris paribus
  - “All else being equal”

[Boutilier, Brafman, Hoos, Poole, UAI99]

CP-nets

- CP statements
  - italian > french
  - italian: cheap > expensive
    french: expensive > cheap
- Directed dependency graph
  - Cyclic or acyclic?

CP-nets

- Various interesting extensions
  - Tradeoffs + CP-nets: “price is more important than weight”
    [Brafman, Domshlak UAI02]
  - Constraints + CP-nets
    [Prestwich, Rossi, Venable, self AAAI05]
  - Multiple agents: mCP-nets
    [Rossi, Venable, self AAAI04]

CP-nets

- Unfortunately dominance testing in CP-nets is computationally hard
  - PSPACE-complete

[Goldsmith, Lang, Truszczynski, Wilson IJCAI05]

- Various approximations proposed
  - E.g. where optimality is linear, dominance testing is NP-complete

[Prestwich, Rossi, Venable, self AAAI2005]

Approximating constrained CP-net

- A > B iff exist flipping sequence of improving flips from B to A
  - Each outcome in chain feasible
- Turn into set of hard constraints, *opt*(P)*

[Prestwich, Rossi, Venable, self AAAI2005]

We’ve said a little about representing preferences.

How do we combine them?

Combining preferences

- Use voting
- But what is a “good” voting rule
  - Condorcet’s paradox
  - Arrow’s impossibility theorem
  - Gibbard-Sattertwhaite theorem

Condorcet’s paradox

- Who should win?
  - Voter1: A>B>C
  - Voter2: B>C>A
  - Voter3: C>A>B
  - Majority prefer A to B, C to A, B to C

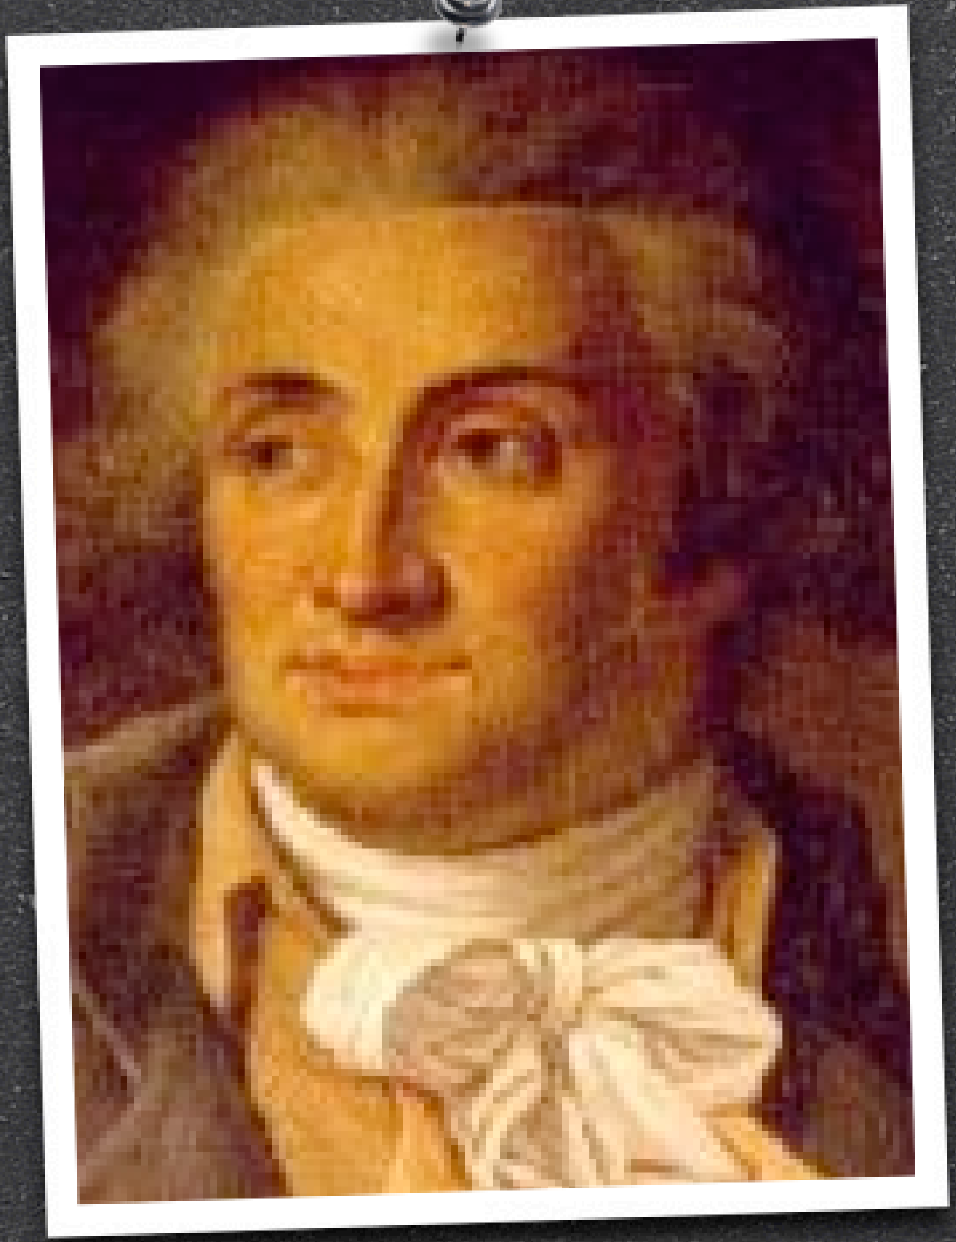

[Marquis de Condorcet 1785]

Arrow's theorem

- Impossible for a voting rule to be "fair"
  - 3 or more candidates
  - Rule is monotonic and independent to irrelevant alternatives
  - Then the rule is dictatorial

[Kenneth Arrow 1951]

Gibbard Satterthwaite

- All voting rules are “manipulable”
  - 3 or more candidates
  - Voting rule is onto (everyone can possibly win) but not dictatorial
  - Then you may need to vote tactically to get the result you want

Manipulation

- Generally considered a “bad” thing
  - Not transparent to electorate
  - Need sophisticated and informed voters
  - Result hard to predict

Avoiding manipulation

- Most voting rules are manipulable
- So use one where it is NP-hard to work out the manipulation

[Bartholdi, Tovey, Trick 89]

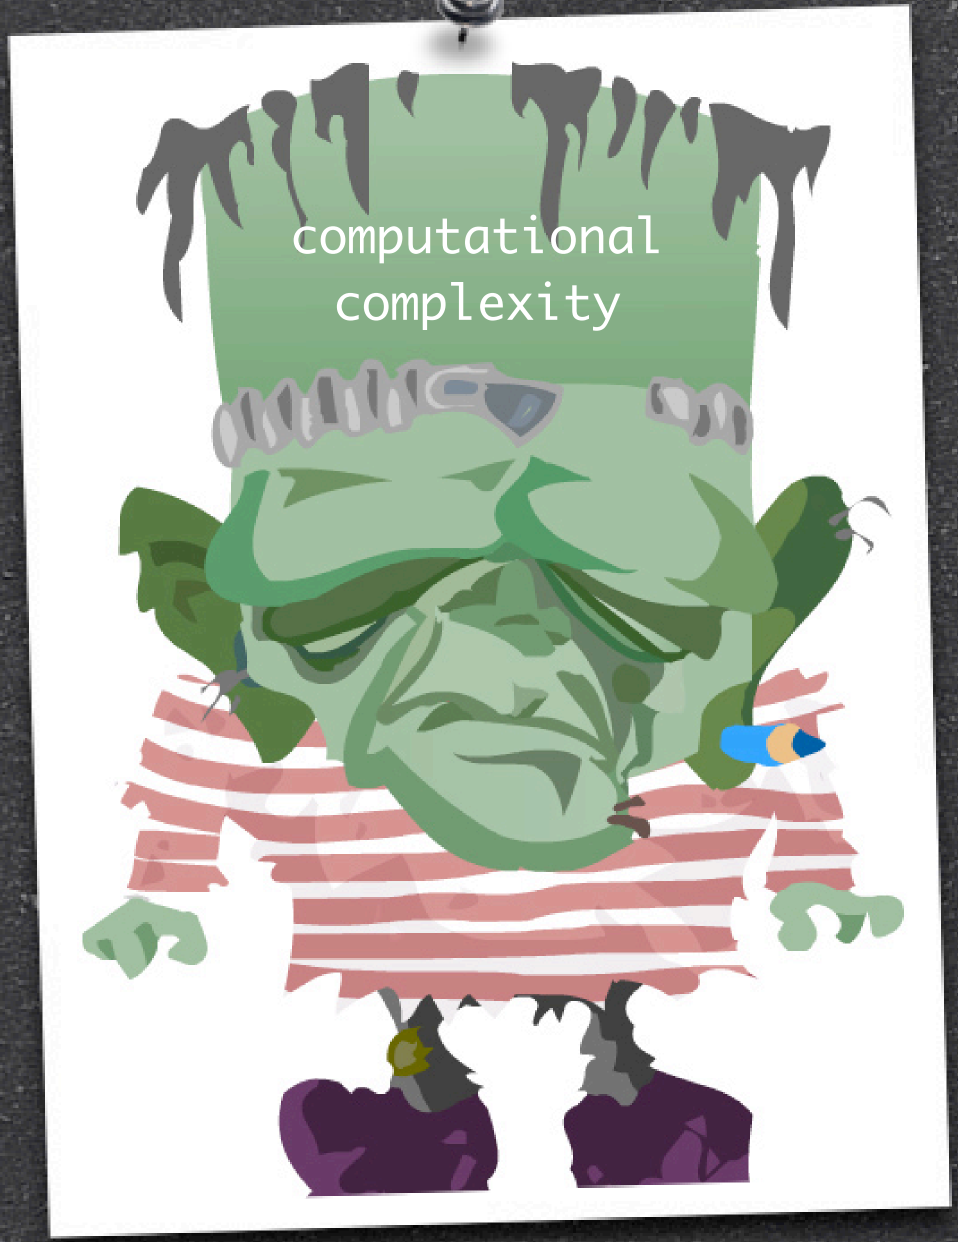

Some manipulable voting rules

- Scoring rules
  - Weight vector: $(\alpha_1,\alpha_2,..,\alpha_n)$
  - If voter ranks candidate in ith place, they receive score of α_i
  - Candidate with highest score wins

Some manipulable voting rules

- Scoring rules
  - Plurality has weight vector (1,0,..,0)
    - With 2 candidates, majority
  - Veto has weight vector (1,..,1,0)

Some manipulable voting rules

- Cup (aka “tournament”)
  - Knockout tournament of pairwise majority elections
- Single transferable vote (STV)
  - Eliminate weakest candidate and “transfer” their votes until there is a winner

Plurality

- Well known that plurality may encourage strategic voting
  - You might want A>B>C but as A has no hope, you vote B>C>A
- And easy to work out manipulation
  - Assuming you know other votes!
  - Consider uncertainty shortly

STV

- Manipulable
  - Satisfies conditions of GS
- NP-Hard to manipulate
  - But proof requires large number of candidates

[Bartholdi, Orlin 91]

Manipulation

- Small domain
  - Only polynomial number of possible votes
  - Can try them all in polynomial time
- Large domain
  - May not turn in social choice but does turn up in AI!

Weighted votes

- Equivalent to coalition voting same way
- Can be NP-hard to manipulate
  - Even with small domain
- Weighted votes used in practice
  - Shareholder meetings, elected assemblies, ...

[Conitzer, Sandholm AAAI02]
[Conitzer, Lang, Sandholm TARK03]

Weighted votes

- Weighted case informs uncertain case
  - Thm: if NP-hard to manipulate with weighted votes then NP-hard with unweighted but uncertain votes
  - Weights like probabilities ..

[Contizer, Sandholm AAAI02]

STV & few candidates

- NP-hard to manipulate STV with weighted votes
  - With as few as 3 candidates

[Conitzer, Sandholm AAAI02]

- Manipulation now by a coalition
  - [Bartholdi, Tovey, Trick 89] considered just one “strategic” voter

Cup rule

- Knockout tournament
  - Sequence of majority comparisons
- Agenda of matches
  - Who plays against who?

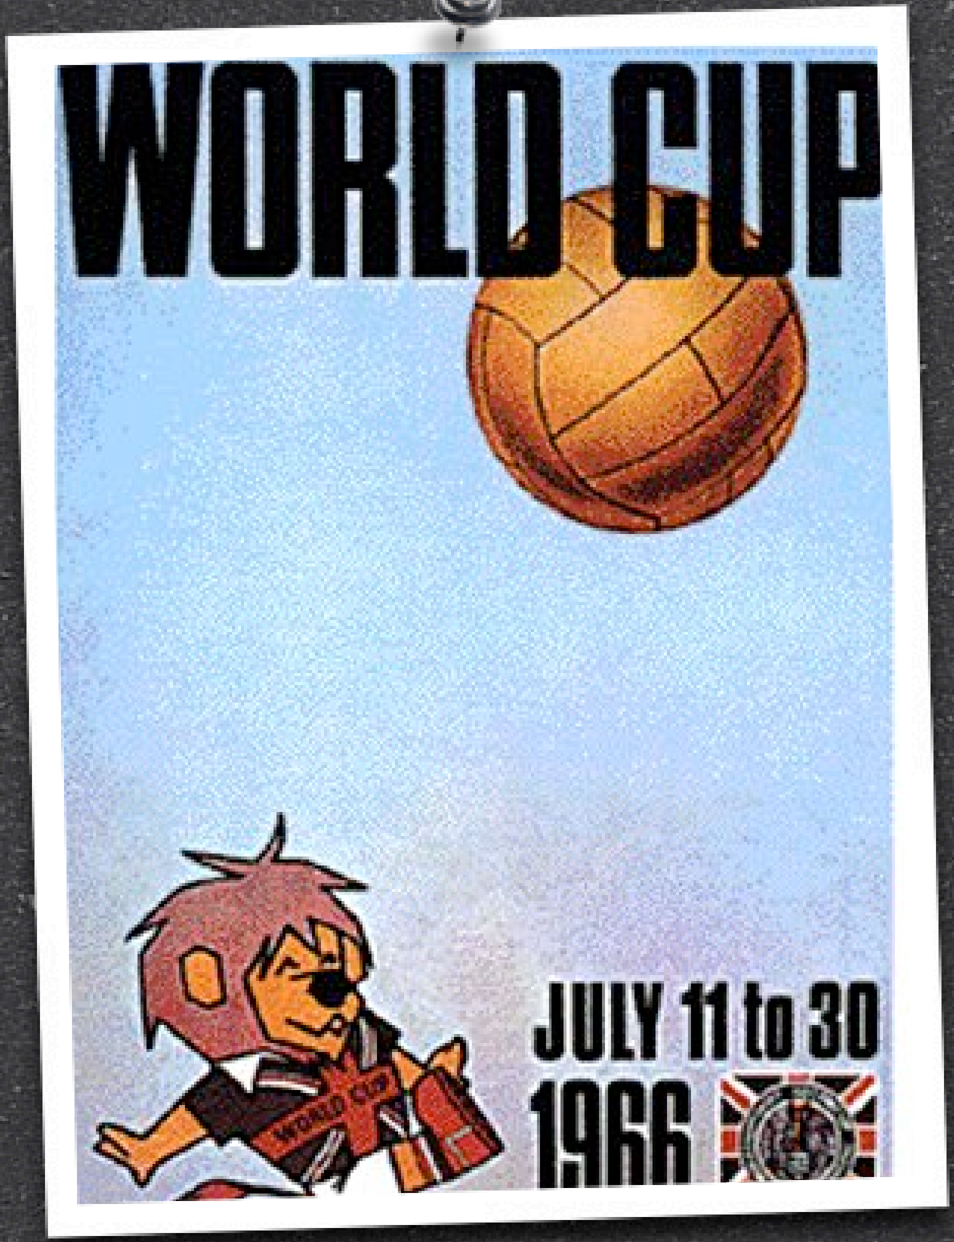

Cup rule

- Fixed agenda
  - Easy to manipulate
- Random agenda
  - NP-hard to manipulate with 7 or more candidates

[Conitzer, Sandholm AAAI02]

Cup rule

- Uncertain agenda
  - Chair tries to manipulate result by choosing agenda
  - Unbalanced tournament: polynomial to manipulate
  - Balanced tournament: open (NP-hard from weighted majority graphs)

[Lang, Pini, Rossi, Venable, self IJCAI07]

Manipulation

- Who is manipulating result?
  - One strategic voter (but one voter can rarely change result!)
  - Coalition of voters
  - Chair (via agenda)
  - ..

Manipulation

- Suppose we can only manipulate certain individual preferences?
  - Bribery
  - Campaigning
  - ...

"You can persuade me to vote for Kerry in front of Gore,
but I'll only ever put Bush last on my ballot!"

Cup rule

- Fixed agenda
  - Manipulating by coalition of votes is polynomial [Conitzer, Sandholm AAAI02]
  - Manipulating of individual preferences is NP-hard
    - 3 or more candidates, weighted votes [Self, unpublished 07]

Elicitation

- Can we declare winner?
  - If we can no longer manipulate election, elicitation can be terminated
  - Manipulation is NP-hard implies terminating elicitation is NP-hard

[Konczak, Lang, 05]

Elicitation

- Can we declare winner of Cup rule?
  - Polynomial if we elicit whole votes
  - NP-hard if we elicit individual preferences

Elicitation

- Motivates elicitation strategy
  - For Cup rule, collect whole votes not individual preferences!
  - Don't ask each voter: "Do you prefer Bush to Gore?"
  - Do ask each voter: "What is your complete preference ranking?"

Other manipulations

- Adding/deleting candidates
  - Partitioning candidates
- Adding/deleting voters
  - Changing agenda
- Bribery
  - Given a particular pot of money

Other manipulations

- Constructive manipulation
  - Ensuring a particular candidate wins
  - Typically P or NP-hard
- Destructive manipulation
  - Ensuring a particular candidate doesn't win

Other manipulations

- Destructive manipulation
  - Ensuring a particular candidate doesn’t win
  - Typically P or coNP-hard
  - Can be easier than constructive manipulation
  - Veto is NP-hard to manipulate constructively but P destructively

Incomplete votes

- Possible winners
  - Can win in some (transitive) completion
- Necessary winner
  - Must win in any (transitive) completion

[Konczak & Lang 05]

Possible & necessary winners

- Closely related to manipulation
  - A $\in$ possible winners $\equiv$ constructive manipulation for A
  - A $\neq$ necessary winner $\equiv$ destructive manipulation for A

Hybrid rules

- Can hybridize voting rules to make them hard to manipulate
  - Plurality easy to manipulation
  - But begin with one round of Cup then it is NP-hard

[Conitzer & Sandholm IJCAI 03]

Hybrid rules

- General method to hybridize voting rules
  - Run k steps of 1st, then execute 2nd rule on remaining candidates
  - E.g. k rounds of Cup then plurality
  - Hybrid is often NP-hard to manipulate

[Elkind & Lipmaa 05]

Only worst case?

- All worst-case complexity results
  - Manipulation/ termination/... might be easy for preferences met in practice?
- Consider single peaked preferences

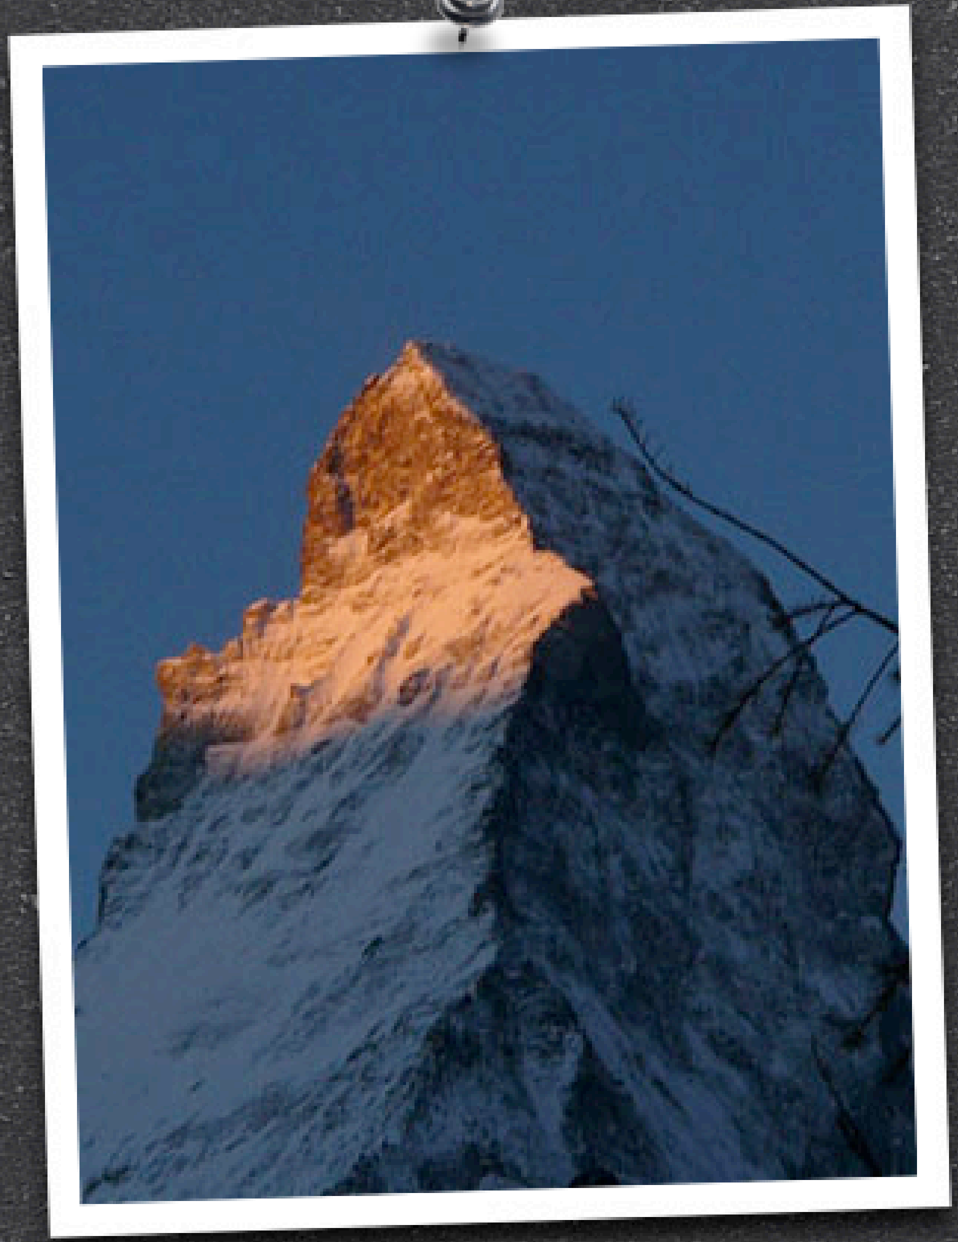

Single peaked preferences

- Occur in practice (e.g. price)
- Defeat Arrow’s theorem
  - Voting rules can be fair!
- Manipulation results often continue to hold
  - STV is NP-hard to manipulate with 3 or more candidates

[Self AAAI07]

Hard on average?

- Several negative results
- Scoring rules and general "junta" distributions
  - On average, likely to find destructive manipulation in polynomial time
  - Applies to "uniform" distribution

[Procaccia & Rosenschein JAIR 07]

Hard on average?

- Any weakly monotone voting rule
  - If manipulator can make either of exactly 2 candidates win
  - Then manipulation can be found in polynomial time

[Conitzer & Sandholm AAAI06]

My impression

- Single round rules tend to be easy on average
- Multiple round rules (like STV or Cup) may introduce difficult balancing problem
    - Good enough to get through to final but bad enough to win
    - May therefore be hard on average?

Conclusions

- Representing and reasoning preferences is an active area of research in AI
  - Some fresh challenges compared to social choice
  - E.g. large domains, computational complexity, constraints, ...
- Much still to be done!

Questions?